

Oracle
Real-Time Scheduler

Business User Guide

Release 2.3.0.3.0

F11061-01

October 2018

Oracle Real-Time Scheduler Business User Guide

Release 2.3.0.3.0

F11061-01

October 2018

Documentation build: 10.12.2018 14:39:57 [M1_1539373197000]

Copyright © 2018, Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners. Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

Contents

Business User Guide.....	10
Overview.....	10
User Interface Standard Features.....	11
Page Components.....	11
The Application Toolbar.....	12
Home Button.....	12
Menu.....	12
Admin Menu.....	13
Back Button.....	14
History Button.....	15
Forward Button.....	15
About Button.....	15
User Menu.....	16
Menu Item Search.....	16
Help Button.....	17
Page Title Area.....	17
Page Actions Toolbar.....	18
Object Display Area.....	22
Multiple Tabs.....	22
Text Box.....	23
Date Only Field.....	24
Date and Time Field.....	25
Search Button.....	26
Context Menu Button.....	26
Go To Button.....	27
Broadcast Button.....	27
Embedded Help Icon.....	27
Work Lists.....	28
Scroll Bars with Dynamic Height.....	28
Sorting Functionality.....	29
Get More Functionality.....	29
Adding and Removing Rows Using Arrows.....	30
Reorganizing Items Using the Up and Down Arrows.....	30
Trees.....	31
Required Field Indicator.....	32
Accordion Windows.....	32
Dashboard Area.....	33
Script Area.....	33
Portal Pages.....	34
Portals Are Made Up Of Zones.....	34
Users Must Be Granted Security Access To Each Zone.....	35
Zones May Appear Collapsed When A Page Opens.....	35
Some Portals Allow User Customization.....	36
Dashboard Portal.....	36
Bookmarks Zone.....	37
Current To Do Zone.....	37
Favorite Links Zone.....	38
Favorite Scripts Zone.....	39
Favorite Searches Zone.....	39
Switch Language Zone.....	39
To Do Summary Zone.....	39
Work List Zone.....	39
Zone Types and Characteristics.....	40
Common Characteristics of Zones.....	40
Explorer Zones.....	40
Map Zones.....	46
Timeline Zone.....	47
Common Base Portals and Zones.....	49

All-In-One Portal.....	49
Separate Query and Maintenance Portals.....	49
Process Flow.....	51
Fixed Pages.....	51
Maintenance Pages.....	51
The Unique Identifier Area.....	51
Confirmation & Error Messages.....	51
Warning If Unsaved Changes Exist.....	52
Query Pages.....	52
The Search Area.....	52
User Preferences.....	52
User - Main.....	53
User - To Do Roles.....	54
User - Access Security.....	54
User - Bookmarks.....	55
User - Portal Preferences.....	55
User - Favorite Links.....	56
User - Favorite Scripts.....	57
User - Characteristics.....	57
User - Miscellaneous.....	57
Initiating Scripts.....	58
Shortcut Key Summary.....	59
Activities.....	62
About Activities.....	62
Activity Type Controls Everything.....	63
An Activity Is Associated With One Host System.....	63
Activity Workflow.....	63
Activity Identifiers.....	68
Location.....	69
Service Area.....	70
Service Class.....	70
Estimated Work Duration.....	70
Time Windows.....	71
Appointments.....	72
Emergency Activities.....	74
Activity Priority.....	75
Resource Capabilities.....	76
Resource Restrictions.....	77
Scheduling An Activity.....	79
Expired Activities.....	79
Activity Actions.....	80
Dispatching An Activity.....	83
Completion Information.....	84
Activity Cancellation By Host.....	84
Depot Related Activities.....	85
Handling Keys.....	85
About Complex Activities.....	86
A Complex Activity Is Scheduled Differently.....	86
Complex Activity Workflow.....	89
Visit Task Workflow.....	91
Use Case Example.....	92
Same Crew.....	94
Under Scheduled Alert.....	94
Time Tracking.....	94
Completion Information.....	94
About Activity Dependency.....	95
Job.....	95
Chain.....	96
About Assignments.....	98
Current Assignment.....	98
Assignment Workflow.....	98
Crew Assignment.....	103
Recall.....	106
Concurrent Assignments.....	106

About Activity Attachments.....	107
Attachments Related To The Type Of Work.....	107
Attachments Related To A Specific Activity.....	107
Maintaining Activities.....	108
Request An Assisting Crew.....	109
Maintaining Current Crew Assignment.....	109
Completing Activities with Concurrent Assignments.....	109
Maintaining Complex Activities.....	110
Maintaining Assignments.....	110
Activity Dependency.....	110
Override Time Windows.....	111
About Override Time Windows.....	111
Prohibit Work When Otherwise Allowed.....	111
Allow Work When Otherwise Not Allowed.....	111
Overrides Regenerate Arrival Time Windows.....	111
Enforcing Overrides.....	112
Use Case Examples.....	112
Maintaining Override Time Windows.....	113
Crews.....	113
About Crews.....	113
Work Is Scheduled To Crew Shifts.....	113
Mobile Workers.....	114
Vehicles.....	115
Fixed Crew Resources.....	115
Resource Capabilities May Change.....	116
Crew Locations.....	116
Planned Time Off.....	116
Crew Hierarchy.....	116
Depot Restrictions.....	116
Crew May Be Managed Externally.....	117
Maintaining Mobile Workers.....	117
Maintaining Vehicles.....	117
Maintaining Crews.....	118
My Calendar.....	118
Resource Route Replay.....	118
Crew Shifts.....	119
About Crew Shifts.....	119
Crew Shift Type Controls Everything.....	119
Creating Shifts.....	119
Default Crew Shift Details.....	119
Crew Shift Workflow.....	119
Controlling The Schedule.....	120
Crew Shift Actions.....	125
Planned Unavailability.....	126
Breaks.....	126
Non-Productive Tasks (NPTs).....	129
Meetings.....	130
Locking The Schedule.....	131
Working The Schedule.....	132
About Shift Weekly Templates.....	135
A Weekly Template Is Made Of Shift Templates.....	135
Crew Specific Weekly Template.....	136
Common Weekly Templates.....	137
One Effective Shift Weekly Template For A Crew.....	138
Shift Generation.....	138
Overridden Shift Details.....	138
Draft and Original Templates.....	139
Maintaining Crew Shifts.....	139
Recommended Crew Allocation.....	140
Maintaining Breaks / Non Productive Tasks.....	140
Maintaining Shift Weekly Templates.....	140
Comparing Draft and Original Templates.....	141
Navigating Between Original and Draft Templates.....	141
Maintaining Shift Templates.....	141

Periods of Unavailability (POUs).....	142
About Periods of Unavailability (POUs).....	142
Different Types of Events.....	142
Leave POUs.....	142
Crew Meeting POUs.....	142
Creating POUs.....	144
About POU Templates.....	144
Actual vs. Template POUs.....	144
POU Generation.....	145
Overridden POU Details.....	146
Maintaining Periods Of Unavailability.....	146
Maintaining Period Of Unavailability Tasks.....	146
Dispatchers.....	146
About Dispatchers.....	147
Managing Crews And Activities.....	147
Dispatcher Shifts.....	148
Handling Alerts.....	148
KPIs.....	149
Managing Depots.....	149
To Dos In Scope.....	149
About CDI.....	149
CDI is for Dispatchers Only.....	150
CDI Data.....	150
CDI Map.....	150
CDI Is Customizable.....	151
Maintaining Dispatchers.....	151
Maintaining Dispatcher Shifts.....	152
CDI.....	152
Managing your Dispatcher Shift.....	153
Using the Scheduling Gantt.....	153
Gantt Views.....	153
Crews Pane.....	154
The Schedule Area Pane.....	154
The Toolbar.....	156
Customize Your Gantt Display.....	156
Filter The List of Crews.....	157
Taking Actions.....	157
Plan View.....	158
Tasks View.....	159
Depot View.....	159
Bay View.....	159
Using the Activity Search.....	160
Using the CDI Map.....	160
Monitoring KPIs.....	161
Monitoring Contracted Work.....	161
My Alert Queue.....	161
Managing To Do Entries.....	161
KPI Specific CDI.....	162
Dispatcher To Do List.....	162
Alerts.....	162
About Alerts.....	162
Alert Type Controls Everything.....	163
Alert Queue.....	163
Backup Alert Queue.....	164
Alert Actions.....	164
About Alert Areas.....	165
Hazard Areas.....	165
Vehicle Leave Areas.....	165
Excluded Resources.....	165
Maintaining Alerts.....	165
Maintaining Alert Areas.....	166
Procedures.....	166
About Procedures.....	166
Procedure Type Controls Everything.....	167

Shift Related Procedures.....	167
Activity Related Procedures.....	167
Maintaining Procedures.....	168
Timesheets.....	168
About Timesheets.....	168
Timesheet Period.....	169
Time Usage Summary.....	169
Generating Timesheets.....	169
Review and Corrections.....	170
Maintaining Timesheets.....	170
Mail.....	171
About Mail.....	171
Emergency Messages.....	171
Mail Summary.....	171
Maintaining Mail.....	171
Transfer of Goods.....	172
About Transfer of Goods.....	172
The Goods.....	172
The Transfer.....	173
Load Capacity.....	173
Vehicle Maximum Capacity.....	173
Load Compatibility.....	174
About Depots.....	174
Location.....	174
Depot Time Windows.....	174
Depot Capacity.....	175
Depot Site Delay.....	175
Depot Task.....	175
Depot Profile.....	176
Depot Time Window Overrides.....	176
Crews Can be Restricted to Specific Depots.....	176
Depot Breaks.....	176
Bays.....	176
Bay Usage.....	177
A Bay Represents a Single Queue.....	177
Handling Various Categories of Load.....	177
Bay Scheduling Restrictions.....	177
Arrival Delay.....	177
Clearance Delay.....	178
Scheduling a Task to a Bay.....	178
Bay Scheduling Restrictions.....	179
Depot Replenishment.....	179
About Depot Runs.....	179
Depot Task Duration.....	180
Depot Task Workflow.....	180
Depot Task Actions.....	181
About Depot Cutoff.....	182
Processing A Cutoff Event.....	182
Initiating Cutoff Events.....	183
Scheduling Manifest.....	183
Close Entire Shift.....	183
About Depot Transfer Activities.....	183
A Depot Transfer Is a Job.....	184
Scheduled to Bays, Not to Resources.....	184
Depot Transfers Are Managed Externally.....	184
Total Break Time.....	185
Estimated Arrival Time to the Unloading Depot.....	185
Crew Location.....	185
Transfer Shift.....	186
Transfer Task Duration.....	186
Transfer Job Workflow.....	186
The Dispatch Process.....	187
Transfer Activity Workflow.....	188
Transfer Activity Actions.....	191

Transfer Scheduling Manifest.....	191
About Depot Transfer Assignments.....	191
Loading Assignment Workflow.....	192
Unloading Assignment Workflow.....	193
Maintaining Depot Tasks.....	195
Maintaining Depot Time Windows.....	195
Maintaining Depot Cutoffs.....	195
Contractor Management.....	196
About Contractors.....	196
Shift Based Model.....	196
Capacity Model.....	196
Activity Contracting Restrictions.....	199
Contractor Eligibility Rules.....	199
Committing An Activity To A Contractor Company.....	200
Resource Groups.....	200
CDI and Contractors.....	200
About Capacities.....	200
Capacity Type Controls Everything.....	201
Creating Capacities.....	201
Capacity Workflow.....	201
Controlling The Schedule.....	202
Capacity Actions.....	202
About Capacity Weekly Templates.....	203
A Weekly Template Is Made Of Capacity Templates.....	203
Actual vs. Template Capacities.....	203
Defining and Applying A Template.....	204
Making Changes To An Applied Template.....	204
One Effective Weekly Template For A Contractor.....	204
Capacity Generation.....	204
Overridden Capacity Details.....	205
Draft and Original Templates.....	205
Maintaining Resource Groups.....	206
Maintaining Capacities.....	206
Maintaining Capacity Weekly Templates.....	206
Maintaining Capacity Templates.....	207
Maintaining Contractor Eligibility.....	207
Contractor Work Management.....	207
Managing Your Work.....	208
Mobile Application.....	208
About The Crew's Application.....	208
Overview.....	208
Preview.....	209
The Mobile Device Is Known At Start Time.....	209
Real-time Updates.....	209
Declining Work.....	210
Sending Current Location.....	210
Geographic Map.....	210
Enforce Procedures.....	210
Handling Attachments.....	211
Creating Activities in the Field.....	211
Switch To Offline Mode.....	211
Request For Help.....	212
Sending / Receiving Mail Messages.....	212
The End Of Shift Process.....	212
Reviewing Timesheets.....	212
Switching Devices Mid Shift.....	213
About The Contractor's Application.....	213
Overview.....	213
No Dispatch Process.....	213
Completion Only Updates.....	214
About Remote Messages.....	214
Messages To A Mobile Device.....	214
Messages From A Mobile Device.....	214
Maintaining Remote Messages.....	214

To Do Processing.....	215
The Big Picture of To Do Processing.....	215
To Do Role.....	215
To Do Priority.....	215
Work Distribution.....	215
A User's Current To Do.....	216
Related To Do Entries.....	216
Complete A To Do.....	217
Managing To Do Entries.....	217
To Do Supervisor Functions.....	217
Supervisor To Do Summary.....	217
Supervisor User Summary.....	218
Supervisor To Do Assignment.....	218
To Do List Functions.....	219
To Do Summary.....	219
To Do List.....	220
To Do Entry.....	222
To Do Entry - Add Dialog.....	222
To Do Entry - Main.....	223
To Do Entry - Keys/Values.....	224
To Do Entry Actions.....	224
To Do Search.....	226
To Do Search - Main.....	226
To Do Search Actions.....	228
How To.....	229
How To Assign a To Do to a User.....	229
How To Work on To Do Entries.....	229
How To Add a New To Do.....	230
How To Complete To Dos.....	230
How To Forward a To Do.....	230
How To Send Back a To Do.....	230
How To Add a Log Entry to a To Do.....	231
Reports.....	231
On-line Report Submission.....	231
Report Submission - Main.....	231
Viewing Historic Reports.....	231
Report History - Main.....	232

Chapter 1

Business User Guide

Overview

OracleReal-Time Scheduler provides enterprise software designed to receive, schedule, assign, dispatch, and track work activities performed by your organization's resources and contractors.

The following lists the functions of each of the application's core components.

- Receive requests to perform activities at specific locations from a host application (i.e., other Oracle Utilities application, Customer Relationship Management [CRM], Enterprise Asset Management [EAM], Transportation Management System [TMS], etc.). Update the requesting host application on the status of work activities and return information collected about the work upon completion.
- Provide the ability for the host application to schedule customer appointments based on resource availability and costs.
- Schedule work activities to the available resources based on the type and the location of the work, and the required skills and equipment needed to complete the work. Manage third party contractors, either resources acting like internal resources or based on work capacity.
- Dispatch work activities to the available resources, including necessary customer / location information, work instructions, and driving instructions.
- Provide the administration function for service management to define their work, work calendars, and work resources, including skill sets, certifications, work schedules, and vacation schedules. Service management defines the scheduling rules, including cost parameters for overtime, and meeting appointments and other service level agreements. Service management also administers the users of the application and their roles.
- Provide the capability for service management and dispatchers to track work activities and manage exceptions.
- Provide a mobile application for the resources to use on mobile devices to work and complete the activities, receiving information about the work to be done (including driving directions), and collecting necessary information from the work location.

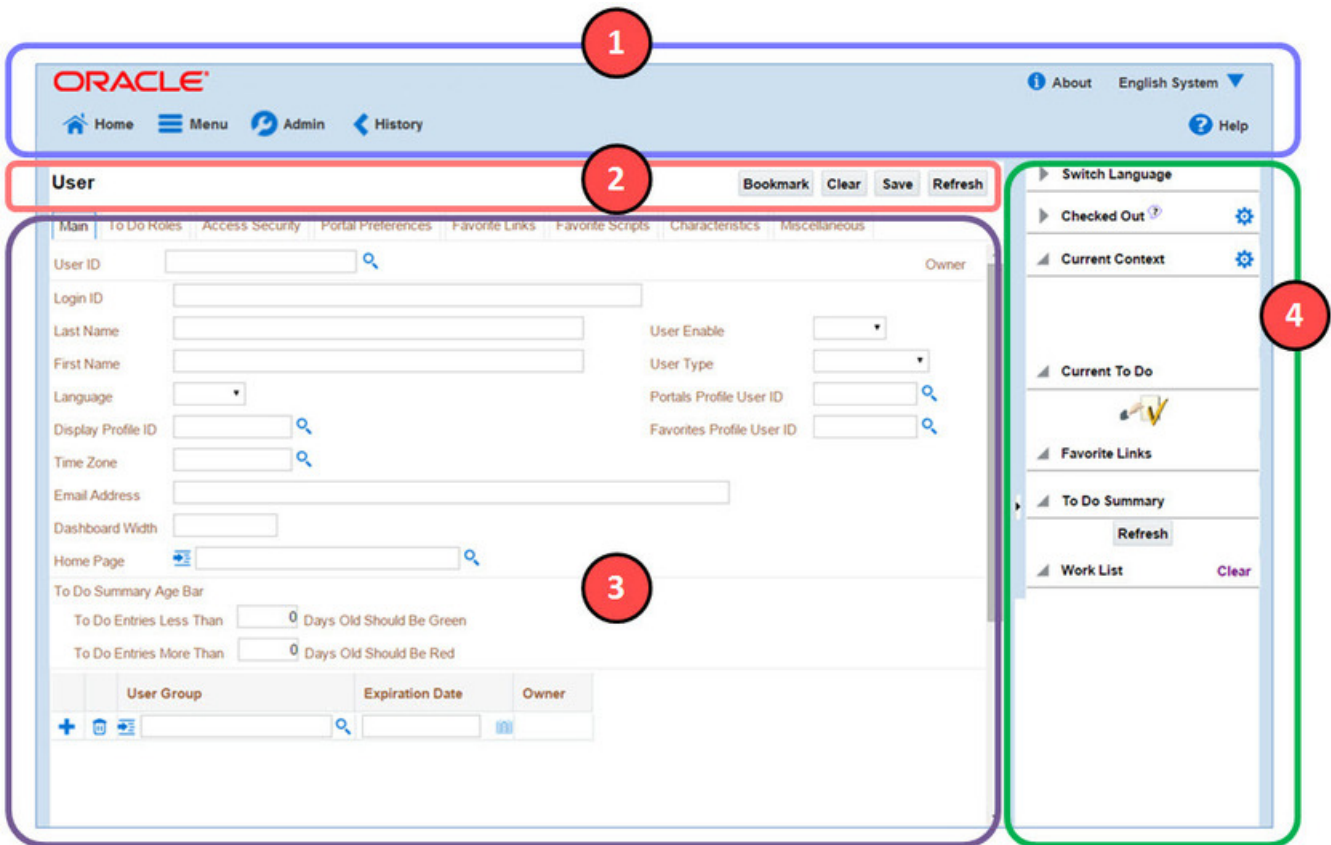
User Interface Standard Features

This section describes basic system concepts, features, and standards of the user interface.

Page Components

Oracle Utilities Application Framework screens are comprised of the following main areas:

1. The Application Toolbar
2. The Page Title Area
3. The Object Display Area
4. The Dashboard Area




NOTE: The look and feel of the application may be modified after the product is installed. See [Custom Look and Feel Options](#) for customization information, including how to change colors, fonts, and other system features. The information provided in this document represents features and functionality available only in the delivered product.

The next topics provide more information about each component and the Script area, which appears in the Object Display area when applicable.

The Application Toolbar

This section describes the features available on the application toolbar.

Home Button

Click the **Home** button,  **Home**, on the Application Toolbar to open your home page. Your home page is defined in [User Preferences](#).

Keyboard Shortcut	Alternate Shortcut
Alt+O	Alt+Shift+O


NOTE: Refer to [Shortcut Key Summary](#) for information about the alternate shortcut.

Menu

Menu is available in the application toolbar to help you navigate to the different pages of the system

Keyboard Shortcut
Ctrl+Alt+M

The menu list is organized by functional area. Clicking the Menu button displays each functional area. Clicking a functional area, in turn, displays a submenu that contains pages within that area.



The pages within each functional area *typically* have two options: Search and Add. If the menu item does not have an Add or Search option, select the menu item itself.

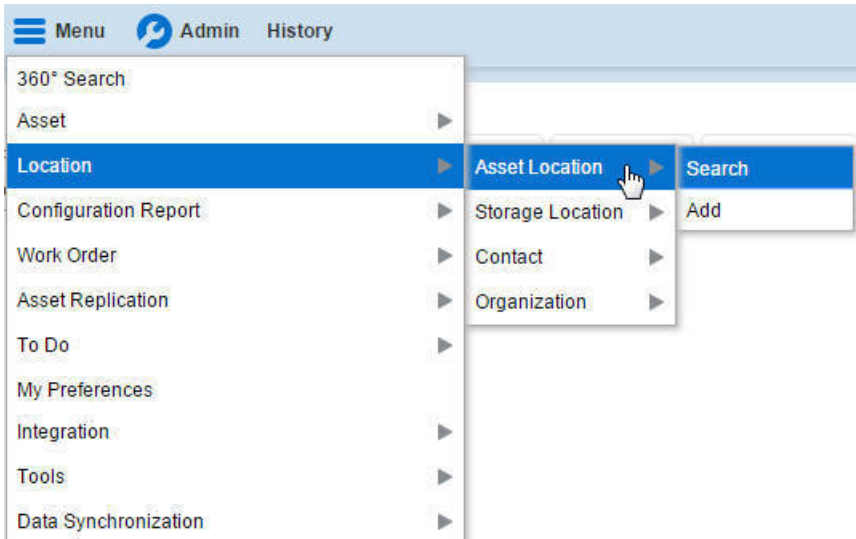
How the Search and Add option behaves depends upon whether the maintenance page is fixed or portal-based.

Fixed:

- **Search:** Displays a pop-up search window for the user to enter the search criteria and select the entity to display. Once the entity is selected, the user is taken to the maintenance map with the data populated in the fields.
- **Add:** Displays the maintenance page with empty fields so that the user can complete the information and save the entity.

Portal-Based:

- **Search:** Displays a search portal where the user can enter the search criteria and select the desired entity. Once the entity is selected, the user is typically taken to a maintenance portal or the information is broadcast to other zones within the same portal.
- **Add:** Either navigates to a page where the user can select the entity type or business object, or directly to the input map where the user can enter and save the entity.



Users may also opt to use the [Menu Item Search](#) input field.

Admin Menu


The Admin dropdown is available in the application toolbar to help you navigate to the different pages of the system.

Keyboard Shortcut

Ctrl+Alt+A

Depending on how your implementation has configured the Admin menu list, it may be organized either by functional area or alphabetically.

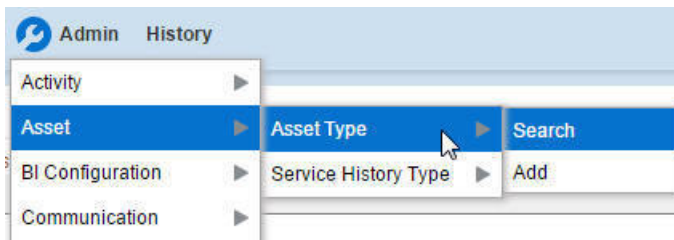
Clicking the Admin button displays each functional area or alphabetical list. Clicking one of the options, in turn, displays a submenu that contains pages within that area. The following is an example of the Admin menu organized functionally.



NOTE: Menu navigation paths referenced in the administrative user guide provide the functional grouping. If your implementation uses the alphabetic grouping, the page can be found in the alphabetic letter that is the first letter of the page name.

NOTE: Refer to [Installation Options](#) for information about the admin menu configuration options.

The menu lines displayed for each menu entry *typically* have two options: Search and Add. If the menu item does not have an Add or Search option, select the menu item itself. See [Menu](#) for more information about these options.



Users may also opt to use the [Menu Item Search](#) input field.

Back Button

Click the **Back** button on the application toolbar to open the previously visited page.

Keyboard Shortcut

Alt+B



The **back** and **forward** navigation buttons, as well as the **History** list, appear in the application toolbar only when page use warrants their appearance.

NOTE: Under certain circumstances, such as revisiting a page that was left unsaved after being loaded with default data, a page may not appear as it did when you left it.

History Button

Click the **History** button to display a list of previously visited pages. When you click on an entry in the list to return to an earlier page, all items above the selected page are removed from the list.



The **back** and **forward** navigation buttons, as well as the **History** list, appear in the application toolbar only when page use warrants their appearance.

Note that if the page has been configured to display an information string in the [page title area](#), that information will also be visible in the History dropdown.



Keyboard Shortcut

Ctrl+Alt+H

Forward Button

Click the **Forward** button to return to the page that was most recently dismissed.

Keyboard Shortcut

Alt+G




The **back** and **forward** navigation buttons, as well as the **History** list, appear in the application toolbar only when page use warrants their appearance.

NOTE: Under certain circumstances, such as revisiting a page that was left unsaved after being loaded with default data, a page may not appear as it did when you left it.

About Button

Click the **about** button to display a window that describes the current and release information for your version of your product.



The about window shows a variety of information including:

- The logged in user ID
- The list of products installed in the application along with the release ID. This information comes from the [installation record](#). Note that in a Cloud installation, only the Cloud product name and release ID are displayed. The details of the products that make up the Cloud installation are not displayed (but are still visible on the installation record).

User Menu

The User menu is available in the application toolbar to provide access to user-specific options.




Through this dropdown, the user can access their [User Preferences](#) or [Logout](#) of their current session.

NOTE: Configurable Menu. Additional menu entries may be visible in your implementation as this menu configuration allows for additional options to be defined.

Preferences

Click the **Preferences** menu item in the User Menu to view or modify your settings in a form that appears in the **Main** tab in the object display area.



Keyboard Shortcut

Ctrl+Alt+P

Logout

Click the **Logout** menu item in the User Menu to complete your current session of the application. Clicking this item logs the current user out and displays the Welcome page so that another user can login.



Keyboard Shortcut

Ctrl+Alt+L

Menu Item Search


The Search Menu input field and search icon are used to allow a user to enter the description of a menu item entry to navigate directly to the corresponding page or BPA script rather than using the menus to navigate to the desired page / script.

Keyboard Shortcut

Ctrl+Alt+F




Typing text in the field shows menu items whose descriptions include the typed text.



Note that only menu items within the [Menu](#) and [Admin Menu](#) that the user has security access for are included in the results. The text included in the search is taken from the menu item description. As described in [Menu](#), menu lines often have two sub item: Search and Add. In most instances, the description of the "Search" menu item is simply the menu entry text, for example **To Do Entry**. The description of the "Add" menu entry includes the word Add in front of the menu entry text, for example **Add To Do Entry**.

Help Button

Click the **Help** button in the Application Toolbar to open context-sensitive Help in a new browser window.



Keyboard shortcut:

Keyboard Shortcut

Alt+F1

Page Title Area

The Page Title Area displays the page title and the page actions toolbar.

For certain pages in the system, the page title may be followed by an information string. It is only applicable to portal-based pages that display a single record. This information is only visible if the portal has been configured to display the primary key's information string. Typically only "stand alone" maintenance portals are configured to show the information string, whereas "all-in-one" portals are not configured this way.

FASTPATH: Refer to [Common Base Portals and Zones](#) for more information about standard base product portals.

NOTE: If an information string is long and would interfere with a page actions button, the information string will automatically be truncated and suffixed by an ellipsis, "...", indicating that more information is available. In this situation, the missing text can be viewed if the user hovers their mouse over the string.

The Page Actions Toolbar displays standard control options, such as Save and Duplicate. More detail about the actions that may appear are described in the following section.

Page Actions Toolbar

In the Page Title Area, control functions appear in the Page Actions Toolbar; such as, Clear, Save, Duplicate, Delete and Refresh buttons.

NOTE:

Buttons in the Page Title Area appear as the function is available to the user; for example, the Delete button will not appear if the user has made any changes to the page's content, or does not have the necessary security rights for this action.

In addition, other functions that are specific to a portal-based page may appear in the toolbar, if configured. For example, the ability to navigate to the Search zone for a specific maintenance object may be visible on the maintenance portal.

This section describes the standard features that are available in the Page Actions Toolbar.

Bookmark Button

Click the Bookmark button on the Page Actions Toolbar to save the current page and context to the bookmarks list.



Keyboard Shortcut

Ctrl+Alt+B


NOTE: Secured. This button is secured using the application service **F1-BOOKMARKS**. Users that do not have security for the application service will not see the Bookmark button (or the corresponding dashboard zone).

After clicking the Bookmark button, you are prompted to provide a name for the bookmark. After saving, it is visible in the [Bookmarks Zone](#), allowing you to return to this page with this context loaded with one click.

NOTE: Unsaved data will not be included in the bookmark. If you view a record and make changes to data without saving the changes and then bookmark the page, the unsaved changes are not included in the bookmark information.

You may change the name of the bookmark or change the sequence of the bookmarks by navigating to the [User - Bookmarks](#) in the user preferences.

Next/Previous Item Buttons



The (**Previous Item** and **Next Item** buttons appear in the Page Actions Toolbar under these conditions:

- You access a fixed page, meaning that it is not portal-based.
- You use a search page to display an item.
- There are multiple items in the search results list *below* or *above* the selected item.

If, for example, you enter search criteria of "Smith" on the User Search page, you might see the following results:

- Smith, Albert

- Smith, Peter
- Smith, Zane

If you select **Smith, Peter**, Peter's user information is displayed on the User page. If you then click the **Next Item** button, **Smith, Zane** will be displayed. In addition, because **Smith, Zane** is the last entry in the search results, the **Next Item** button is no longer needed and will not appear in the Page Actions Toolbar.


If you then click the **Previous Item** button, **Smith, Peter** will be displayed; click again and the **Smith, Albert** record will appear, and, because **Smith, Albert** is the first entry in the search results, the **Next Item** button is no longer needed and will not appear in the Page Actions Toolbar.

Keyboard Shortcut	Alternate Shortcut
Alt+P	
Alt+N	Alt+Shift+N

NOTE: Refer to [Shortcut Key Summary](#) for information about the alternate shortcut.

Duplicate Button

Click the **Duplicate** button in the Page Actions Toolbar to create a duplicate of the currently selected object.



Before the new object is added to the database, a dialog appears in which you must enter the unique identifier for the duplicate object. When you click OK, the system creates the record by copying the original record. It may be helpful to think of the duplication function as the equivalent of a "save as" action.

Keyboard Shortcut

Ctrl+Alt+D

NOTE: Copying ancillary information for a record sometimes requires confirmation during a duplication request. If so, you may also be asked if you also want to duplicate ancillary information. If, for example, you duplicate a user group, besides supplying the ID of the new user group, you are asked if all users in the original user group should be linked to the new user group.

This button does not appear in the Page Actions Toolbar when:

- Duplication is not supported for the current object.
- You are on a portal based page. (For these types of pages, if duplication is supported, the actions are defined in a zone on the portal.)
- You have made changes to the object without committing the changes to the database by clicking Save. (After making changes, you must either discard them using the **Refresh** button, or save them.)

Delete Button

Not all transactions have the capability to be deleted. Deletion is related to the business impact of the action. In some transactions, Delete is present for some scenarios and unavailable for others; for example, you may not be able to delete a record at a certain status level.


NOTE:

If you make a change to an object and do not save the changes to the database, you cannot delete the object until you either save the changes or refresh the object with the contents on the database.


Before you delete an object, the system verifies that the object is not referenced anywhere in the system. If it is, you cannot delete it. If, for example, you attempt to delete a user, the system ensures this user has not been referenced in a user group or on a To Do.

Depending on the type of page that you are using, the option to delete may vary.

- Fixed pages display a Delete button in the page actions toolbar when the option is available.



- Portal-based pages display the Delete button in the object-related map.



- List zones on All in One portals commonly display the Delete button (garbage can icon) next to the entity in the search results.



In all cases, before the object is deleted, a confirmation dialog is displayed.


CAUTION: A timeout situation could occur as the system attempts to verify the use of the deleted object if the object is referenced in many places *and* involves large volumes of production data. If this timeout situation occurs, you must work with your database administration group to delete the offending object.

Clear Button

Click the **Clear** button to remove the values from the currently displayed object's fields. This action is typically used to clear input fields when adding a new object.


Keyboard Shortcut

Ctrl+Alt+C



Save Button

Clicking the **Save** button saves any changes you have made in the data on the current object.



Keyboard Shortcut

Alt+S

Alternate Shortcut

Alt+Shift+S

NOTE: Refer to [Shortcut Key Summary](#) for information about the alternate shortcut.

The button does not appear unless:

- You enter a page in add mode and default values are loaded.
- You make changes to an object after it is displayed (note that you must tab or click out of the field before the Save button appears).


Refresh Button

Click the Refresh button to reload the values on the page with those from the database.

CAUTION: All unsaved changes are discarded.

Keyboard Shortcut

Alt+R



Object Display Area

The Object Display Area contains the details of the object(s) currently being viewed or maintained. For example, if you display a portal, its zones are displayed in this area.

Because many of the page components follow Internet standards of operation, the behavior of these types of components are not explicitly described. This section focuses on special features that are unique to the application framework.

Multiple Tabs

When a page in the application includes a lot of information, the page may be implemented with vertical scrolling to allow the user to view all the information. Or alternatively, the page may be implemented with several tabs, allowing the user to navigate between tabs to view or update information. The following image is a section of the User page showing the multiple tabs.





Figure 1: Multiple Tabs Example

The names of the tabs describe the type of information that will be displayed in the page when the tab is clicked. For example, if you click on the **Favorite Links** tab, the contents of the page show the favorite links configured for the user.







	Sequence	Navigation Option	Security Access
 	10	CI0000000841  To Do Role 	Yes

Figure 2: Multiple Tabs Example (After Favorite Links Tab Is Clicked)

NOTE: Hot key options. Rather than clicking on a tab, you can click F2 or Shift+F2 to scroll through the tabs. F2 moves you through the tabs from left to right. Shift+F2 moves you through the tabs from right to left.

NOTE: Hot key options. Use Alt+ a number to navigate to the tab page whose relative position corresponds with the number. For example, if you press Alt+6 on the above page, the **Favorite Links** tab will be displayed. If there are more than 9 tabs on a page, you must use the mouse or F2 / Shift+F2 to display a tab. Note that the numeric keypad on your keyboard cannot be used to implement this function.

NOTE: Returning to a page that contains multiple tabs. If you return to a page (by clicking the [Back](#) button or the [History](#) button), the page will be positioned on the last tab that was opened.

NOTE: Page versus Tab nomenclature. The documentation frequently refers to the contents of a given tab as a page. However, the page actually comprises the entire collection of tabs.

Text Box

The following is an example of a basic text box:

Description

If the text box is a large field, the text box may be configured to enable scrolling. In this situation, an expanded edit icon, (a pencil), will appear adjacent to it.

Detailed Description

This zone type is used to create an info zone. Info zones contain a list of information about a given object. For example, you could set up an info zone to produce a list of all contacts made by the customer in context.


Info zones offer many features that can be used to change how the list looks "on the fly" (by end users). For example, users can add additional columns to the list, change the sort order, ...

For a more detailed description of how these zones look and behave, open the help and navigate to the index entry titled "info zone configuration".

To set up a new zone, fill in the zone's parameters as per each parameter's Comments. For examples, please refer to the "tips" that appear in the dashboard.


Additional capabilities are provided to display and edit data in these text boxes:

- If the field contains more text than is visible, a pop-up that displays the full text appears when you hover your mouse over the text box. The text in the pop-up can be selected and copied.



The screenshot shows a browser window with a text box titled "Zone Type" and value "F1-DE-SINGLE". The text box contains several paragraphs of text. A mouse cursor is hovering over the text box, and a pop-up window has appeared, displaying the full text of the text box. The pop-up window has a title bar "Detailed Description" and contains the same text as the text box. The text box also has a search icon and a scroll bar on the right side.

- If you click the expanded edit icon, an editable pop-up window appears. If the field is not protected as a result of business rules for the page, you may modify the text in this window.




NOTE: The following table highlights special characters that are converted to equivalent standard characters when entered in a text box.

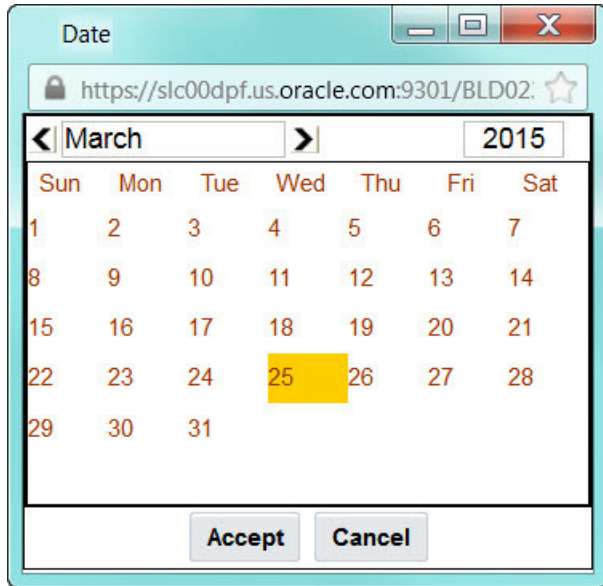
Entered Character Description	Unicode	Converted Character Description
Curly left double quote.	U+201D	Straight double quote.
Curly right double quote.	U+201C	Straight double quote.
Curly left single quote.	U+2019	Straight single quote.
Curly right single quote.	U+2018	Straight single quote.
En dash.	U+2013	Simple hyphen.
Em dash.	U+2014	Simple hyphen.
Ellipse	U+2026	Three dots.
Middle dot (bullet / centered dot)	U+00B7	Asterisk.
Fraction 1/4 (using 1 character).	U+00BC	1/4 (using three separate characters).
Fraction 1/2 (using 1 character).	U+00BD	1/2 (using three separate characters).
Fraction 3/4 (using 1 character).	U+00BE	3/4 (using three separate characters).

Date Only Field

When only the date is required, the Date field looks as follows:



If you click the calendar icon next to the date field, or you double-click inside the date field, the following window appears:



If the date field was blank, the current date is highlighted when the window first appears. If the date field contained a date, it is displayed.

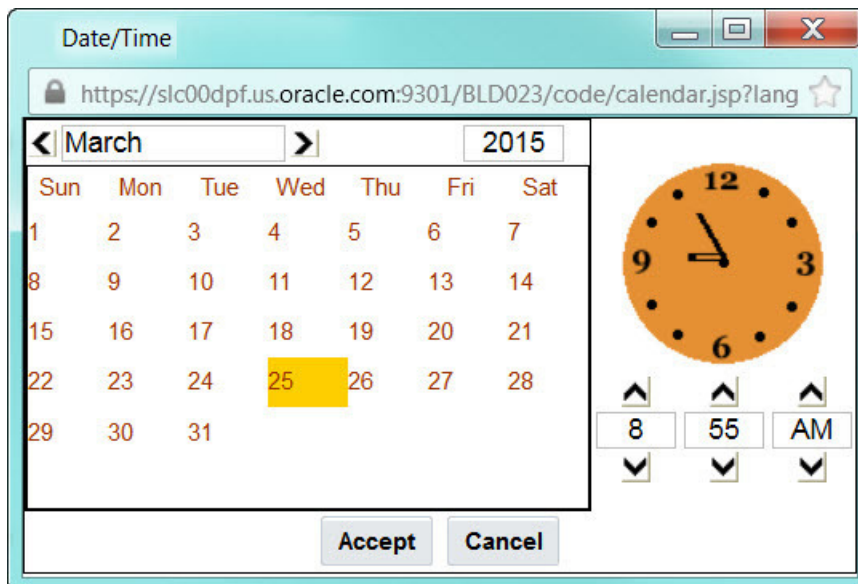
You can double-click the desired date, or you can select a date and click Accept to close the date window and populate the date field.

Date and Time Field

When both the date and time are required, the Date and Time fields typically look as follows:



If you double-click or press the Enter key while the cursor is in a Date and Time field, the following window appears:




Select the date and use the toggle buttons (or enter) the time and then click **Accept** to populate the original fields.


NOTE: The user cannot select and move the hands on the clock to set the time.

Search Button

The search button appears primarily as a magnifying glass.



When using a pop-up search window, the search button can appear as a button with its function labeled.



You can have the system search for an item that matches the data you entered using any of the following methods:

- Enter a value in the search field and press **Enter** while the cursor is in the search field. The system assumes you want to search for the value entered and pre-populates it in the search window's search criteria. If a unique match can be found, the system automatically populates the field with details corresponding to the record found. If a unique match cannot be found (either no matches or multiple matches), a search page opens containing selectable items.
- Click the Search button without entering a criteria. This will typically cause a pop-up search window to appear so that you can further refine your search.
- Enter a value and tab out of the search field. The system looks for a value that exactly matches what you entered. If it cannot find such a value, the search dialog will not be presented.


When an item is selected in the search page, the search page closes and the input field is populated with the value selected.

NOTE: Sorting columns within the search grid. You may click on the column heading of columns in a search grid to cause the contents of the grid to be sorted. Refer to [Sorting Functionality](#) for more information.

Context Menu Button

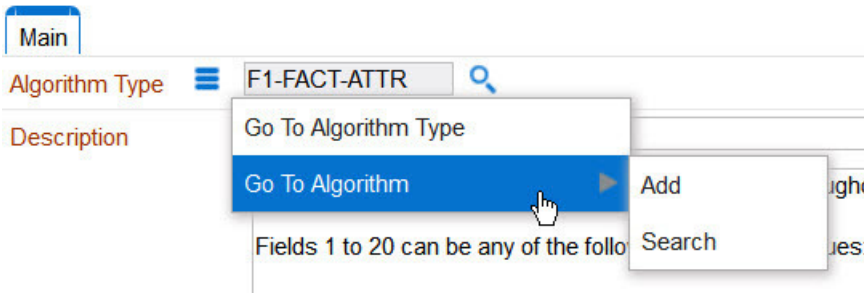
The Context Menu button is a small version of the Menu button from the Application Toolbar. An example is shown below:

Algorithm Type




Context menus exist to help you navigate to other pages and populating the page with the current context (by passing the ID of the related object). In addition, the menu items offered in a Context Menu differ depending on the entity that the menu is related to; for example, when you click the Algorithm Type context menu, a pop-up menu similar to this appears:


Algorithm Type




Different context menus exist for each of the major objects in the system; for example, in Customer Care and Billing, the context menu for Premise contains premise-oriented menu items, whereas the context menu for Account contains account-oriented menu items.

Go To Button

Clicking a Go To button, , opens the page for the related object - it functions like a hyperlink.



Broadcast Button

Some zones have been configured to "broadcast" unique identifiers to other zones within the portal; if so, you'll see a broadcast icon, , in the first column. Clicking the icon causes specific zones to reload - these zones that have been configured to rebuild when the unique identifier is received. For example, you might have configured a zone to display the image of a bill when a bill ID is broadcast from an info zone.


Embedded Help Icon

The icon containing a small 'i' is an Embedded Help icon.



When an embedded help icon is clicked, a tool tip appears containing additional information.

Work Lists


The work list buttons can appear in a zone's column heading, .

Clicking the icon results in populating the dashboard [Work List Zone](#) with the entire column. The user may use the Work List zone to easily process entries one at a time without returning back to the source query zone. The work list stays populated until the user overlays the work list by clicking a different column's icon.


When a user clicks on a work list entry in the dashboard, the system behaves as if the user clicked on the hyperlink in the zone. In addition, the system adds a check mark next to the item showing that the user has "worked on" or visited the entry.

NOTE: The Work List zone is populated only upon user request.

Scroll Bars with Dynamic Height


Sometimes the height of a grid isn't sufficient to display a grid's rows. When this occurs, a vertical scroll bar appears and an "expand button" appears above this bar, . For example, the following grid is not high enough to show all of the available dashboard zones:


Dashboard - 10 of 10 zones selected for display.

Zone	Display	Initially Collapsed	Sequence	Refresh Seconds		Security Access		
Switch Language	<input checked="" type="checkbox"/>	<input type="checkbox"/>		<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="text" value="20"/>	<input type="text" value="0"/>	Yes
Bookmarks	<input checked="" type="checkbox"/>	<input type="checkbox"/>		<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="text" value="0"/>	<input type="text" value="0"/>	Yes
Checked Out	<input checked="" type="checkbox"/>	<input type="checkbox"/>		<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="text" value="0"/>	<input type="text" value="0"/>	Yes
Current Context	<input checked="" type="checkbox"/>	<input type="checkbox"/>		<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="text" value="0"/>	<input type="text" value="0"/>	Yes

If you want to view all of the person's phone numbers without using the vertical scroll bar, you can click the expand button. When clicked, the height of the grid expands to show all rows in the grid:

Dashboard - 10 of 10 zones selected for display.

Zone	Display	Initially Collapsed	Sequence	Refresh Seconds		Security Access		
Switch Language	<input checked="" type="checkbox"/>	<input type="checkbox"/>		<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="text" value="20"/>	<input type="text" value="0"/>	Yes
Bookmarks	<input checked="" type="checkbox"/>	<input type="checkbox"/>		<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="text" value="0"/>	<input type="text" value="0"/>	Yes
Checked Out	<input checked="" type="checkbox"/>	<input type="checkbox"/>		<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="text" value="0"/>	<input type="text" value="0"/>	Yes
Current Context	<input checked="" type="checkbox"/>	<input type="checkbox"/>		<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="text" value="0"/>	<input type="text" value="0"/>	Yes
Current To Do	<input checked="" type="checkbox"/>	<input type="checkbox"/>		<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="text" value="0"/>	<input type="text" value="0"/>	Yes
Favorite Links	<input checked="" type="checkbox"/>	<input type="checkbox"/>		<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="text" value="0"/>	<input type="text" value="0"/>	Yes
Favorite Scripts	<input checked="" type="checkbox"/>	<input type="checkbox"/>		<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="text" value="0"/>	<input type="text" value="0"/>	Yes
Favorite Searches	<input checked="" type="checkbox"/>	<input type="checkbox"/>		<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="text" value="0"/>	<input type="text" value="0"/>	Yes
To Do Summary	<input checked="" type="checkbox"/>	<input type="checkbox"/>		<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="text" value="0"/>	<input type="text" value="0"/>	Yes
Work List	<input checked="" type="checkbox"/>	<input type="checkbox"/>		<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="text" value="0"/>	<input type="text" value="0"/>	Yes


Notice, after the grid is expanded, the grid's vertical scroll bar disappears and the button in the upper corner becomes a Collapse Button, . If you click the collapse button, the grid will return to its original height.

NOTE: Search grids. Expand all functionality is not available in search grids.

Sorting Functionality

Grids have a heading row that contains labels for the columns. By clicking a column heading, you cause the rows to be sorted in the order of the selected column. By clicking on a column more than once, you cause the rows to be sorted in reverse order. Clicking the heading a third time removes the sorting function and returns the grid to its default display organization.

The following shows how information look after organizing the grid, clicking the column heading, by algorithm type (notice the icon in the column name).

Algorithm Code	Description	Algorithm Type 
F1-ACTIVEWS	Web Service - Verify Active Operations	F1-ACTIVEWS
F1-ATCHM-INF	Attachment Information	F1-ATCHM-INF
F1-BAT-LSDEF	Batch Level of Service Algorithm	F1-BAT-LVSVC
F1-BNDL-INFO	Bundle Information	F1-BNDL-INFO
F1-BNEXVALID	Bundle Export Validation	F1-BNEXVALID
F1-BOSTATFS	Take BO Status/Reason Final Snapshot	F1-BOSTATFS
F1-BUNEXPBUN	Prepare bundle XML	F1-BUNEXPBUN
F1-BUNIMPAPP	Apply a bundle	F1-BUNIMPAPP
F1-BUNLOGUSR	Log user updating an export bundle	F1-BUNLOGUSR
F1-CMPSNAPSR	Compare Sync Request Snapshots, populate status reason if discarded.	F1-CMPSNAPSR
F1-COMPSNAPS	Compare Initial and Final Snapshots	F1-COMPSNAPS
F1-DECRT-INF	Decision Report Information	F1-DECRT-INF
F1-FINALSNAP	Take Final Snapshot	F1-FINALSNAP

NOTE: Sorting on dropdown values. If you attempt to sort a column that contains a dropdown value, the information is sorted in the order of the unique identifier associated with the dropdown value, not the description you see on the screen; for example, if you sort on a column containing unit of measure, the information will be sorted in the order of the unit of measure code, not the description that appears in the drop down. This means that the resultant sort will appear incorrect and therefore we discourage you from doing this.

Get More Functionality

Grids that can potentially contain a large number of rows have special processing that we refer to as Get More functionality. We'll use an example from Oracle Utilities Customer Care and Billing to explain - consider the grid that contains all payments that have been added to a cash drawer. By the end of the day, this grid could contain hundreds of rows (one for each payment remitted during the day). Rather than show every payment in the grid, the system retrieves the first X payments (where X varies from grid to grid, and from release to release) and appends a Get More button to the bottom of the grid as illustrated below:

	Tender Amount	Tender Type	Account	Check N
	\$123.90	Cash	4172150879 - Rodehaver, Gary - Residential - M-D1	
	\$69.76	Cash	5983360951 - Nicholson, Edvard Paige - Res - M-D1	
	\$10.00	Cash	4438465447 - Richie, Kate - Residential - - ALOP	
	\$10.00	Cash	2609586025 - Richie, Harry - Residential - - ALOP	
	\$10.00	Cash	8593641982 - Richie, Girlie - Residential - ALOP	
	\$10.00	Cash	0401455188 - Richie, Iris - Residential - - ALOP	


Get More...

When you click the **Get More** button, the system appends the next X rows to the bottom of the grid.


NOTE: The absence of a "get more" button. If no button appears at the bottom of the grid, this means you are seeing all of the rows.

NOTE: Search grids. The Get More function is not available in search grids.

Adding and Removing Rows Using Arrows

In some editable grids, a new row can be added to another editable grid (the grid on the left) by clicking the Add item button, , of the desired row in the grid on the right:

	Sequence	Description		Sequence	Description
	10	Winter kWh charge, first 500		30	Winter subtotal
	20	Winter kWh charge, remaining		60	Summer subtotal
	40	Summer kWh charge, first 300			
	50	Summer kWh charge, remaining			
	70	Employee discount			

When the Remove item button is clicked, , the row is transferred from the grid on the right to the grid on the left. This change will not be committed to the database until you save the object being maintained

Reorganizing Items Using the Up and Down Arrows

Occasionally, rows can be repositioned in an editable grid by clicking up and down arrows. The following is an example of such a grid:

		Merge Type	Sequence	Step Type	Description
		Original	10	Edit data	Edit data - Perform State Transition
		Original	20	Edit data	Edit data - Perform Main Processing
		Original	30	Label	Label - End of script

Notice the following about such a grid:

- The up arrow is dim in the first row. This is because it can't be moved further to the top.
- The down arrow is dim in the last row. This is because it can't be moved further down.
- The up and down arrows are usable in all other rows. Each time you click one of these arrows, the respective row is moved up or down one row.

NOTE: Reminder. When you move rows into a grid, the system does not apply the change to the database until you click the Save button (which commits the entire object to the database).

Trees

Trees are used to illustrate complex relationships between objects. The following is an example from Oracle Utilities Customer Care and Billing of a page with a tree that is expanded:



The following points describe how to use trees:

You can expand and contract the nodes in a tree by clicking on the icon that prefixes the node (you can also click the + and - button to expand and contract a node).

NOTE: Expand all. If you press `Ctrl` while clicking on an icon that prefixes a node, all child nodes (and their children) expand.

You can click the context menu button to cause a context menu to appear. You can navigate to any of the pages that are presented on this menu. Refer to [Context Menu Button](#) for more information.

You can click on the information that appears in an underlined node to drill into the page in which the object is maintained. For example, if you click on the account information, you will be transferred to the Account - Main page.

NOTE: Important! If you make changes after expanding a tree's nodes (e.g., by drilling down on a Premise and changing the premise's address), the tree will contract to the first node if you return to the page.

Nodes in trees fit into two broad categories:

- **Information nodes.** Information nodes highlight the existence of other data. For example, on the Control Central - Bill Payment Tree, a node exists to highlight the existence of new financial transactions that were created since the last bill was completed.

- **Data nodes.** Data nodes show information about objects in the system. For example, data nodes exist on the Control Central - Bill Payment Tree that show bills and payments.

Required Field Indicator

An asterisk (*) is used to indicate that a field is required. The system determines this based on the field's configuration. When a field becomes required because of an algorithm and not the field's configuration, however, the system will not add an asterisk.

NOTE:

If a page includes a grid, fields within a row may be marked as required.




In this situation, the field(s) within the grid are **only** required if the row is being populated.

Accordion Windows

Accordion windows are used when the data held in an object's rows does not comfortably fit into a grid.

Each row in an accordion has a header that summarizes a row's details. When you click a header, the row expands to show the details. The following is an example of an accordion where all headers are collapsed (meaning that only summaries of the rows' details are shown):



When you expand a header, its details are exposed. For example, the following shows the same accordion after the third row is expanded:

The screenshot shows a software interface with an accordion containing three steps:

- Step 10:** Edit data - Get the Physical BO for the MO. This step is expanded, showing:
 - Step Sequence: 10
 - Step Type: Edit data
 - Text: Get the Physical BO for the MO
 - Edit Data Text:


```
move "parm/mo" to "F1-RetrieveMOOption/input/mo";
move 'F1PB' to "F1-RetrieveMOOption/input/moOptionType";
move 'false' to "F1-RetrieveMOOption/input/isMultiple";
invokeBS 'F1-RetrieveMOOption' using "F1-RetrieveMOOption";
```
- Step 20:** Edit data - Get the prime keys of the BO
- Step 30:** Edit data - Read the BO

NOTE: You must expand to update. In order to update a row in an accordion, you must expand it.

Row Summary and Chunk Management

At the top of an accordion you'll find a summary of its rows. The format of this information differs depending on the number of rows and the size of the accordion:

- If the accordion contains every row, a summary of the total number of rows is displayed.
- If the accordion contains a subset of rows, you are presented with a summary of the rows AND you are provided with options to display different "chunks" of rows.
 - Click **First** to display the first "chunk" of rows (the size of a chunk differs depending on the accordion). This option is dimmed if the first chunk is in the accordion.
 - Click **Previous** to display the previous chunk of rows. This option is dimmed if the first chunk is in the accordion.
 - Click **Next** to display the next chunk of rows. This option is dimmed if the last chunk is in the accordion.
 - Click **Last** to display the last chunk of rows. This option is dimmed if the last chunk is in the accordion.

Dashboard Area


The dashboard is a [portal](#) that contains zones. These zones appear regardless of the object being displayed in the [Object Display Area](#). Refer to [Dashboard Portal](#) for a description of these zones.

NOTE: **User configurable.** Every user has control over the zones that appear in their dashboard. Refer to [Portal Preferences](#) for more information about how a user can customize the contents of the dashboard.

Script Area

When a [script is initiated](#), the [object display area](#) shifts down to make room for the script area.

The information in the script area takes you through the steps in the selected business process (provided by your implementation team).



Portal Pages

This section describes using portals and zones to find and display information. This is the preferred technique used by the base product for building user interfaces.

Additionally, portals and zones may be configured by your implementation to supplement the base package provided functionality.

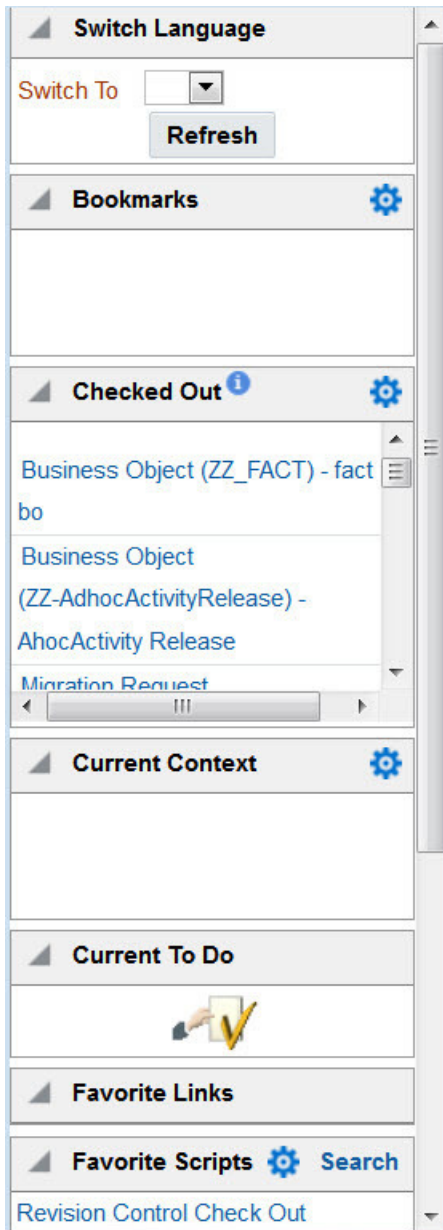
The contents of this section provide background information about portal page behavior.

Portals Are Made Up Of Zones

The term "zone" is used to describe a section on a portal. For example, the following is a portal with seven zones:

- Switch Language
- Bookmarks
- Checked Out
- Current Context
- Current To Do
- Favorite Links
- Favorite Scripts

NOTE: This example may not be available in your application; it is used for illustration purposes only.





Users Must Be Granted Security Access To Each Zone

A user must be security granted access rights in order to see a zone on a portal.

FASTPATH: Refer to [Granting Access to Zones](#) for information .

Zones May Appear Collapsed When A Page Opens

Zones may be designed to appear collapsed (or minimized) when the portal is initially displayed. To view the information in a collapsed zone, the user can press the zone's expansion button, , to either expand or collapse the zone:



A zone's data is only retrieved when it is expanded.

Some Portals Allow User Customization

There are some portals in each base product that are configured to allow user customization. In this case, the portal is visible on the user's portal preferences. For these portals each user has control over the following:

- Which zones appear on their portal
- The order in which the zones appear
- Whether the zones should be "collapsed" (i.e., minimized) when the portal opens.

Users define their preferences by clicking the [User Preferences](#) button. Pressing this button transfers the user to the [User Preferences](#) page on which their [Portal Preferences](#) can be defined.

Portals used for base product query and maintenance pages are typically configured to not allow user customization.

Dashboard Portal

The dashboard is a special portal that always appears on the desktop. Its zones contain tools and data that are useful regardless of the object being displayed.

NOTE:

Minimize the dashboard. You can minimize the dashboard by clicking on the vertical bar on the left side of the dashboard. This causes the dashboard portal to disappear allowing more of the page area to display. The vertical dashboard bar is then displayed on the right side of the page. Clicking the dashboard bar restores the display of the dashboard portal. It also suppresses the refresh of the zones within the dashboard based on updates to the context. Maximizing the dashboard then refreshes the zones with the latest appropriate content.

Minimize the dashboard when launching the system. If desired, a URL parameter may be used to minimize the dashboard when launching the system. The following URL highlights the syntax. **`http://system-server:1234/cis.jsp?minDashboard=true`**

Keyboard Shortcut

Alt+J

The contents of this section describe the zones that are provided by the framework product for this portal. Please note the following additional information about the dashboard portal that may impact how it appears.

- **Additional zones may be available in other products.** The zones described below are available in every base product. Please refer to the product-specific documentation for additional dashboard zones available in your product (open the help and navigate to the index entry labeled **Dashboard**).
- **Context Sensitive Zones may appear.** Some pages may define additional zones that should appear in the dashboard only if that page is displayed. These are referred to as Context Sensitive Zones.
- **Dashboard zones are configurable.** The dashboard portal is configured to allow a user to define preferences for it. Refer to [Portal Preferences](#) for more information.

NOTE: Security. As a general rule, dashboard zones each refer to a unique application service so that a security administrator may configure which users are allowed to have access to which zones. Note that the user portal preferences indicates which zones a user has security access to.

Bookmarks Zone

After clicking the [Bookmark](#) button and saving a bookmark, it is visible in the **Bookmarks** zone. Clicking one of the links in the zone navigates to the page where the bookmark was saved and loads the information related to the record that was displayed when the bookmark was saved.

Click the trash can icon to delete the bookmark.


Note that a user may change the name of the bookmark or change the sequence of the bookmarks by navigating to the [User - Bookmarks](#) in the user preferences.

Current To Do Zone

The **Current To Do** zone is designed to streamline the process of working on To Do entries from assignment to completion. It displays key information about the user's current To Do entry and provides functions that help minimize the number of steps it takes to resolve the entry and potentially other entries related to it.

The following sections provide more detail about the zone.

Next To Do Assignment

Users may have the option to click the  To Do icon to request the next To Do entry they should work on. This functionality is optional and can be configured to follow your organization's business rules in determining the next task for a user. For example, if your organization practices work distribution "on demand", this feature can be configured to find the highest priority entry a user can work on and assign it to the user. Refer to [work distribution](#) for more information on this business practice.

NOTE: Next Assignment Rules Are In A Plug-in. Refer to [Next Assignment Algorithm](#) for more information on how to configure the system to take advantage of this function.

If the system is successful in determining the next entry the user should work on, this entry becomes the user's current To Do and its information is displayed on the zone.

User's Current To Do

The user's current To Do may be set using the Next Assignment algorithm, or by drilling in on a To Do from the To Do List or To Do Search page as described in [A User's Current To Do](#). Once a current To Do has been set, the zone displays details related to the To Do.

The **Message** describes why the entry exists. The message description allows the user to drill into the page on which the entry's respective object is displayed. For example, if the entry is associated with an asset that needs review, the user is

brought to the asset page to view the information and potentially take action. If the [To Do type](#) has been configured to associate a [script](#) with this To Do entry's message number, then instead of being taken to the respective page, the associated script is launched.

Info displays the standard To Do information. An Info icon  appears if there are comments on the To Do and a Description icon  appears if there is a log entry of type **Forwarded**, **Sent Back** or **User Details** on the To Do. The description drills into the [To Do entry](#) page allowing the user to view more information about the entry or make changes to it as needed.

The user may use the **Complete** button in the zone to complete the current entry.

Next / Previous To Do Buttons

The **Previous To Do** and **Next To Do** buttons are enabled on this zone under the following conditions:

- The user selected an entry either on the [To Do List](#) page or the [To Do Search](#) page.
- At least one other entry exists *below* (for Next) or *above* the entry selected (for Previous).

Clicking **Next To Do** or **Previous To Do** will bring the user to the transaction indicated on the To Do Entry's message, where its respective object is displayed. It also makes that To Do the Current To Do displayed in the zone.

Keyboard Shortcut	Comments	Alternate Shortcut
Alt+Y	Previous To Do	
Alt+Z	Next To Do	Alt+Shift+Z

NOTE: Refer to [Shortcut Key Summary](#) for information about the alternate shortcut.

NOTE: When a user drills down on a To Do entry from the To Do List page, the system automatically changes the status of the entry to **Being Worked**. If **Next To Do** or **Previous To Do** is subsequently used to scroll through this list, the system does *not* change the To Do entry's status. To change the status of these entries, the user must navigate to the [To Do List](#). On this page, multiple entries may be selected and their status changed with a single user-interface action.

Related To Do Entries

If the To Do entry has non-complete [related entries](#) a summary **Related To Do** description appears. The description states the total number of related entries, how many are still in **Open**, how many are assigned to current user and how many assigned to others. This hyperlink navigates to the [To Do Search](#) query allowing the user to review the related entries. At this point the user may decide to assign entries that can be positively identified as being caused by the same problem as the current To Do also to himself.

If the user determines that all related To Do entries have been resolved with the appropriate action taken, the **Complete All** button may be used to complete in one step the current entry and all its truly related entries assigned to the current user.

If the related entries are not resolved with one action, the user may opt to use the **Complete** button to complete the current entry alone and work on related entries one at a time. Using the related To Do information hyperlink the user may navigate to the To Do search page and drill into the first related To Do to work on using the message hyperlink. This entry now becomes the user's current To Do and the **Next To Do** and **Previous To Do** action buttons are set to scroll through the related entries list.

Favorite Links Zone

Each line in the **Favorite Links Zone** corresponds with one of the user's favorite links (e.g., transactions). Clicking on a line invokes the transaction. A number appears adjacent to the first nine entries to highlight that these links can be invoked

by pressing an accelerator key (Ctrl + a number). Note that the numeric keypad on your keyboard cannot be used to implement this function.

Users define their favorite links by clicking the [User Preferences](#) button. Clicking this button transfers the user to the [User Preferences](#) page on which their [Favorite Links](#) can be defined.

Favorite Scripts Zone

The Favorite Scripts zone lists the business process assistant (BPA) scripts that you have defined on your [user preferences](#). You can click on a script to [execute](#) it. You can click the **Search** button in the zone to search for other scripts.

NOTE: If you do not see the Favorite Scripts zone in your dashboard, navigate to User Menu and click the **Preferences** menu item to display the [User - Portal Preferences](#) page. Under the Dashboard Portal row, select the **Display** option for the Favorite Scripts zone.

Favorite Searches Zone

If you marked a saved search as a favorite using [Saved Searches](#) functionality, the search is visible in the **Favorite Searches** zone. Clicking one of the links in the zone navigates to the page where the search was saved and populates the filter values and executes the search.

Switch Language Zone

The **Switch Language** zone is used to switch the application's language. Only languages that are supported by your implementation will appear.

To switch the application's language, choose the desired language and press the **Refresh** button.

NOTE: If your implementation uses a single language, you should suppress this zone by [not granting security rights](#) to the zone's application service.

To Do Summary Zone

The **To Do Summary Zone** has several purposes:

- It presents a summary of the [To Do entries](#) that have been assigned to you and are **Being Worked On**.
- The colored "age bars" highlight the age of the To Do entries. The colors red, yellow and green are used to highlight the relative number of old, middle-aged and new To Do entries. The age associated with each of these colors is [user-configurable](#). Note, you can hover the cursor over an "age bar" to see the number of To Do entries.
- If you click on the To Do Type description or on an "age bar", you'll be transferred to the [To Do List](#) transaction where these To Do entries will be displayed. From this page, you can drill down to the specific entries.

NOTE: Time saver! After you drill down on entry from the [To Do List](#) page, if you have the [Current To Do](#) zone configured in the dashboard, you can use the **Previous To Do** and **Next To Do** buttons in the zone to scroll through additional To Do list entries without returning to the [To Do List](#) page.

Work List Zone

The **Work List** zone provides quick access to an entire column of an info or query zone. When a user clicks on a work list entry, the system will do whatever it would have done if they'd clicked on the hyperlink in the zone. In addition, the system

ticks off that they've worked on the entry (this feature is mostly so users know that they've at least looked at something). Refer to [Work Lists](#) to better understand how to configure query columns to support work lists.

Zone Types and Characteristics

The topics in this section describe common zone types, function, and common characteristics.

FASTPATH: Refer to [The Big Picture of Portals and Zones](#) for a description of general portal and zone functionality.

Common Characteristics of Zones

The following topics describe common characteristics of all zone types supported by the base package.

Zone Visibility

A zone may be configured to dynamically be visible or completely hidden depending on the information being displayed. When hidden, not even the header is displayed.

For configuration information, refer to [Zone Visibility Service Script](#).

Control Zone Headers

A zone may reference a UI Map to define and control its header. This feature provides the means to achieve a uniform look and feel for the application as well as add functionality to the zone header beyond being just a title bar to open and close the zone. Refer to [UI Map](#) and to its context sensitive tips zone for more information on **Fragment** UI Maps and how they are used to construct zone headers.

Explorer Zones

The product provide several zone types that are used to find or display a set of information. These are called explorer zones. There are two types of explorer zones:

- **Info zones.** An info zone displays a list of information about an object. These types of zones often display information without requiring any input from the user.
- **Query zones.** A query zone allows users to find objects matching a given search criteria.

Virtually every aspect of an info or query zone is controlled by its zone [configuration](#). However, most share basic common components described in the subsequent sections.

Search Criteria / Filter Area

For query zones, the Search Criteria Area (also called Filter Area) contains the criteria used to find objects. The design of the zone controls the number and type of filters that appear in this area. After search criteria is entered, clicking the **Search** button causes the system to look for records that match the criteria. The results are shown in the Results Area.

For info zones, the zone's configuration indicates whether or not a Filter area is applicable and if so, whether or not it is open or closed by default. If the filter area is configured, the [Show Filter](#) icon is visible in the zone header. Click this icon to expose the Filter Area, if it is not visible and click the icon to hide the Filter Area if desired. Use this area to filter the information that appears in the zone by different values. Click the **Search** button to rebuild the Results Area using the specified filters.

For example, the following shows an info zone listing user groups that do not have access to a given Application Service where a filter area has been configured to allow users to show user groups specific to a given user by ID or name.

The info zone includes configuration to indicate whether the filter area, when exposed, should be shown at the top of the zone (above the results) or at the bottom of the zone (below the results). For query zones, the search criteria is always displayed at the top.

NOTE: Press Enter. Rather than clicking the **Search** button, you can press Enter to trigger the search / info zone refresh.

Description Bar

The Description Bar describes the search criteria that were used to build the Search Results Area.


Results Area

The Results Area displays the appropriate list of objects. For an info zone, this is the list of appropriate objects with any input filter criteria considered. The following is an example of the results for an info zone. For a query zone, this is the list that matches the search criteria defined in the Filter Area. The following is an example of the results set of an Info Zone.

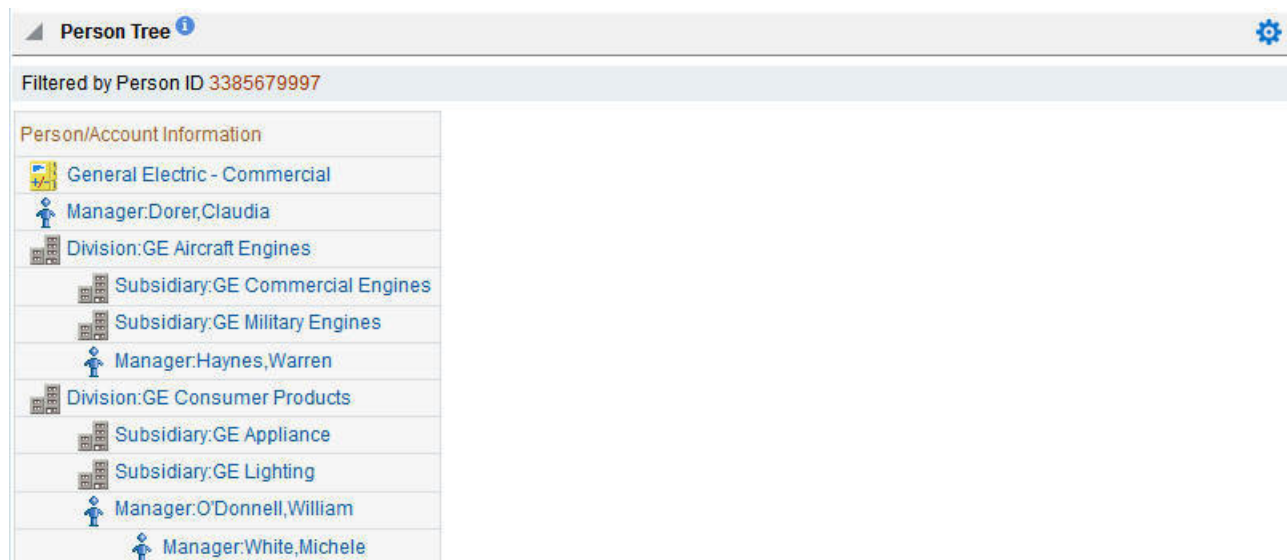
Business Object	Description	Statuses
1 F1-GenericBISyncRequest	Generic BI Sync Request	Discarded, Cancel
2 W1-ActivityEquipRequirement	Equipment for Activity	Fulfilled, Canceled
3 W1-ActivityGenerator	Activity Generator Lifecycle	Canceled
4 W1-ActivityLaborRequirement	Labor for Activity	Fulfilled, Canceled
5 W1-ActivityMaterialRequirement	Material for Activity	Fulfilled, Canceled
6 W1-ActivityOtherRequirement	Other Resource for Activity	Fulfilled, Canceled
7 W1-AdHocActivityCore	Ad Hoc Activity Core	Completed, Discarded
8 W1-ApprovalRequest	Approval Request	Determine Approvers, Approval In Progress, Approved, No Appro
9 W1-BCHeader	Blanket Contract Header	Created, Awaiting Activation, Active, Closed, Suspended
10 W1-BCLine	Blanket Contract Line	Created, Ready for Revision, Revised, Complete
11 W1-CompletionEventRoot	Completion Event Root	Executed, Discarded

For a query zone, the columns that appear may differ depending on the search criteria. For example, if you search for a customer by their driver's license number, the resulting columns may be different than the ones displayed when you search by customer name.

The following points highlight details related to viewing and working with the result set:

- **Broadcast.** If the results set is configured to display a broadcast icon, , it means that at least one other zone on the portal is configured to display data related to the record selected in the results list. Clicking this icon will cause the other zone or zones to be built or refreshed to display related data.
- **Hypertext.** One or more columns in the zone may be configured as hypertext. Clicking the hypertext will result in navigating to another page to display data related to that column based on the configuration for that column.
- **Maximum height.** The maximum height of the report area is controlled by the zone's configuration. A vertical scroll will appear when more rows exist than fit into the height.
- **Maximum rows.** The maximum number of rows that can appear is controlled by the zone's configuration. Excess rows are truncated before display. Please note, you can configure the zone to [download more rows to Excel](#) than appear in the Results Area.
- **Control column sorting.** By default zones are designed to allow a user to sort the results list by clicking any column heading. Clicking once sorts ascending. Clicking again sorts descending. However, a zone may be configured to not allow this type of sorting.
- **Rearranging columns.** A zone may be configured to allow a user to rearrange columns and add or remove columns visible in the Result Area. Refer to [Column Configuration](#) for the details.


The type and variety of columns that appear in the Results Area is controlled by the zone's configuration. You can even set up an info zone to display information in hierarchical trees. For example, the following info zone from Oracle Utilities Customer Care and Billing has been configured to show a person's related persons.



Title Bar

The Title Bar contains the zone's description along with other icons and functions described in this section.

In the corner of the Title Bar, you'll see an indicator that tells you whether the zone is collapsed (e.g., minimized) or expanded. Each time you click the indicator, the zone toggles between expanded and collapsed. The following image shows a collapsed zone.



NOTE: Recommendation: For zones that are configured to show on portal preferences, set up [Portal Preferences](#) to collapse zones that aren't needed every time a portal is displayed. Rationale: The system doesn't build collapsed zones


when a portal is displayed. Instead, collapsed zones are only built when a user expands them. Thus, indicating a zone is collapsed improves response times.

If the zone is configured with help text, a help icon appears to the right of the zone title. Click this icon to view the help text.

If a [saved search](#) is open for a zone, the name of the saved search is displayed after the zone title.

The following sections describe additional icons and tools visible on the title bar. Note that in addition to the information described in this section, additional actions that are specific to a zone may appear based on configuration. The actions may appear in the form of hypertext, an icon or a button.

Explorer Zone Menu

Click the Explorer Zone Menu icon  to list available tools. The following topics describe each tool.

Clear Filters

Click the **Clear Filters** link in the Explorer Zone menu to clear the filter values and zone data.

Export to Excel

Click the **Export to Excel** link in the Explorer Zone menu to download the zone's data to Excel. This link is hidden if the zone has been configured to not download rows to Excel.

Print

Click the **Print Zone** link in the Explorer Zone menu to print the zone.

Save Search Options

Query zones and info zones that include filter areas, allow a user to capture and save filter criteria so that the same search can be more quickly executed in the future. Users may save many different search options for the same zone. Users also have the option to nominate one search as the default so that this search is automatically loaded upon entering the zone. In addition, users may choose to set one or more searches as a Favorite, which causes the search to appear in the [Favorite Searches](#) zone in the dashboard. This allows a user to navigate to the portal for that zone and load the saved search in one click.

The Explorer Menu includes various options for supporting Saved Searches that are described below.

Note the following additional points about this functionality:

- **Click Search before Saving.** When a user wants to save a search, the search criteria should be entered and the **Search** button should be clicked prior to saving the search.
- **Criteria is saved, not results.** Although a user must click **Search** prior to saving the search, the system saves the criteria, but not the results. When returning to the query or info zone in the future and opening a given saved search, the results displayed will be the current records in the system that match the saved search criteria.
- **Multi-Query Zones.** When a user is on a multi-query zone, the searches are saved at the multi-query zone level so that all saved searches for the various search dropdowns for a given multi-query zone are visible in the Open Saved Searches menu option.
- **Functionality is Secured.** If for any reason an implementation wishes to restrict which users are allowed to save searches, the functionality is securable. Only users with **Inquire** access to the **F1_USRFAVSCH** application service will see these options in the explorer menu. This is the same application service that is used to provide access to the Favorite Searches dashboard zone.

Save As

The **Save As...** option in the Explorer Zone menu allows a user to save the current entered search criteria. This should be clicked after executing the search.

The user is prompted to provide a name for the search. This search is now available from the Open Saved Searches option. If desired, a user can also choose to make this search the Default or may choose to Add to Favorites.

This option is visible whether or not an existing saved search is loaded. When a saved search is loaded, this option allows a user to make changes to the search criteria and save a new search with that criteria.

Open Saved Search

The **Open Saved Search...** option in the Explorer Zone menu is visible if there are any saved searches for the zone that is currently visible. Select this menu entry to view the current list of saved searches for the zone. For multi-query zones, any saved search for any of the query zones that are part of the multi-query zone are visible. Choose the desired saved search. The criteria is updated with this information and the search is performed.

Save

The **Save** option in the Explorer Zone menu is visible if a Saved Search is already loaded. This option allows a user to save changes to the criteria of an existing saved search (overwriting the previously saved criteria).

Set As Default

The **Set As Default** option in the Explorer Zone menu is only visible if a saved search is loaded and it is not currently the default. Click this option to nominate this search to be the one that is loaded when navigating to this zone from the menu.

Remove As Default

The **Remove As Default** option in the Explorer Zone menu is only visible if a saved search is loaded and it is currently the default. Click this option to no longer load this search automatically when navigating to this zone from the menu.

Add To Favorites

The **Add To Favorites** option in the Explorer Zone menu is only visible if a saved search is loaded and it is not currently included in the favorite searches zone. Click this option to include this search in the [Favorite Searches](#) zone. Once a search is in the Favorite Searches zone, you can use the link to the search in the dashboard to navigate to the portal for this zone and load the saved search in one click.

NOTE: Not available for dependent zones There are some portals that include a zone that is not the main query or info zone on the portal, but is rather a zone providing additional information for a record that is displayed. If a zone like this has filter options or configurable columns, a user is able to save searches and mark one as a default, but will not be able to add the saved search to favorites because the zone is dependent on other data on the page that is not captured for the individual zone.

NOTE: Dashboard Zone Preferences. Be sure that your [portal preferences](#) are configured to include the Favorite Searches dashboard zone.

Remove From Favorites


The **Remove From Favorites** option in the Explorer Zone menu is only visible if a saved search is loaded and it is currently included in the [Favorite Searches](#) zone. Click this option when you want to keep the saved search but you no longer want it visible in your Favorite Searches zone. (Deleting a saved search that is also marked as a favorite will automatically remove the search from the favorite searches zone.)

Delete

The **Delete** option in the Explorer Zone menu is only visible if a saved search is loaded. Click this option to remove the currently opened saved search from the list of saved searches.

Note that if the saved search had been configured as favorite, deleting the saved search remove it from the favorite searches zone.

Show Filters

Click the **Show Filters** icon  to show the [Filter Area](#). This button only appears if the zone has been configured to have filters.

NOTE: If the zone is configured to have filters and the drag and drop (column configuration) area, only one or the other is shown at any given time. It means that a user can either change the filter criteria and refresh the search results or change the column configuration and refresh the results.

Show Column Configuration Area

Click the **Column Configuration** icon  to show the [Column Configuration Area](#). This button only appears if the zone has been configured to allow columns to be dragged and dropped into the Results Area.

NOTE: If the zone is configured to have filters and the drag and drop (column configuration) area, only one or the other is shown at any given time. It means that a user can either change the filter criteria and refresh the search results or change the column configuration and refresh the results.

The Open Zone Maintenance Icon

Click the "wrench" icon  to switch to the zone's [zone maintenance](#) page. On this page, you can change the parameters that control the zone's behavior.

NOTE: Debug mode only. This icon only appears when the system operates in [Debug](#) mode.

Configurable Columns


Info zones and query zones may be configured to allow users to rearrange columns and to remove or add columns to the output. This is also referred to as "drag and drop" of columns.

When a zones is configured to support this functionality, the [Show Column Configuration Area](#) icon is visible in the zone title area. Click this icon to expose the Column Configuration Area, if it is not visible and click the icon to hide the Column Configuration Area if desired.

If configurable columns are allowed, a user may rearrange the order of columns currently displayed in the Results Area by clicking the column heading and "dragging" it to the desired location.

In addition, the column configuration area may display columns that are not currently in the Results Area. These are indicated by a cream background. A user can drag one of these columns to the results area and then click **Refresh** to rebuild the zone. To remove a column that is currently in the results area, click the trash can icon next to the column name in the column configuration area and then click **Refresh** to rebuild the zone.

The following is an example of how the Column Configuration Area looks:



Multi-Query Zones

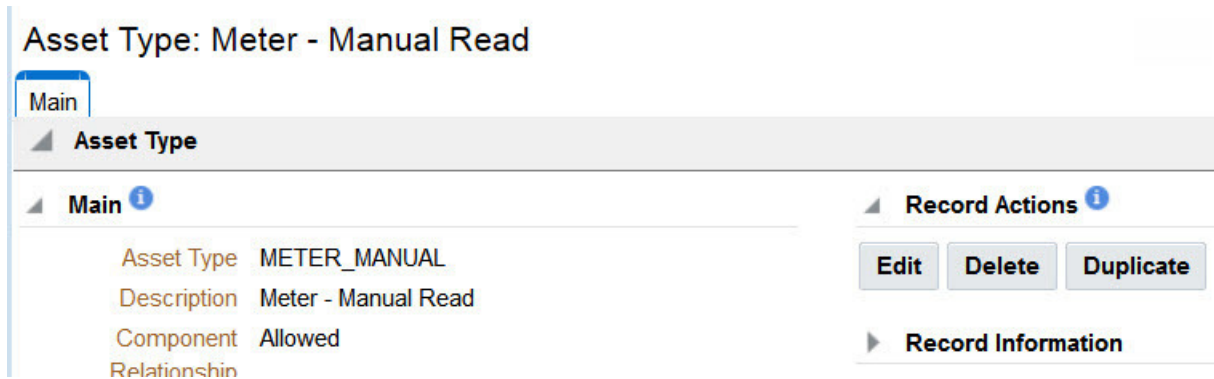
This type of zone combines individual Query Zones into a single query zone. A **Search By** dropdown, built from the zone descriptions of the individual Query Zones, allows you to select the desired query to execute. Selecting a specific Query Zone display the filters of that query in the filter area. At this point the zone simply executes the selected Query Zone, as if it was executed on a standalone zone. By default, the first query zone option is selected.



Map Zones

A map zone displays information related to one or more objects. For objects that are maintained using portal-based pages, the standard display of the object is using a map zone. Map zones may be configured to display an explicit map or may dynamically display the map based on the configuration of the record being displayed.


Because the information displayed in a map zone is configured using HTML, virtually any format supported by HTML is possible. The following is an example of a typical map zone used to display information about an object.



A [tool tip](#) can be used to display additional help information for the details shown in a map zone. The tool tip may appear adjacent to the section header or adjacent to individual elements on a map.

NOTE: You can customize the help icon or disable map help using the [Custom Look and Feel](#) feature configuration.

Open Map Maintenance Icon

Click the "hammer" icon  to switch to the [UI map](#) maintenance page. On this page, you can view the HTML configuration for the zone. Note that for map zones that display a business object configured without a map, the elements displayed are automatically rendered. In this case, no "hammer" icon is visible.

NOTE: **Debug mode only.** This icon only appears when the system operates in [Debug](#) mode.

Timeline Zone

A timeline zone is used to show events over time for multiple types of data. This type of zone is helpful for providing a consolidated view of separate events occurring in time that may in fact impact each other. For example, in the billing system, a timeline can show bills, payments and collection processes for the same time period.

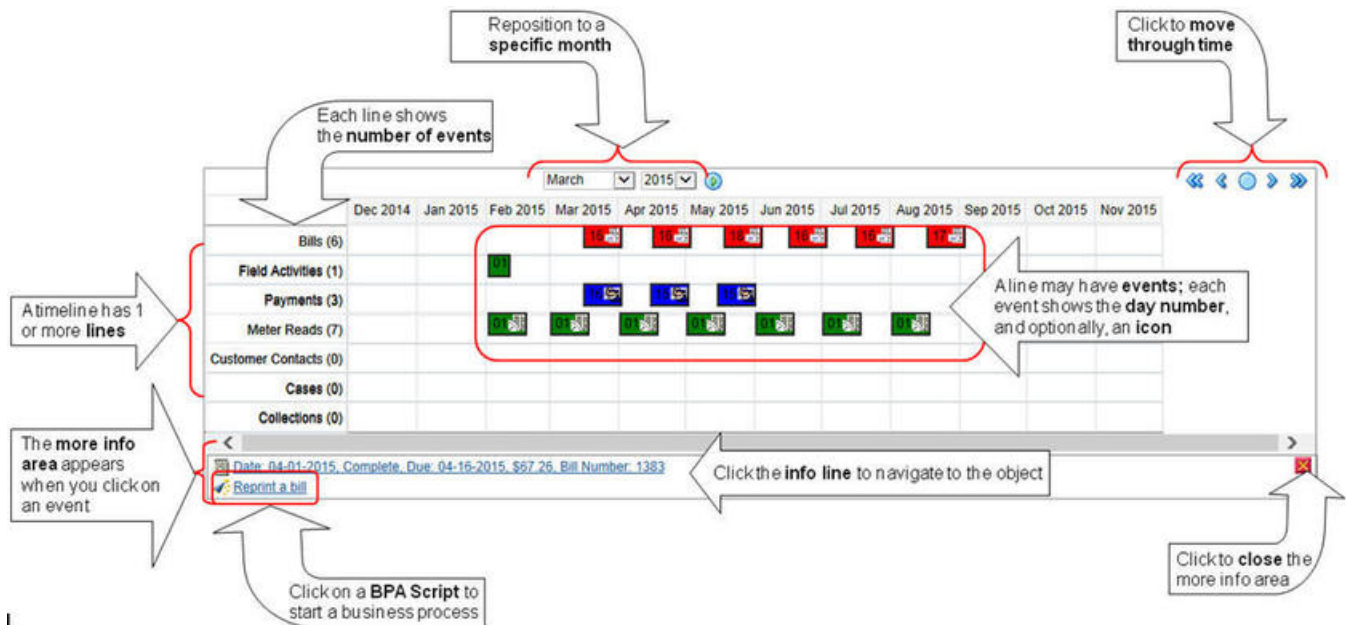
The topics in this section describe the rich functionality available in timeline zones.

Timelines Zones Are Configured By Your Implementation Team

The product provides support for creating timeline zones. However, whether timeline zones are configured for your implementation and on which portals they may be visible will depend on your specific configuration. Refer to [Configuring Timeline Zones](#) for general information about configuring timeline zones. Refer to your specific edge application's documentation for details about what type of timeline information is provided out of the box.


The Anatomy Of A Timeline Zone

The following illustration highlights the various elements in a timeline zone. These elements are described in the remaining topics.



You Can Move Through Time

You can click the controls at the top of a timeline zone to change the date-range of the zone's information:



The following points describe how to use these controls:

- To reposition the timeline to a specific date, selected the desired month and year and click the search button.
- To go back one year, click the double-left arrow.
- To go back one month, click the single-left arrow.
- To go to today, click the red line (between the arrows).
- To go forward one month, click the single-right arrow.
- To go forward one year, click the double-right arrow.

Timelines Can Have Many Lines

Depending on how a timeline zone is configured, it will typically have one or more "lines" that show when significant events have occurred. Each line represents a specific type of data. The data may be related to a specific maintenance object, for example, payments may be displayed in one line and bills in another. It may also be possible to create multiple lines for different types of records for the same maintenance object. For example, if your product provides an algorithm to display cases, it may support defining a different entry for each case type.

Each Line Shows Events

Each line on a timeline may contain zero or more events where each event shows the date when the event occurs. For example, the payment line in a timeline has a separate event for every payment received from the customer. Each line's description contains the number of events on the line.

CAUTION:

If a line has more events than can fit onto a timeline, the line will show the first "chunk" of events and a message will appear in the "more info area" explaining that some events have been truncated. If this happens and the truncated events are in a later period, you can reposition the timeline's base period to show the truncated events. Refer to the You Can Move Through Time section above for more information

An Event's Color And Icon Can Convey Information About The Event

The timeline zone is driven by algorithms and the algorithms may be configured to highlight different types of events in different colors and / or icons. For example, a timeline that shows payments may use different colors and/or icons to visually differentiate between canceled payments and frozen payments.

An Event's Hover-Text Can Contain Additional Information

The timeline zone supports providing hover text that appears when a user hovers the mouse over an event. Each type of event can have different hover text. Typically the standard information string is shown for each event.

Clicking On An Event Shows More Information In The Detail Area

When you click on an event, additional information appears in the zone's detail area (at the bottom of the zone). The following information may appear:

- The event's common "information string" appears. This information is hypertext to allow for easy access to the transaction on which the object is maintained.
- Additional information may appear. This is dependent on the particular timeline algorithm. Refer to each timeline algorithm's description to find out if additional information appears in the detail area under any conditions.
- If the algorithm has configured BPA scripts that can be executed to perform business processes on the object, the BPA script description appears prefixed with a "wizard's hat" icon.

Common Base Portals and Zones

The following sections describe common standards followed by the product for base query and maintenance dialogue.

All-In-One Portal

The all-in-one portal is used for low volume objects, where a small number of records are expected.

For example, many control tables use this style of portal. In this portal, rather than separate query and maintenance portals, a single portal allows a user to find a record and maintain the record. A zone listing all existing records is displayed when first entering the portal. A user can perform actions on existing records or add new records from there. More details about the list zone are provided below.

Refer to [Maintenance Portal](#) for information about zones commonly found on maintenance portals.

The all-in-one portal will typically provide the ability to add a new record. This may be a button in the Page Action area or a hyperlink defined directly on the list zone's header area.

List Zone

This info zone appears when first entering the portal and displays a row for each record.

The following functions are commonly available in a list zone:

- A broadcast icon is included in each row. Clicking this icon opens other zones that contain more information about the adjacent record.
- An **Edit** icon may be included in each row. Clicking this icon launches the maintenance BPA script to modify the adjacent record.
- A **Duplicate** icon may be included in each row allowing the user to duplicate the adjacent record without having to first display the record.
- A **Delete** icon may be included in each row allowing the user to delete the adjacent record without having to first display the record.
- If the record has a simple BO lifecycle of **Active** and **Inactive**, the list zone may include a button to **Activate** or **Deactivate** the adjacent record.

Separate Query and Maintenance Portals

For high volume objects where more sophisticated search criteria is needed to find a particular record, separate portals are provided to search for the record and to maintain it.

The query portal contains a zone to search of the object that you wish to maintain. Additional information about the query portal's search zone is provided below. Clicking an appropriate record in the search results navigates to the object's stand-alone maintenance portal.

The query portal will typically provide the ability to add a new record. This may be a button in the Page Action area or a hyperlink defined directly on the list zone's header area.

Query Portal Search Zone

The query portal contains one zone used to search for the object.

The zone may be a standard query zone or may be a multi-query zone providing one or more ways to search for your record.

Once the search results are displayed, click on the hypertext for the desired record's information string to navigate to the stand-alone maintenance portal.

Maintenance Portal

A maintenance portal refers to a page that is used to view and maintain a specific record. These types of records are governed by a business object, which defines the user interface layout for both the display of the record and the add or update of the record.

An all-in-one portal is a type of maintenance portal where the maintenance zone is displayed once a record has been broadcast from the list zone. For the separate query and maintenance portal dialogue, the maintenance portal includes only zones relevant for the specific record selected.

Maintenance portals may have multiple tabs to display information relevant for the object. Often, if the object has a Log, the log is on a separate tab. (However, there may be some maintenance portals that include the log zone on the Main tab.)

The stand-alone maintenance portal is typically configured with the ability to navigate back to the query portal using either a **Search** button in the Page Actions area or using a **Go To Search** hyperlink in the main zone's header area. The stand-alone maintenance portal may also include the ability to add a new record using an Add button in the Page Actions area. (If not, a new record can be added by navigating through the menu or by navigating to the query portal where an Add hyperlink may be found.)

The following sections describe zones that are common to portals used to maintain a record. Note that additional zones may exist for specific all-in-one portals or stand-alone maintenance portals based on the particular functionality. Refer to the documentation for each portal to understand the specific zones that are provided.

Object Display Zone

This is a zone that displays the object being maintained.

An object being displayed may or may not allow for actions or updates to be made. Typically there is a Record Actions section in the display area that presents the valid actions. Depending on the functionality, the actions may prompt the users to input values required for that action.

Log Zone

If the maintenance object for the record includes a status, it typically has a log that records important events in its lifecycle.

This log is shown in an info zone on the portal, either on the Main tab or on a separate Log tab. A link appears in the zone's title bar to allow the user to add a new manual log entry.

Process Flow

This portal is used to guide a user through the steps needed to complete a process flow. If the process flow type is configured to keep a record of a completed process flow then this portal may also be used to display the details of such record.

Refer to [Understanding Process Flows](#) for more information.

Each type of process flow is assumed to be configured with designated custom menu options that allows you to launch a new process flow of that type as well as review saved off and completed records. When reviewing existing process flow records, you are brought to a query portal with options for searching for a specific record. You may **Resume** a saved off but not yet finalized process flow or review a complete one.

The **Process Flow** zone on the portal's **Main** tab page assists the user in completing the process flow as well as displays the information captured on a completed process flow.

Fixed Pages

There are some pages in the system that are built using a fixed page metaphor. All user interfaces implemented in this way function in a similar manner. This section describes functionality to these types of pages.

NOTE: Not all pages in the application use this metaphor. The preferred metaphor for pages in the system is to use portal-based pages. For more information see [Common Base Portals and Zones](#).

Maintenance Pages

Maintenance pages are used to maintain (add, change, etc.) the information in the system using a fixed page metaphor. This section focuses on functionality common to all maintenance pages using this metaphor.

The Unique Identifier Area

Every object in the system must have a unique identifier. For some objects, typically administrative or control table data, the unique identifier is defined by the user creating the record. Other objects, typically master or transaction entities, the unique identifier is assigned by the system after the record is saved.

Where the identifier is displayed depends upon whether the ID is user-defined or system generated.

Pages displaying objects with a user defined identifier (such as To Do Type) typically display the identifier on the left along with a search button. For master and transaction tables, (such as To Do Entry) the unique identifier is typically found in the upper right corner of the screen along with a search button and is labeled as the entity's ID. In this case, if there is an information string associated with the object, it is displayed on the left.

If the object being maintained has an associated context menu, a [Context Menu Button](#) is displayed appear.

Confirmation & Error Messages

Whenever a database modification action is requested, the system first validates the data to be put on the database.

First, the system validates the information. If there are errors, the first error encountered is displayed in a window. In addition to the error message, the error window also contains the error number and the identities of the programs that detected the error. If you don't understand an error and you have to call product support, please supply them with this information.

If there are no errors, the system is modified.

Warning If Unsaved Changes Exist

If you make any changes to an object and you don't commit these changes to the database, the following warning appears: "You have unsaved changes that will be lost if you continue. Press OK to continue anyway."

If you want to commit the changes, click **Cancel**. You will then return to the page and can make further changes or press Save. If you want to discard your changes and transfer to the new page, click **OK**.

CAUTION: This warning is not displayed if you are working on a *new* object on a fixed page and you navigate to a different page. Information about new objects is preserved in memory in anticipation of your return to the page. For example, if you are in the middle of adding a new user and you want to quickly reference information on a different page, you will not receive the above warning because your information will be saved. If you use the back button to return to the user page, all of the information you were adding will be present and you can enter the rest of the new user's data.

Query Pages

Query pages are used to display record sets in the system. Most query pages function in a similar manner. This section focuses on functionality common to all of the system's query pages.

The Search Area

The upper right corner of the screen is almost always the identifier of the object associated with the search. In Oracle Utilities Customer Care and Billing for example, the unique identifier on the Account Financial History page is the Account ID.

To the right of the unique identifier is a [Search button](#). This button is used to search for a different object than the one whose data is being displayed in the page.

If the object being displayed has an associated context menu, a [Context Menu Button](#) is displayed.

User Preferences

The Preferences (or My Preferences) transaction and the User Maintenance transactions are the same except that with My Preferences, only the following items can be modified:

- Main Tab
 - Language
 - Display Profile
 - Time Zone
 - Email Address
 - Dashboard Width
 - Home Page
 - To Do Summary Age Bar Parameters
- Bookmarks
 - All information can be modified

- Portal Preferences
 - All information can be modified
- Favorite Links
 - All information can be modified
- Favorite Scripts
 - All information can be modified

User - Main

This page lets you view or define user information and access settings.

To open the page, use one of the following navigation options:

- The **Admin > Security > User > Search** menu item.
- The **Menu > My Preferences** menu item.
- The **User Menu > Preferences** menu item in the application toolbar.

Note that some fields are only modifiable when the page is opened from the from the **Admin** menu.

Description of Page

Basic user identification information is viewed or set in the **First Name**, **Last Name**, **User ID**, and **Login ID** fields, all of which are modifiable only when the page is opened from the **Admin** menu. User and logon IDs must be unique. The Login ID is the id that the user uses to log in to the system. The User ID is the ID that is used throughout the application when users are assigned to tasks or when the user ID is stamped on records when logging add or update actions.

Language Code defines the language in which the user's screen prompts and messages appear.

Set **User Enabled** to allow the user to access the system.

NOTE: This field, available only when the page is accessed through the **Admin** menu, provides special field-level security. A user can change the field only if the **Enable/Disable** access mode is set.

Display Profile ID defines the [display options](#) for the user.

The next group of fields is used to view or set user preferences. For example, you can set up [portal preferences](#) on a "template user" and indicate that specified users will inherit their portal preferences from the template. Doing so provides consistency in the application of account information and dashboard zones. This fields are modifiable only when the page is accessed through the **Admin** menu.

- Set **User Type** to **Template User** if the user's preferences are to match those of other users.
- If a user's [portal preferences](#) are to be inherited from a "template user", use the **Portals Profile User ID** to define the user ID of the template user. If applied, the specified user will inherit portal preferences from the **Portals Profile User ID** and will be unable to customize those preferences.
- If a user's [favorite links](#) are to be inherited from a "template user", use the **Favorites Profile User ID** to define the user ID of the template user. If applied, the specified user will inherit favorites preferences from the **Favorites Profile User ID** and will be unable to customize those preferences.

The **Time Zone** associated with the user may be defined. Implementations may use this information for features and functions where the user's time zone impacts display of data.

Email ID captures the user's email address.

Dashboard Width defines the number of pixels to use for the user's [dashboard](#) area. A setting of 200 is recommended as a minimum. Depending on the width of a user's monitor, a higher setting (such as 250) will display more information without compromising the space on the main area of the application. Setting the value to 0 will suppress the dashboard.

Home Page defines the page that appears when the user starts a session or clicks the [Home](#) link.

The To Do entry fields control how To Do entries are aged in respect of the "age bars" that appear on [To Do Summary](#), [To Do Supervisor Summary](#) and the [To Do Summary Dashboard Zone](#).

- Use **Lower Age Limit for Yellow Bar** to define the number of days old a To Do entry must be to be considered "yellow" in the age bar. To Do entries below this limit are considered "green" in the age bar.
- Use **Upper Age Limit for Yellow Bar** is the high limit for the "yellow" bar. To Do entries older than this are considered "red" in the age bar.

The **User Group** collection defines the group(s) with which the user is associated and the date (if any) when the user's association expires. This feature is modifiable only when the page is accessed through the **Admin** menu. For more information, see [The Big Picture Of Application Security](#). Assignments of users in a user group can also be established or maintained through the [User Group - Users](#) page.

NOTE: The ability to maintain the user groups on this page requires special field-level security. A user can change the field only if the **System Administration** access mode is set.

If your implementation has configured [Object Erasure](#) for the User maintenance object, the **Object Erasure Schedule** dashboard zone is visible displaying the erasure status for this user record, if applicable. Refer to [Erasing User Information by Obfuscation](#) for more information on object erasure for User records.

User - To Do Roles

Open the User page and then navigate to the **To Do Roles** tab to define the To Do roles assigned to the user.

The information on this page cannot be modified if accessed through **My Preferences**.

FASTPATH: A user's To Do roles control the types of **To Do** entries the user can access. For additional information on To Do roles, see [To Do Entries Reference A Role](#).

Description of Page

The grid contains the user's To Do roles. To modify a To Do role linked to the user, simply move to a field and change its value. To remove a To Do role, click the - (minus) button. To add a new To Do role, click the + (plus) button and enter the role.

Where Used

Every **To Do** entry references a role. A To Do role has one or more users (and a user may belong to many To Do roles). Users who are part of the To Do role assigned to a To Do entry may work on the To Do entry. For more information on this feature, see [The Big Picture Of To Do Lists](#).

User - Access Security

Open the User page and then navigate to the **Access Security** tab to define a user's security rights. You cannot modify this page if you accessed it through **My Preferences**.

FASTPATH: Refer to [The Big Picture of Application Security](#) for more information about data access roles and access groups.

Description of Page

Use the **Default Access Group** to define the access group that is defaulted on new records added by this user that are subject to [row security](#).

The scroll area contains the **Data Access Roles** to which this user belongs. A user's data access roles play a part in determining the accounts whose data they can access.

To add additional data access roles to this user, click the + button and specify the following:

Enter the **Data Access Role**. Keep in mind that when you add a **Data Access Role** to a **User**, you are granting this user access to all of the accounts linked to the data access role's access groups.

Use **Expiration Date** to define when the user's membership in this data access role expires.

NOTE: You can also use [Data Access Role - Main](#) to maintain a data access role's users.

The tree highlights the **Access Groups** to which the above **Data Access Roles** provide access.

User - Bookmarks

This page is used to manage a user's bookmarks that are created when a user clicks the [Bookmark](#) button and saves a bookmark. From this page the user can delete a bookmark, change the name of the bookmark or change the sort sequence.

Open the User page and then navigate to the **Bookmarks** tab to modify or delete a bookmark.

Description of Page

Sequence is the sort sequence associated with the bookmark. Bookmarks with the same sequence are then sorted alphabetically.

Name is the user-defined name assigned to the bookmark. This is the text that is visible in the [Bookmarks](#) dashboard zone.

NOTE: Favorite Links vs. Bookmarks. Favorite links are configured through the user page and define pages that a user can navigate to. They do not define context (although some pages may use information currently in global context when navigating). Some users may not be allowed to configure their own favorite links and would see the favorite links associated with a profile user. Bookmarks are created using a bookmark button available on each page. The bookmark saves navigation information, context information and also captures some information about the state of a page. Bookmarks are not governed by any user profile configuration.

User - Portal Preferences

The base product contains several [portals](#) that allow users to customize them via portal preferences. Portal preferences allow users to control:

- Which zones appear on the portal
- The order in which the zones appear
- Whether the zones should be "collapsed" (i.e., minimized) when the portal opens.

NOTE: You may not be able to change your portal preferences. If a note appears immediately before the list of portals, a system administrator has configured your user ID to reference a **Portals Profile User ID** (this is defined on the **Main** tab). Preferences set in this way cannot be modified. System administrators do this in order to enforce a common look-and-feel throughout the user community.

Open the User page and then navigate to the **Portal Preferences** tab to modify a user's preferences.

Description of Page

The [accordion](#) contains a row for every portal configured to show on user preferences and to which you have access. To change how a portal's zones appear, expand the respective row and change the elements accordingly. The remainder of this section describes how you can configure how a portal's zones appear.

Zone lists all of the [zones](#) in the portal.

Place a check under **Display** if the zone should be displayed on the portal.

Place a check under **Initially Collapsed** if the zone should be minimized when the portal is initially opened. Setting zones as initially collapsed is a good idea to save space and to help pages to load more quickly. Refer to [Zones May Appear Collapsed When A Page Opens](#) for more information.

CAUTION: Recommendation. It is recommended that you set your preferences to collapse zones that you don't use often. This is because the system doesn't perform the processing necessary to build collapsed zones until you ask to expand a zone. Opening a page with zones collapsed accelerates the response times of portal.

Use **Sequence** to control the position of the zone in respect of the other zones. This field is disabled if **Display** is not checked. Duplicate **Sequence** numbers are allowed.

For zones displaying data that is changing often, use **Refresh Seconds** to define in seconds how often the zone is refreshed. The minimum valid value is 15. The maximum valid value is 3600 (1 hour). A value of 0 indicates no automatic refresh.

Security Access indicates whether the user has access rights to a zone. Refer to [Granting Access To Zones](#) for more information.

User - Favorite Links

This page is used to define a user's favorite transactions and URLs. Each such link appears as an entry in the [Favorite Links Zone](#) in the [Dashboard](#) area.

NOTE: A hot key can be used to execute a favorite link. Rather than clicking a button in the favorite links zone, you can invoke your first nine favorite links by pressing `Ctrl + a` number (where the number corresponds with the relative position of the transaction in your list of favorites). Note that the numeric keypad on your keyboard cannot be used to implement this function.

Open the User page and then navigate to the **Favorite Links** tab to modify a user's favorite links.

Description of Page

NOTE: You may not be able to change your favorite links. If a **Note** appears beneath your **User ID**, your user id has been setup to reference a specific **Favorites Profile User ID** as defined on the **Main** tab. If a system administrator has setup your user ID in this way, your preferences may not be changed. System administrators may apply this feature in order to enforce a common look-and-feel throughout the user community.

Each row references a [Navigation Option](#). The navigation options supplied with the base package correspond with the system's menu items. This means you can setup a favorite link to open any transaction in the system (in either **add** or **update** mode).

NOTE: Favorite links can point to other places. The navigation options supplied with the base package always open a transaction's **Main** tab. If you find yourself continually accessing a different tab page, you can have your implementation team create a new navigation option that references the desired tab page. Once this navigation option is set up, you can reference it as one of your favorite links. In addition, it's also possible for your implementation team to setup navigation options that open a URL outside of the system. For example, your implementation team could setup a navigation option to a credit rating company's website. Once this navigation option is set up, you can reference it as one of your favorite links.

Use a **Sequence** to control the order in which the navigation option appears in the [Favorite Links Zone](#).

Security Access indicates if you have been given security rights to the respective transaction. Refer to [The Big Picture Of Application Security](#) for more information.

NOTE: Favorite Links vs. Bookmarks. Favorite links are configured through the user page and define pages that a user can navigate to. They do not define context (although some pages may use information currently in global context when navigating). Some users may not be allowed to configure their own favorite links and would see the favorite links associated with a profile user. Bookmarks are created using a bookmark button available on each page. The bookmark saves navigation information, context information and also captures some information about the state of a page. Bookmarks are not governed by any user profile configuration.

User - Favorite Scripts

This page is used to define a user's favorite scripts. Each script is displayed in the **Favorite Scripts** zone on the **Dashboard** portal where it can be used to [initiate a script](#).

NOTE: A hot key can be used to execute a favorite script. Rather than selecting a script from the business process assistant, you can invoke your first nine favorite scripts by pressing `Ctrl + Shift + a number` (where the number corresponds with the relative position of the script in your list of favorites). Note that the numeric keypad on your keyboard cannot be used to implement this function.

Open the User page and then navigate to the **Favorite Scripts** tab to modify a user's Favorite Scripts.

Description of Page

NOTE: You may not be able to change your favorite scripts. If a **Note** is displayed on top of the grid, your user id has been setup to reference a specific **Favorites Profile User ID** as defined on the **Main** tab. If a system administrator has setup your user ID in this way, your preferences may not be changed. System administrators may apply this feature in order to enforce a common look-and-feel throughout the user community.

Each row in the grid contains one of your favorite **Scripts**.

Use the **Sort Sequence** to control the order in which the scripts appear in the [business process assistant menu](#).

Security Access indicates whether you have access to the script's application service. Refer to [The Big Picture Of Application Security](#) for more information.

User - Characteristics

Select **Admin > Security > User > Search** and navigate to the **Characteristics** tab to define characteristic values for the user. You cannot modify this page if you accessed it through My Preferences.

Description of Page

Use the characteristics grid to capture additional information about the user.

You can only choose **Characteristic Types** defined as permissible for the user record. Refer to [Setting Up Characteristic Types & Their Values](#) for more information.

Enter a valid **Characteristic Value** for the selected characteristic type. You may enter more than one characteristic row for the same characteristic type, each associated with a unique **Sequence** number. If not specified, the system defaults it to the next sequence number for the characteristic type.

User - Miscellaneous

Select **Admin > Security > User > Search** and navigate to the **Miscellaneous** tab to view additional information associated with a user.


User Id, First Name and Last Name are displayed. Depending on your implementation, additional zones may appear on the User Portal

If additional zones appear, use the embedded help for information about the fields managed on the zone.

Initiating Scripts

Your implementation team can set up business process assistant (BPA) [scripts](#) to walk you through business processes. Once scripts have been set up, you can initiate a script using any of the following methods:

- You can define a list of your "favorite scripts" on your [user preferences](#). You can then invoke any of the scripts from the Favorite Scripts zone on the dashboard portal. You can also invoke any of your first 9 favorites by pressing `Ctrl + Shift + a number` (where the number corresponds with the relative position of the script in your list of favorites). For example, if you press `Ctrl+Shift+2` and you've setup your "favorite scripts" to be **Start Autopay** and **Stop sending marketing information**; the **Stop receiving marketing information** script is invoked. Note that the numeric keypad on your keyboard cannot be used to implement this function.
- You can configure a menu item to launch a script by referencing a [navigation option](#) that is a **script** navigation option type.
- The system creates [To Do entries](#) to highlight tasks that require attention (e.g., bills in errors, accounts without bill cycles, etc.). Users can complete many of these tasks without assistance. However, you can set up the system to automatically launch a script when a user selects a To Do entry. For example, consider a To Do entry that highlights a bill that's in error due to an invalid mailing address. You can set up the system to execute a script when this To Do entry is selected. This script might prompt the user to first correct the customer's default mailing address and then re-complete the bill. Refer to [Executing A Script When A To Do Is Launched](#) for more information.
- You can indicate that a script should be [launched upon entering the system](#). This functionality enables an external system to launch the system and immediately launch a script.
- The following [hot keys](#) cause an interactive menu is displayed:
 - `Ctrl+Alt+Q`
 - `Ctrl+Shift+S`



The menu allows you to search for a script using any word in its description. Note, if you've turned on **Search only Eligible Scripts**, your search is limited to [eligible scripts](#). Note, the system displays an indication if each script in the grid is **Eligible** or **Ineligible** in this section.

You can initiate a script from this area by clicking on it.

When a script is initiated, this menu closes and the [object display area](#) shifts down to make room for the script area.

NOTE: Securing scripts. Only scripts the user is allowed to execute are displayed. Refer to [Securing Script Execution](#) for more information.

NOTE: Customize the script area. Using the [Height](#) step type, an individual script may be configured to use more or less of the script area than the standard script area. It is also possible to configure a script to hide the script area altogether.

The information in the script area takes you through the steps in the selected business process (note, your implementation team controls the information that appears in the script area).

The following points highlight useful information common to all scripts:

- **Return focus to script area.** Often during a script the focus moves from the script area to another section of the page. For example, you might be required to fill in data for a new record being added. When you are ready to return to the script area, rather than using the mouse to return to the area, you can press `Ctrl+Shift+F`. The focus returns you to either the first available input field or to a "default" button if no input field is available.
- **The Continue button.** Many steps in a script ask you to click the **Continue** button when you're ready to restart a script. For example, a step may ask you to confirm a customer's automatic payment information and then click **Continue** when you're ready for the script to restart. Rather than using the mouse to click this button, you can press `Ctrl+Shift+C`.
- **Selected button changes color.** When you click a button in the script area, the button selected changes color so that you can easily see which option you selected.

FASTPATH: Refer to [The Big Picture Of Scripts](#) for more information about scripts.

NOTE: You can abort a script at any time by clicking the red X button in the upper right corner of the script area (or by pressing `Ctrl+Shift+X`).

Shortcut Key Summary

The following table summarizes the keyboard shortcuts (hot keys) available in the system:



NOTE: It is possible that your specific product supports additional shortcut keys. Refer to your product documentation for more information.

NOTE: It is possible that some of the following system accelerator keys may duplicate shortcuts, or hot keys, that are predefined within your browser. If this is the case, the browser's functionality commonly takes precedence.

NOTE: It is possible that your implementation uses one of the accelerator keys listed below for a different purpose. For example, in some languages, a shortcut key may be used for a special character in the alphabet not included on all keyboards. It is possible to turn off any of the `Alt+letter` shortcuts and `Ctrl+Alt+letter` shortcuts listed below using a properties file entry. The property is `ouaf.shortcut.ignore.altKeys=` where one or more letters may be entered separated by a comma. When setting this property, the system automatically enables a 'potential alternate key' for certain functions (see the table below). Refer to the *Server Administration Guide* for more information about the property.

Accelerator Key	Potential Alternate Key (see Note above)	Function
<code>Ctrl+Alt+A</code>		Opens the Admin Menu button.
<code>Alt+B</code>		Returns to the previous page on which you were working. This is the equivalent of clicking the Back button.

Accelerator Key	Potential Alternate Key (see Note above)	Function
Ctrl+Alt+B		Launches the Bookmark dialog. This is the equivalent of clicking the Bookmark button.
Alt+C, Ctrl+Alt+C, Alt+Shift+C		Clears the currently displayed object from the page. This is the equivalent of clicking the Clear button.
Ctrl+Shift+C		If the current step in the Script area has a Continue button, this shortcut is equivalent to clicking Continue .
Ctrl+Alt+D, Alt+D		Creates a duplicate of the object being displayed. Equivalent to clicking the Duplicate button.
Ctrl+Alt+F		Positions the user's cursor to the Menu Item Search input field.
Ctrl+Shift+F		If the Script area is open and you are working on a different area of the page, this key returns focus to the script area. Focus returns to either the first available input field or to a "default" button if no input field is available.
Alt+G		Returns to the most recent page that was displayed when the Back button was clicked. Equivalent to clicking the Forward button.
Ctrl+Alt+H		Opens a dropdown list of recently-visited pages, the equivalent of clicking the History button.
Alt+J		Minimizes the dashboard . It also expands the dashboard, if it was minimized.
Alt+K		Transfers you to To Do Entry - Main for the current item in the To Do list. Pressing Alt+K while on this page takes you to the first To Do entry on the list.
Ctrl+Alt+L		Logs the user out of the system. This is equivalent to clicking the Logout menu item.
Ctrl+Alt+M		Opens the Menu button.
Alt+N	Alt+Shift+N	Displays the next entry in the search results. Equivalent to clicking Next Item .
Alt+O	Alt+Shift+O	Navigates to the user's home page. Equivalent to clicking the Home button.
Ctrl+Alt+P		Opens your User Preferences page. Equivalent to clicking the Preferences menu item.
Alt+P		Displays the previous entry in the search results. Equivalent to clicking Previous Item .
Ctrl+Alt+Q		Opens the business process assistant , where you can select a script that walks you through a business process.
Alt+R Ctrl+Alt+R		Refreshes the page with the last saved version of the object being displayed in the page. Equivalent to clicking the Refresh button.
Alt+S	Alt+Shift+S	Saves any changes that you've made on the database. Equivalent to clicking the Save button.
Ctrl+Shift+S		Opens the business process assistant , where you can select a script that walks you through a business process.
Alt+X	Alt+Shift+X	Opens the To Do Summary .
Ctrl+Shift+X, Ctrl+Shift+Z		If the Script area is open, equivalent to clicking the Close button in the script area. If a pop-up window is open, this closes the window.

Accelerator Key	Potential Alternate Key (see Note above)	Function
Alt+Y		Displays the previous item in the To Do list. Equivalent to clicking the Previous To Do button in the Current To Do zone.
Alt+Z	Alt+Shift+Z	Displays the next item in the To Do list. Equivalent to clicking the Next To Do button in the Current To Do zone.
Alt+Insert		Inserts a new row into an editable grid if the insertion point is in the grid. Equivalent to clicking the Insert button.
Alt+Delete		Removes a row from an editable grid if the insertion point is in the grid. Equivalent to clicking the Delete button.
Alt+F1		Invokes online Help for the application. Equivalent to clicking the Help button.
Alt+1 ...Alt+9		These keys are only applicable when a page / portal with multiple tabs is open. The shortcuts open the tab page whose relative position corresponds with the shortcut number. For example, if you press Alt+3 on a page with multiple tabs, the third tab is displayed. If there are more than nine tabs on a page, you must use the mouse or F2/Shift+F2 to display a tab. Note that you cannot use the numeric keypad for these shortcuts.
Ctrl+1 ...Ctrl+9		Invokes the favorite link whose relative position corresponds with the shortcut number. For example, if you press Ctrl+3, the transaction for your third "favorite link" is invoked. Note that you cannot use the numeric keypad for these shortcuts.
Ctrl+Shift+1 ... Ctrl+Shift+9		Invokes the favorite script whose relative position corresponds with the shortcut number. For example, pressing Ctrl+Shift+2 selects the second of your "favorite scripts". The numeric keypad cannot be used for these shortcuts.
Alt+left arrow		Populates a scroll area with the previous entry. Equivalent to clicking the scroll control  .
Alt+right arrow		Populates a scroll area with the next entry. Equivalent to clicking the scroll control  .
Enter when in a field with an adjacent Search button		Invokes the search. Equivalent to clicking the Search Button .
Enter when a row in a search result grid is highlighted		Selects the item and closes the search. Note that after selecting an item from the search results, you can press Alt+N or Alt+P to scroll up or down through other items in the result grid without returning to the search results.
Enter when a menu item is highlighted		Selects the menu item.
Enter when the insertion point is within a tree node		Equivalent to clicking on the item in the tree: <ul style="list-style-type: none"> • If the insertion point is on a node image, the node is either expanded or collapsed, dependant on its current state. • If the insertion point is on a context menu, the context menu is opened and normal menu processing applies. • If the insertion point is on a "link" element, normal navigation processing occurs.

Accelerator Key	Potential Alternate Key (see Note above) Function
Tab when a row in a search result is highlighted	Highlights the next row.
Shift + Tab when a row in a search result is highlighted	Highlights the previous row.
Space when the insertion point is on a button	Equivalent to pressing the button.
Space when the insertion point is in a date/time field	Opens the date/time selection window.
Space when the insertion point is in a check box	Toggles the check box between checked and unchecked.
Page Up while in a list grid	Scrolls the contents of the list grid up approximately one page.
Page Down while in a list grid	Scrolls the contents of the list grid down approximately one page.
Up or Down Arrow when in a dropdown list	Moves to the next or previous value in the list. You can also press a letter key to locate the next value starting with that letter.
Up or Down Arrow keys when in search results	Highlights the next/previous row in search results.
Up, Down, Left or Right Arrow when using a menu	Highlights other menu items.
Esc when using a menu	Closes the menu.
F2	Opens the next tab page when viewing a page / portal with multiple tabs.
Shift+F2	Opens the previous tab page when viewing a page / portal with multiple tabs.

Activities

About Activities

An activity is a task to be performed by a crew at a particular geographic location within a time period.

The system prioritizes the activity relative to other activities, schedules it to appropriate resources in a timely manner, and reports back completion information to the requesting host system.

Details provided with the activity control key aspects of how it is managed, such as

- The type of work to be performed
- Where work is to be performed?
- When work should be performed?
- How long is it estimated to take?
- Who may or may not be assigned to perform it. The activity may pose restrictions around the crew that may be considered to perform the work. For example the assigned crew may need to possess certain skills or equipment in order to carry it out.
- And much more

This section describes concepts related to managing activities.

Activity Type Controls Everything

Each activity references an activity type that describes the type of work to be performed as well as many settings that control how activities of that type should be scheduled, prioritized, dispatched, worked and a lot more.

Refer to Understanding Activity Types for more information.

An Activity Is Associated With One Host System

An activity is requested by a host system. All communications related to an activity are exchanged solely with the host system that requested it, which is defined on the activity's type.

Refer to The Activity's Host System for more information.

The host system sends in the activity along with the activity's [unique identifier](#) in that system. The host identifier of an activity is required for tracking purposes as well as for all communications from and to the host system related to the requested activity.

There are situations where activities are initiated within the system. For example, crews may use their mobile device to [initiate an activity in the field](#) and dispatchers may create an activity to [assist](#) another crew with their activity in the field. The host system has to be aware of all work in the service territory, including activities it did not initiate. In these situations, the system notifies the host system defined on the [activity type](#) of the activity that is being created within the system (the assist or activity created in the field). In response, the host system creates a corresponding activity in that system and reports back its unique host identifier for further tracking and communications. This ensures that regardless of where an activity was initiated, all messages such as status and completion information are sent to the appropriate host system.

Activity Workflow

Activity states are depicted in the following diagram:

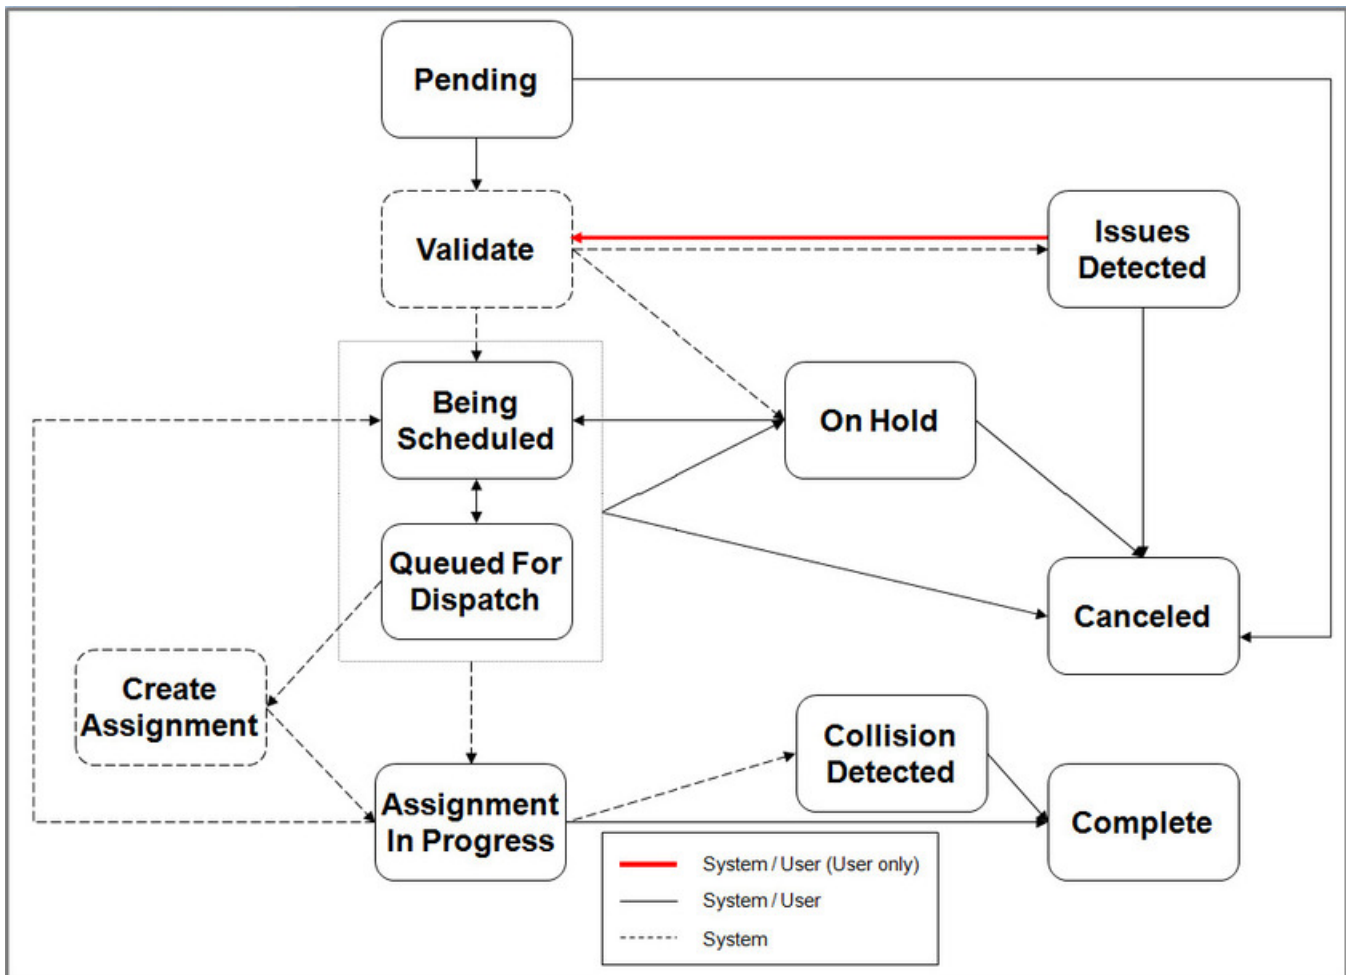


Figure 3: Activity Lifecycle

This section describes key stages in the overall lifecycle of an activity.

Refer to [About Complex Activities](#) for more information about complex activities and their unique workflow.

Validation

When an activity is received from the host system (or created internally), it is validated to ensure it has all the information needed for properly handling it. Basic validation ensures the system accepts and sends to the scheduler well formed activities.

If the activity is valid, the activity is automatically set to the **Being Scheduled** state and made available for [scheduling](#). If critical errors are found, the activity is rejected and not saved at all.

The system will accept the activity if all the errors are such that they can be corrected by a user within the application. In this situation, the activity is saved in an **Issues Detected** state. The activity stays in this state until all errors are corrected. Refer to the [Issues Detected](#) state for more information.

Once all errors are fixed by a user, the activity can now be [scheduled](#) and is automatically set to the **Being Scheduled** state.

Custom rules may be implemented to automatically place a valid activity [on hold](#) for various business reasons instead of directing it to be scheduled.

Issues Detected

When an activity is added and does not have all the necessary details to be properly scheduled, it is set to this state along with a list of the issues detected.

NOTE: If critical errors are found, the activity is rejected and not added at all. Only if the validation algorithms determine that the issues can be resolved by a user within the application, the activity is allowed to be added in this state.

A To Do Entry is raised to notify a user of the situation. When the user corrects the information, the activity is [validated](#) again.

Being Scheduled

An activity in this state is available for scheduling but its schedule is not finalized yet.

The activity typically starts in this state when it is initially received and dispatched to a specific crew (or contractor company) to be worked and then proceeds along its lifecycle towards completion. However, there are times when the assigned resource declines the work and the activity is set back to this state to be rescheduled.

Refer to [Scheduling An Activity](#) for more information about this process.

While the activity is in the **Being Scheduled** state

- The scheduler keeps improving on its scheduling decisions. During this continuous optimization phase, the activity's schedule may frequently change between being assigned, reassigned to another shift (or contracted capacity), its order changed within its currently assigned shift's route, as well as become unassigned again. This goes on until it is time to [queue the activity for dispatch](#) so it can be communicated and dispatched to the assigned crew (or contracting company).
- The activity's time windows are checked and if the activity has expired, it may be automatically extended or [canceled](#) as per configuration on the activity type. Refer to [Expired Activities](#) for more information.
- The activity's time windows are recalculated to take into account new or canceled applicable override time window events, if any exist. Refer to [About Override Time Windows](#) for more information.
- The activity's preference ranking of contractor companies are kept up to date. Refer to [Contractor Eligibility Rules](#) for more information.
- If the activity type is configured to advance allocate work of this type, then such activity may be allocated to the currently assigned crew shift some time prior to the shift's planned start. Refer to [Advance Allocation](#) for more information.
- Business rules may put the activity [on hold](#).

NOTE: Many of these business rules rely on the task being periodically monitored by the **Task Monitor** batch process. Refer to [M1-TSKTR](#) batch control for more information.

Queued For Dispatch

Depending on shift configuration, when the crew starts their shift, they may be dispatched their entire schedule for the day or only the first few tasks. Shift configuration may also allow the crew to preview a few of these tasks in advance, prior to starting the shift.

Refer to [Locking The Schedule](#) for more information.

When it is time to make the activity available for preview or dispatch it to the crew's mobile device, the system queues it for dispatch so that its schedule is locked and can no longer change once communicated to the crew.

A crew's work schedule is communicated to them in the order they should work their tasks. Therefore activities are queued for dispatch in that order.

Setting the activity to the **Queued For Dispatch** state is the first step in the dispatch process. The scheduler adds the activity to the end of the dispatch queue and stops further optimization of its schedule. The activity retains its current schedule from this point on, unless the activity is explicitly [rescheduled](#).

The following business events queue activities for dispatch:

- If a crew's shift is configured to support [Advance Dispatching](#) the configured number of first tasks they may preview are queued for dispatch at a predefined time prior to the shift's planned start time.

- When a crew starts their shift, the scheduler automatically queues for dispatch (if not already queued in advance) their first few scheduled tasks, all of them or none, as per the shift's drip mode configuration. If a drip horizon is used, each time the crew starts or completes a task the scheduler proactively queues the next one for dispatch. Refer to [Controlling Auto-Dispatching](#) for more information.
- All of the activities assigned to a contractor's capacity are queued for dispatch and committed to the contractor company at the predefined notification time before the capacity's start time. Refer to [Advance Notification](#) for more information.
- A dispatcher may manually mark an activity as ready for [dispatch](#). For example, the dispatcher may decide that the crew should perform the activity earlier than originally planned by the scheduler for some reason.

NOTE: Queuing an activity for dispatch should not be confused with the actual process of dispatching the activity to the crew's mobile device (or communicating it to the contractor company). These are consecutive but separate steps.

Certain conditions have to be met before the activity is actually dispatched. The activity remains in the queued for dispatch state until the following conditions are met to proceed with the actual dispatch.

- If the activity is assigned to a crew's shift the following conditions apply:
- The activity cannot be dispatched until the scheduler locks the activity's sequence within the route, i.e. stops further optimization of the activity's schedule, so that it can be communicated to the crew.
- This condition is only applicable in situations when a user has [manually dispatched](#) an activity and the schedule has not yet processed that request. Once the scheduler processes the request it locks the activity's schedule.
- If the activity was auto dispatched, i.e. placed in the queue by the scheduler itself, this condition is always met since the scheduler always locks the activity's schedule at the same time.
- Actual dispatching can only happen once the crew has started their shift and is ready to take on work.
- If the activity was placed in the queue in advance, prior to the crew starting their shift, then it would wait in this state until the shift starts.
- If the activity is queued after the shift has already started, then this condition is always met.
- If the activity is assigned to a contractor capacity, it is communicated to the contractor (dispatched) right away as it enters the queue with no additional wait time. Refer to [Committing An Activity To A Contractor Company](#) for more information.

NOTE: When conditions are met, the activity would be automatically dispatched without relying on any background process to move it along. Even when the activity is advance dispatched, it would automatically be dispatched when the shift starts.

When ready to be dispatched, the system creates an assignment for the activity and its currently assigned crew shift (or contracted capacity), and sets the activity to the **Assignment In Progress** state. At this point the crew (or contractor company) is in charge of [working the activity](#).

Refer to [The Dispatch Process](#) to better understand the big picture of the dispatch process.

Assignment In Progress

When it is time for an activity to be dispatched to the assigned crew (or capacity contractor), the system creates an assignment for the activity and the assigned shift (or contracted capacity), and sets the activity to the **Assignment In Progress** state.

An assignment is a copy of the activity information created by the system to track the assigned resource's progress in accomplishing the work. The assignment is created on the server and a copy of it is dispatched (or sent) to the assigned resource. All updates provided by the assigned resource are communicated via their copy of the assignment.

Refer to [About Assignments](#) for more information.

The activity is considered dispatched when the assigned resource acknowledges receiving their copy of the assignment and thus accepting the work.

Refer to [The Dispatch Process](#) to better understand the big picture of this process.

While the activity is in this state the ownership of the work is with the assigned resource and the system waits for their resolution. Either the resource performs the work and completes their assignment, or declines it and returns their assignment.

NOTE: If due to dispatcher interventions and [recall](#) attempts the activity was sent to multiple crews the system would wait for the resolution from all these assignments before it can assess the activity's state.

Possible resolutions may be as follows"

- If the resource (crew or contractor company) performs the work and completes their assignment, the activity is automatically [completed](#). The activity is transitioned to the **Completed** state and the completion information is sent to the requesting host system.
- If the resource declines the work and returns their assignment, the activity automatically transitions back to the **Being Scheduled** state to be [rescheduled](#).
- If due to dispatcher interventions and recall attempts more than one crew have actually indicated that they completed their assignment, the system has no way of knowing which crew completion information to use. The activity is set to a **Collision Detected** state and a To Do Entry is raised to notify a user of the [collision](#) situation. A dispatcher should review the information and manually choose one of the completed assignments to be used to complete the activity.

Completed

The activity is set to the **Completed** state when:

- The crew (or contractor company) performs the work and completes their activity assignment along with all the necessary completion information. Refer to [Completion Information](#) for more information.
- The crew completes their activity assignment but marks it as incomplete because actual work could not be performed; for example, when details on the activity are invalid or there was some problem getting access to the work site. A completion message is sent to the host indicating that the system will no longer attempt to perform the work. This allows the host system to cancel their activity, or correct and resend it as needed. Refer to [Completion Information](#) for more information.
- The host notifies the system that the activity has been completed in the host system. This is not a typical situation but there are business practices where the crew may prefer to key in the completion information directly in the host system application. This may happen when all communication with the crew is handled by the host system and not by this system. Refer to [Crew May Be Managed Externally](#) for more information. In either situation, since the host has the completion information already, there is no need to send them a completion message.

Collision Detected

Typically only one crew completes the work requested by an activity and the completion information they have entered on their assignment is sent to the host system. However, due to user intervention and recalls, the activity may be assigned concurrently to more than one crew.

If for some reason more than one of the concurrent crews assigned to work the activity claim they have worked it then the system has now multiple sources for completion information and has no way to determine which to forward to the host system. This collision situation requires user intervention.

While in the [Assignment In Progress](#) state waiting for open assignments to be finalized, if the system detects more than one completed assignment for the activity, it sets the activity to this state and raises a To Do entry for a user to review the assignments and manually choose the one to use to complete the activity.

Refer to [Concurrent Assignments](#) for more information on this situation.

Canceled

Canceling an activity is typically initiated by the host system for the purpose of preventing the work from being performed for some reason. The system can only confirm a cancellation when the activity is not yet dispatched. Once dispatched, the system attempts to cancel the activity but only the crew has the latest information as to its current state.

Refer to [Activity Cancellation By Host](#) for more information about how such a request is handled.

[Expired activities](#) may also be automatically canceled based on the activity type configuration.

On Hold

There are situations where the activity is not yet ready to be scheduled as it waits for some business event to occur or for proper timing, etc. An activity that is placed on hold is not scheduled; it is not sent to the scheduler and as such, would not be automatically assigned a resource to perform it.

An activity is automatically placed on hold when:

- The activity is [created in the field](#) and the crew indicated they intend to work it themselves. In this situation there is no need to request the scheduler to schedule it as the crew intends to complete the work themselves. If the crew eventually does not perform the work, the activity no longer stays on hold and is sent back for [scheduling](#).
- The activity is [created in the field](#) and the crew does not plan to work it but the activity requires an appointment to be made. The system sends the new activity to the host system and waits for the host system to arrange an appointment with the customer. When the host system sends in the confirmed appointment, the activity no longer stays on hold and is sent back for [scheduling](#).
- Custom business rules are added to prevent the activity from being scheduled.
- A dispatcher may manually place an activity on hold as needed. When the hold is no longer needed, the user can manually transition the activity to be [scheduled](#).

Activity Identifiers

This section describes the various identifiers an activity may be associated with, including its internal unique activity identifier.

Host External Identifier

The host external identifier of an activity uniquely identifies a non canceled activity in the system by its identifier in the host system.

The identifier is provided by the host system when the activity is interfaced into the system and is used thereafter to identify subsequent host requests to update the activity's details or query its current status. Similarly, the system includes this identifier in all interim and final completion messages related to the requested activity that are reported back to the host system.

If the host system cancels an activity they have previously requested, the same host external identifier may be used to create a new activity in the system.

Host Alternate Identifier

The host may associate an activity with an additional non unique identifier for reporting purposes only. The system does not use this alternate identifier for any processing but also includes it on messages sent back to the host system.

Alternate Identifier

A short identifier provided by a dispatcher to easily communicate an activity to a crew in situations where that activity failed to be sent to the mobile device.

Refer to [Field-Referenced Activities \(FRAs\)](#) for more information.

Chain Related Identifiers

When the host sends in a set of activities in a [chain](#) dependency, each activity is associated with the host identifier of the entire chain. The system assigns a corresponding internal system generated identifier of the chain and posts that identifier on all activities in the chain.

Refer to [Chain](#) for more information.

Job Related Identifiers

When the host sends in a set of activities in a job dependency, each activity is associated with the host identifier of the entire job. The system assigns a corresponding internal system generated job identifier for each host job identifier and posts that identifier on all the activities in the job.

Refer to [Job](#) for more information.

Location

This section describes various aspects related to the location of where the work is to be performed.

Site Address

An activity must specify the address of the actual entrance to the site of service. This may be different than the [service address](#), where the work is performed.

The scheduler uses the site address to plan for the crew's travel time to and from the activity's location on the scheduled shift. It may use more or less travel time based on the [relative speed](#) of the vehicle(s) used by the crew on that shift.

The scheduler does not use the address itself but rather its geocoded value. Without the geocode of its site address an activity cannot be scheduled. If a geocode is not provided when the activity is created or when the address is changed the system attempts to geocode the address using the [geocoding service](#) defined on the installation options expecting an exact match.

NOTE: Logic to geocode the site address resides in algorithms associated with the activity's business object. Review these algorithms and adjust their match quality parameter as needed.

Site Identifier

The scheduler is encouraged to assign activities at the [same site](#) to the same crew shift. An activity is considered as being in the same site as another activity if they share the same site identifier.

In some cases, it is desirable for activities to be scheduled to the same crew when the activities are in the same area (site), but have different addresses (locations). For example, a campus, an apartment complex, or an office park would have various units with different addresses but are in close proximity.

The host system may send in a site identifier along with the activity details to group these activities into one site. If not explicitly provided, the system composes an activity's site identifier by concatenating a predefined list of the activity's site address constituents. By default all site address constituents are used.

Refer to [Site ID Configuration](#) for more information.

Site Delay

Site delay instructs the scheduler to factor in additional time (the amount of the site delay) immediately after the activity is planned to start, when planning the route of the assigned crew shift. For example, a crew may need time to park or gain access to a locked building after they travel to the site, but before they actually start work. Site delay is factored once for all activities [scheduled at the same site](#).

NOTE: Site delay cannot be used for a [complex activity](#).

Scheduling Activities At Same Site

When scheduled to a crew shift, same site activities are scheduled one after the other, with the scheduler considering the arrival, start, and completion times for each individual activity, but incorporates the site delay time, if specified, only once. Travel time between locations is calculated as usual for same site activities in different locations. Breaks can be scheduled between activities at the same site.

Service Address

By default the address for which service is provided is assumed to be the same as the site address. However there are situations where they may differ. For example, a large apartment building may have all meters located in a single location in the building and not inside the individual apartments. The crew would be directed to the meters' site address whereas the address of the individual apartment for which the work was requested would be provided as the service address for information only purposes.

Time Zone

The activity's time zone is the time zone of its geographic location, i.e. where work needs to be performed. All date and time information on the activity and its assignment portals is entered and displayed in that time zone.

Service Area

An activity is associated with a single service area that contains its location. The host system may send in the service or let the system default it based on the activity's address and type of service.

NOTE: Logic to determine the service area of an activity resides in a "Determine Service Area" algorithm on the activity's business object. Click [here](#) to see the algorithm types available for this system event.

Refer to [Understanding Service Areas](#) for more information.

Service Class

An activity's type may be associated with one or more service classes. A service class represents a broad category for work that may be used for various purposes.

Refer to [Designing Activity Types](#) for more information.

Estimated Work Duration

An activity specifies the estimated amount of time it should take a crew to perform the requested work. If not explicitly specified, this information is defaulted from the activity type.

The scheduler uses the estimated duration to book enough time on the assigned crew's shift for them to complete the work. The scheduler may allot more or less time based on the [relative efficiency](#) of the crew on that shift. Skill efficiency may also affect the scheduled time when applicable. Refer to [Skill Efficiency](#) for more information.

The nature of the work requested by a standard activity is such that its duration is short enough to be completed by a crew on a single shift. Longer work that requires multiple visits by one or more crews over several shifts is referred to as a complex activity.

Refer to [About Complex Activities](#) to better understand how complex activities differ from standard activities.

Time Windows

This section describes the various time windows associated with an activity.

At Least One Effective Time Window (ETW) Is Required

The host system requesting the activity must specify at least one time window during which the activity is in effect and may be scheduled. More than one ETW may be specified to state additional time periods during which the activity can be worked.

The earliest time window's start time defines the earliest time the activity can be scheduled and the latest time window's end time represents its [due date](#).

If the activity is an appointment it may have one ETW only reflecting the booked [appointment](#) time window.

If the activity is not worked by the end of the latest ETW it is considered [expired](#).

Due Date

There are times when the host system only states the earliest time an activity may be worked and relies on the system to automatically default when it is due. In this situation a single open ended ETW is sent with the activity and the system defaults its due date.

Activity type configuration controls the amount of time which activities of this type should typically be completed. If the activity type lacks the proper set up to support such defaulting, the activity is rejected.

Refer to [Default Due Date](#) for more information.

Multiple ETWs

By default all [ETWs](#) provided for an activity are equally preferable for scheduling. The host may specify a window cost with each ETW to apply a cost of using the time window relative to other time windows, thus establishing preference between the time windows.

Plan To Arrive Or To Complete?

By default an activity is scheduled such that the crew is planned to arrive at the activity location and start the work during an [ETW](#). The system does not necessarily account for the duration of the work to also be within the time window.

When it is critical that an activity is planned to start and complete within its [ETW](#) the host should send the activity's ETWs stating that completion time should be used as their Time Window Mode.

Preferred Time Windows (PTW)

The host system requesting the activity may optionally provide one or more preferred time periods when it is ideal to schedule an activity.

A preferred time window may not start earlier than the earliest [ETW](#) nor end after the latest. If specified, the scheduler is encouraged to prefer scheduling the activity during those times, and applies added cost when scheduling an activity outside of its PTW. The cost is defined on the activity type as the Preferred Time Factor.

A [PTW](#) may specify a service level agreement(SLA) flexibility time in minutes to define time before and after the time window within which an activity avoids incurring added cost.

For example, if the SLA Flexibility is set to 20 minutes, then scheduled arrival must be no more than 20 minutes before the start of this preferred time window or 20 minutes after the end of this window; otherwise, the Preferred Time Factor cost is applied.

When Can The Activity Be Actually Scheduled?

The host system provides the overall time frames during which the activity may be scheduled in the form of one or more [Effective Time Windows \(ETW\)](#).

The following restrictions may prevent the activity from being scheduled during some part of these periods:

- Certain types of work may not be scheduled on [public holidays](#).
- Certain types of work may only be performed on specific days of the week and within certain business hours. Refer to [work profile](#) for more information.
- Due to some upcoming event, for example an extreme temperature forecast, the business is required to prevent certain type of work within certain areas from being performed throughout the duration of the event. For example, there may be a government mandate that electricity meters cannot be cut off during days of extreme heat or extreme cold. Refer to [Override Time Windows](#) for more information on such events.

Due to such restrictions, the system has to identify portions of each ETW during which the crew can actually arrive to work the activity. These are known as [Arrival Time Windows \(ATW\)](#). When no restrictions apply, an ATW covers the entire ETW.

In the same way the system has to identify portions of each ATW that overlap with any [Preferred Time Windows \(PTW\)](#) requested. These are known as [Service Time Windows \(STW\)](#).

ATWs and STWs are automatically generated and stored along with the activity when the activity is created and are regenerated when time information or restrictions change.

The scheduler only considers an activity's arrival time windows and service time windows.

To visualize when an activity can be scheduled, the Schedule tab on the Crew Shift screen depicts all of the activity's time windows in a Gantt view.

NOTE: Logic to calculate an activity's arrival and service time windows resides in a "Update Time Windows" algorithm on the activity's business object. Click [here](#) to see the algorithm types available for this system event.

Arrival Time Window (ATW)

An ATW is a generated time window representing a period of time within an activity's ETW during which the activity is allowed to be scheduled, based on applicable restrictions.

The scheduler never schedules an activity earlier than the start time of an ATW.

The scheduler can however, schedule an activity with a late arrival time (outside of an ATW), though it is not likely. Unlike a too early schedule, lateness is cost driven and not a hard limit. The scheduler's cost control parameters are typically configured to discourage lateness globally.

An activity may further specify a cost of arriving late to this specific activity relative to other activities, to further discourage planned late arrival time. For example, if an activity has a relative late cost of 3 then the cost of late arrival for this particular activity is three times more than the late cost defined globally.

Service Time Window (STW)

An STW is a generated time window record representing time within an activity's [arrival time window](#) during which it is ideal to schedule the activity. An STW is a time overlap between an arrival time window and a preferred time window.

Appointments

An activity may represent an appointment made with a customer and therefore needs to be scheduled within the agreed-upon appointment window. Such activity is sent in from the host system with an indication stating it is an appointment along with a single [Effective Time Window \(ETW\)](#) set to the appointment window selected by the customer.

The following sections describe topics related to appointment activities.

Booking An Appointment

Before sending in an appointment activity, the host system calls the [appointment booking service](#) to request a list of available appointment windows for that activity. The list is presented within the host application and the customer chooses one of the available time slots. The host system then sends in the activity marked as an appointment along with the selected time window.

By default, time slots offered by the appointment booking service are those defined by the [default appointment booking group](#) referenced on the type of activity being requested. The host may override this and request to use another appointment booking group.

If the activity being booked is part of a [job](#), the request must include all of the activities in the job.

If the activity lists multiple depots, the scheduler offers appointments for crews associated with the closet depot only.

Conditional Booking

By the time the user makes a selection for a specific appointment window there is a chance that the selected window may no longer be available.

Therefore the host system sends in the activity with the selected time window but marks it a "conditional booking" pending the availability of the selected appointment window.

If the time window is still available, the activity is created as usual. However, if it is no longer available, then the activity is not created and the caller is notified accordingly.

Re-Booking

The system allows users to change an existing appointment in the host application. When the user requests a change to the appointment, similar to requesting an appointment for the first time, the host system works with the appointment booking service to present the user with a new set of available appointment windows to match the new booking details.

Typically the request to rebook an existing appointment is sent from the host as conditional upon the ability to accommodate the new booking instead of the original booking. As such, if the new booking is no longer available or the old booking can no longer be canceled, the rebooking attempt fails and the user is notified of the situation. Once rebooking is confirmed, the system cancels the original booking and schedules the new one instead.

If the original booking is already dispatched to a crew, the system cannot guarantee its successful recall from the crew as this requires human intervention. If the rebooking request states that it must be fully guaranteed, the rebooking attempt fails and the user is notified of the reason. If however the rebooking request can tolerate it not being guaranteed, then the system alerts a dispatcher to follow up with the crew on the cancellation of the original appointment but at the same time also schedules the new booking. If the dispatcher manages to cancel the original appointment in time, then it is canceled. If it is not canceled in time, both the original and the new booking are scheduled temporarily until the dispatcher cancels the new one.

If the activity was manually allocated to a specific shift or locked to its depot run by a dispatcher, a subsequent rebooking change frees the activity from these allowing it to be fully rescheduled.

When rebooking activities that are part of a job, the new job composition must fully replace the old job composition so that any activity that used to be in the old job and is no longer part of the new one needs to be explicitly listed as being canceled.

Reserve Shift Time For Appointments

In order to accommodate customer appointments up until the day before, crew shifts can be configured to reserve a certain percentage of their time for appointments so that other low priority work does not fill up the shifts.

While the scheduler can be configured to [displace low priority work](#) and give precedence to appointment work, last minute changes in the schedule may cause unnecessary disruption that can be prevented by setting appropriate shift reserves for appointment work.

Refer to [Reserve Capacity](#) for more information.

Displace Low Priority Work

The scheduler can be configured to ignore low priority activities currently scheduled to shifts so that new appointments can be offered to customers in their place.

Refer to [Activity Priority](#) for more information.

The scheduler may displace work that:

- Is non appointment work with a scheduling priority lower than a specific threshold. The appointment booking request as well as the conditional booking request may explicitly specify this threshold but if not specified, the threshold is defaulted from a corresponding scheduler parameter.
- Has enough remaining days to be completed before it expires. This is defined as a scheduler parameter.

Emergency Activities

Emergency activities must be scheduled and worked as soon as possible. They typically represent a hazardous situation that has to be addressed immediately. As such, an emergency activity is scheduled to an available crew at all costs and dispatched to their mobile device using an emergency queue that bypasses any other non emergency task.

The following sections describe topics related to emergency activities.

Emergency Queue

If the activity's queue is set to emergency, the activity is considered an emergency activity. This information is defaulted from the activity type but may be overridden at the activity level. This allows activities of non emergency type to be considered an emergency for any business reason, and vice versa.

The mobile device polls for messages placed in the emergency queue before any in the normal queue. All messages sent on behalf of an emergency activity are placed in the emergency queue so they would reach the device prior to any non emergency queued messages.

Emergencies Are Scheduled Differently

The scheduler handles emergency activities differently:

- They do not have a [scheduling priority](#) as they are more important than the highest priority work.
- The time window specified on the activity is irrelevant.
- The scheduler may be allowed by configuration to assign an emergency activity to a crew that is already working another task if they can address the situation faster than any other crew.
- The scheduler ignores crew shift logon time reservation when scheduling an emergency activity to this shift allowing the crew to address the emergency as soon as possible.

NOTE: Emergency activities are not eligible to be scheduled to capacity contractors. They may only be scheduled to internal crews.

While emergencies take precedence over non emergency work, if there is a lot of emergency work, unnecessary same day changes and disruption to crews' schedules can be avoided by reserving capacity on shifts for emergency work.

Refer to [Reserve Capacity](#) for more information.

Unacknowledged Emergency Alert

The system may be configured to raise an alert when a crew fails to acknowledge taking on an emergency activity. In these situations every minute counts. Instead of waiting for the alert to be raised by the alert type monitor background process, the

system generates the alert in advance, right when the activity enters the system. However, the alert is set to only take effect some number of minutes in the future, allowing some time for the crew to acknowledge the activity (which also closes the alert). If the crew fails to acknowledge the activity before the alert's effective time, the alert appears on a dispatcher's alert queue ready to be addressed in a timely manner.

NOTE: Raising such an alert is optional. An alert is only created if an active alert type is configured and referenced as the **Unacknowledged Emergency Alert Type** [global configuration](#) parameter.

Refer to [About Alerts](#) for more information.

Activity Priority

The scheduler prioritizes non emergency activities based on their scheduling priority.

Scheduling priority is a numerical value from 0.01 to 10 that establishes the importance of the activity relative to other activities, where 10 is the highest priority. Activities with a higher priority are scheduled versus not scheduled, scheduled to an earlier shift, or scheduled earlier in the shift.

For example, if there are two connect activities, one at a hospital with a higher priority and one at an seldom used warehouse with a lower priority, the higher priority hospital activity will get scheduled today and earlier in the day, whereas the lower priority one may not get scheduled on that day at all.

The following sections describe further aspects of activity prioritization.

Scheduling Priority Over Time

An activity's scheduling priority is typically [defaulted](#) by the system when the activity is created but can be adjusted based on business rules.

The scheduler applies a [factor](#) onto the activity's scheduling priority to calculate its priority over time. It is this calculated scheduling priority over time, and not the initial value, that is used by the scheduler to prioritize work. The calculated scheduling priority is known to the scheduler only.

Default Scheduling Priority

An activity may optionally capture a business priority number representing its business value in terms of cost, other currency amount, number of customers impacted, relative customer importance, etc.

The system uses the activity's type configuration as well as that business priority number, if provided on the activity, to determine the activity's scheduling priority based on real business numbers.

Refer to [Determine Activity Priority](#) to better understand how the system sets the default scheduling priority for an activity.

Though not typical, the host system may send in the scheduling priority along with the activity in which case the system would simply use the provided value, ignoring the business priority number if also provided. Note that the scheduler only acts upon activity's scheduling priority. When an activity is created, the priority is defaulted from the activity type unless it is specified by the host. However, the priority on the activity itself is what the scheduler ultimately considers.

This also applies for appointment booking requests. If both values are included, the scheduling priority overrides the business priority number, otherwise the scheduling priority is defaulted as described.

Priority Profile

The scheduler applies a factor onto the activity's scheduling priority to calculate its priority over time. The factor is defined as a function over the activity's broad time window to make it stronger, weaker, or unchanged as needed.

For example, short-cycle, SLA-oriented work is critical at the beginning of the window and diminishes after that. Maintenance work, such as meter inspections, is low priority in the beginning of the window and becomes critical toward the end.

The activity's type references the Priority Profile that defines the function appropriate for activities of this type.

Refer to [Defining Priority Profiles](#) for more information.

Mandatory vs. Optional Work

An activity states whether or not it is mandatory to schedule it by its due date regardless of cost. If not mandatory, scheduling the activity is considered optional and is driven by cost only. As such, optional work may eventually remain not worked by its due date and expires.

Emergency activities are always considered mandatory.

Mandatory work is treated optional until the activity's end time window is within the scheduler horizon, except if the activity is a single day activity. Mandatory single day activities are never treated as optional.

This indication is defined on the [priority profile](#) referenced on the activity's type but can be overridden on the activity.

Resource Capabilities

An activity may require the mobile workers of the assigned crew to have certain work related capabilities in order to perform the work. These may be skills or activity specific requirements such as the need for security clearance, building access, drug testing, permits, language proficiency, etc.

In the same way, an activity may require the assigned crew's vehicle to have certain pieces of equipment in order to perform the work. For example, the crew may need a bucket truck to perform line pole work.

The scheduler assigns an activity to a crew that is capable of carrying it out on a specific shift.

Refer to [Crew Capabilities](#) for more information.

The following topics describe how capacities may be used to influence the activity's crew assignment.

Required Capabilities

An activity may require the crew to possess a specific capability, worker or vehicle related.

The scheduler only considers crews that possess this capability on a specific shift when looking to find a compatible crew shift for an activity. If none exists the activity remains not assigned.

If however the activity is manually allocated to a non compatible crew on a specific shift, the scheduler honors that allocation.

Preferred Capabilities

An activity may prefer the crew to possess a specific capability, worker or vehicle related.

The scheduler is encouraged by cost configuration to prefer crews that possess this capability on a specific shift when looking to find a compatible crew shift for an activity.

Being a preference only and not a requirement, the scheduler consideration is cost based. The scheduler may assign the activity to a crew shift that does not possess the preferred capability if none exist or the selected shift is all in all less costly than a shift with the preferred capability.

Not Allowed Capabilities

There are situations where an activity may require that the crew not possess a specific capability, worker or vehicle related.

NOTE: This type of requirement does not support worker skills.

The scheduler only considers crews that do not possess this capability on a specific shift when looking to find a compatible crew shift for an activity. If none exists the activity remains not assigned.

If however the activity is manually allocated to a non compatible crew on a specific shift, the scheduler honors that allocation.

Default Capabilities

Resource capability requirements are typically the same for all activities of the same type and as such are configured on the activity type. By default, an activity's resource capability requirements are those defined on its activity type. A specific activity may [override](#) the default list by either extending it or completely replacing it.

Override Capabilities

An activity states whether it uses the default resource capabilities as defined on the activity type, extends the default list, or completely overrides and replaces the list.

When extending, the resource capabilities specified on the activity are used by the system in addition to those defined on the activity type. For example, if a customer prefers a mobile worker that speaks Spanish, the activity for that customer may be sent in with an indication that resource capabilities should be extended to also include a preference for a "Speaks Spanish" skill. When scheduling the activity, the system considers all resource capability requirements specified on the activity type as well as those on the activity.

When overriding the default list altogether, the system only uses the resource capabilities specified on the activity to schedule the activity, ignoring any defined on the activity type. For example, assume a business is installing various products, yet the skills and equipment needed to install products of category A is different from those needed to install products of category B. Rather than having to set up separate installation activity types by product category, it may be efficient to simply define a generic "Install Product" activity type and specify the capabilities needed to install the specific product based on its category on the individual activity.

Full override can also be used to indicate that the specific activity requires no resource capabilities at all when the activity type specifies some.

Resource Restrictions

An activity may specify certain restrictions that further limit the resources that may be considered for performing the work. The following describes various options affecting the resources the scheduler considers for an activity.

Allowed Crews

Typically the host system relies on the scheduler to find an appropriate crew to work an activity.

There are situations, known to the host system, where an activity may only be worked by a specific list of crews. In this case the activity is sent in with those crews specified as the allowed crews.

If there are no crews listed, all compatible crews are considered by the scheduler as usual. When specified, the scheduler narrows the list of compatible crews to those in this list. If none of the allowed crews are compatible with the activity, the activity remains not assigned.

If however the activity is manually allocated to a crew's shift where the crew is not on that list, the scheduler honors that allocation.

Excluded Mobile Worker

There may be situations, known to the host system, where one or more mobile workers are to be prevented from being assigned an activity. For example, a negative experience a customer had in the past with a specific worker.

When specified, the scheduler excludes from consideration any shift that lists the specified worker(s) as its crew members.

Excluded Shifts And Contractor Capacities

Over time, an activity may accumulate a list of crew shifts (or contractor capacities) it was once assigned to but was explicitly removed from and should not be assigned to going forward.

These are automatically added to the activity by the system when

- A crew on a specific shift declines taking on the activity for various reasons.
- A dispatcher takes action to remove an activity from its currently assigned shift (or contractor capacity).

When scheduling the activity, the scheduler excludes from consideration any shift (or contractor capacity) specified on that list.

Manually allocating the activity back to an excluded shift (or contractor capacity) automatically removes it from the exclusion list.

NOTE: Exclusions may not be sent in along with the activity from the host system.

Allowed Contractors

By default, for activities that are eligible to be worked by contractors, the scheduler may also consider compatible resources from various contractor companies, in addition to internal company crews in the scheduling process.

There are situations, known to the host system, where such activity may only be worked by a specific list of contractor companies. In this case, the activity is sent in referencing those companies in the allowed contractors list. When specified, the scheduler narrows the list of compatible crew shifts (or contractor capacities) associated with the allowed contractor companies. If no compatible resource is found, the activity remains not assigned.

If however the activity is manually allocated to a contractor crew's shift (or a contractor capacity) that is not associated with the allowed contractor companies, the scheduler honors that allocation.

Refer to [About Contractors](#) for more information.

Preferred Resource Groups

An activity that is eligible for contracting may list the preferred order for scheduling the activity between the company and the contractor companies .

Refer to [Contractor Eligibility Rules](#) for more information.

NOTE: Preference is determined and captured periodically while the activity is in the "Being Scheduled" state. Logic to determine these preferences for an activity reside in a "Contractor Eligibility" algorithm on the activity business object. Click [here](#) to see the algorithm types available for this system event. Preference may not be sent in along with the activity from the host system.

Excluded Contractors

Over time, an activity that is eligible for contracting may accumulate a list of contractor capacities it was once committed to but should not be considered again for various reasons.

A contractor company may be excluded from working an activity when

- A capacity contractor company did not approve taking on the activity in a timely manner or they have explicitly declined taking on the activity.
- A dispatcher takes actions to [exclude](#) an activity from its currently assigned contractor company.

When scheduling the activity, the scheduler excludes from consideration any contractor company specified on that list.

Manually allocating the activity back to a crew shift or capacity affiliated with an excluded contractor company automatically removes it from the exclusion list.

Refer to [About Contractors](#) for more information.

Scheduling An Activity

The purpose of the scheduling process is to assign an activity to a compatible crew on a specific shift (or a contracted capacity offered by contractor companies). Refer to [Scheduling Work To Capacities](#) for more information.

To plan the best work schedule and route for each shift, the scheduler matches activity details such as location, duration, the time frame available and capabilities the assigned crew may or not need to have in order to perform the work with information about the location, capabilities and availability of crews. A dispatcher may override scheduler decisions at any time.

When an activity is scheduled to a crew shift it is also assigned a work sequence representing the order in which the crew is expected to work their schedule on that shift.

Common reasons for why an activity remains unscheduled may be

- No scheduler covers the activity's location
- The activity is placed on hold or has data issues preventing the system from sending it to the scheduler
- No compatible resource is available that matches the activity's scheduling requirements
- The activity has expired and the activity type does not allow further time window extensions
- The activity is optional and may be pushed out by higher priority work

Until the activity is queued for dispatch, its schedule is constantly being optimized by the scheduler. The activity may be scheduled to a shift (or a contractor capacity) for a while, rescheduled to another, become free, rescheduled again, etc.

Refer to [Dispatching An Activity](#) for more information.

At any point in time if the scheduler cannot schedule an activity that otherwise should be scheduled, the activity is disabled from scheduling and the scheduling error is reported in the scheduler's registry record for the activity. The system provides information on the activity portal schedule tab to assist you in determining why an activity is not scheduled or disabled. Once the issue is resolved, the activity is automatically rescheduled.

For more persistent errors, the scheduler monitor batch process converts registry errors to To Do entries for a user to follow up. For dispatchers, the [To Dos In Scope](#) dashboard zone lists these entries for activities they are monitoring.

NOTE: The **Scheduler Monitor** batch process needs to run at least once a day to convert scheduling errors to To Do entries.

Expired Activities

If an activity has not been worked by its due date the scheduler stops attempting to schedule it.

Depending on configuration and business practices, activities of certain types may be configured to handle the situation differently (for example, the activity's time window can be automatically extended).

NOTE: The scheduler may still schedule an expired activity if its time windows overlap with a compatible shift. While scheduled to a shift the activity's time windows will not be extended

Refer to [Handling Expired Activities](#) for more information.

Activity Actions

A dispatcher or a business process may perform various actions on an activity that may impact the activity's state and schedule.

The following section describes these actions and the business flows associated with them.

Allocate

A dispatcher or a business process can allocate an activity to a specific crew's shift (or a contractor's capacity). While allocated, the scheduler honors the decision even if it violates resource compatibility rules, and keeps the activity assigned to this shift (or capacity) only. The scheduler may still optimize the order of the activity within the allocated shift.

If the activity is currently dispatched to another shift (or contractor capacity), the dispatcher is notified of the situation as it involves a recall of the activity from the currently assigned resource. Refer to [Forced Recall](#) for more information.

If the activity is [fixed](#) to a different crew than the requested shift's crew, the user is warned that the activity would no longer be fixed to that crew once allocated.

Activities that are allocated to a specific crew shift have a special visual cue when displayed in the CDI Gantt.

NOTE: Logic to handle the business rules associated with manual allocation resides in an "Allocate to Shift" algorithm on the activity's business object. Click [here](#) to see the algorithm types available for this system event.

Allocating an Activity Using "Chooser"

[Allocating](#) an activity using the "Chooser" service allows you to choose the shift (or contractor capacity) from a list recommended by the scheduler. The list is sorted by scheduling cost from lowest to highest, but you can click on any column to sort as needed.

You can relax some of the scheduling restrictions so that more shifts (or contractor capacities) are considered by the scheduler. For example, removing equipment or attribute compatibility allows the scheduler to return shifts that do not possess the worker attributes or equipment that were indicated as required to perform the activity. A maximum number of shifts (or contractor capacities) to return can be specified to prevent the scheduler from showing an unmanageable number of results.

Remove Allocation

A dispatcher or a business process can remove an existing [allocation](#) restriction from an activity making it no longer allocated to a specific shift. If the activity was [locked to a depot run](#), that restriction is also removed.

Removing an allocation restriction does not impact the current schedule of an already dispatched activity. If the activity is ever rescheduled for some reason, the scheduler would consider any compatible crew.

Fix To Crew

A dispatcher or any business process can fix the activity to a specific crew with the purpose of restricting the scheduler to assign the work to shifts of this crew only. In the same way, they may remove such restriction (by leaving the crew blank) and indicate the activity as no longer fixed to any crew.

Locking an activity to a specific crew overrides any other restriction around which crews are [allowed](#) to perform the work, if specified. That means an activity can be fixed to a crew that is not in the allowed list.

Fixing an activity to a crew or removing such restriction does not impact the current schedule of an already dispatched activity. It only impacts the activity if it is ever rescheduled.

If the activity is [allocated](#) to a shift of a different crew, the user is warned that the activity would no longer be allocated to that shift once it is fixed to this crew.

A different flavor of this same functionality allows crews to [keep an activity](#) with them for future shifts.

Activities that are fixed to a crew have a special visual cue when displayed in the CDI Gantt.

NOTE: Logic to handle the business rules associated with fixing an activity to a crew resides in a "Fix To Crew" algorithm on the activity's business object. Click [here](#) to see the algorithm types available for this system event.

Manual Dispatch

Activities are automatically queued for dispatch based on crew shift (or contractor capacity) configuration. The queue represents a list of activities ready to be sent to the crew once they start their shift, if not already started (or communicated to the contractor company at the agreed upon notification time).

A dispatcher may manually place an additional activity in that queue that is otherwise scheduled to a crew's shift (or contractor capacity), but has not yet been placed in the queue by the scheduler. This action is also applicable to an activity that is allocated to a crew shift but the scheduler has not yet scheduled the activity to that shift.

The scheduler responds to a manual dispatch request by adjusting the shift's schedule such that the manually dispatched activity is added last to the queue of already dispatched tasks.

Refer to [queue for dispatch](#) for more information.

Reschedule

This action may be used to reschedule an activity that is currently in the [queued for dispatch](#) state. It effectively reverses a dispatch action taken on an activity, and indicates to the scheduler that the activity is no longer locked to its current location in the schedule.

If the activity is manually allocated, the scheduler cannot schedule the activity to a different shift but can still optimize its location in the schedule. If the activity is not manually allocated, the activity can now be scheduled to any crew shift (or contractor capacity).

Unassign

A dispatcher may unassign an activity from the crew shift (or contractor capacity) it is currently assigned to and send it back to be [rescheduled](#).

In order to prevent the scheduler from assigning the activity back to the same shift (or contractor capacity) from which it was taken, system adds the shift (or contractor capacity) to the activity's list of [excluded shifts and contractor capacities](#).

If the activity is currently manually allocated to that shift (or contractor capacity), the allocation is removed.

If the activity is currently dispatched, the dispatcher is notified of the situation as it involves a recall of the activity from the currently assigned resource. Refer to [Forced Recall](#) for more information.

NOTE: Logic to handle the business rules associated with unassigning an activity resides in a "Unassign Activity" algorithm on the activity's business object. Click [here](#) to see the algorithm types available for this system event.

Defer

A dispatcher may need to put off the scheduling of an activity to a later date. The dispatcher would enter a deferment date and reason. The deferment date must be prior to the end of the effective time window; you cannot defer an activity beyond the end date.

If the activity is currently dispatched, the dispatcher is notified of the situation as it involves a recall of the activity from the currently assigned resource. Refer to [Forced Recall](#) for more information.

To un-defer an activity, the same action is used providing no deferment date.

NOTE: Logic to handle the business rules associated with deferring an activity resides in a "Defer Activity" algorithm on the activity's business object. Click [here](#) to see the algorithm types available for this system event.

Place On Hold

A dispatcher may need to place an activity on hold if something needs to be done before work can be started. Putting a hold on an activity removes it from consideration for scheduling until the hold is manually removed.

Refer to the activity [On Hold](#) state for more information.

Lock Activity To Depot Run

This action is only relevant to a depot related activity.

The system locks a depot related activity to its currently assigned depot task when a user or a process closes the depot task. A depot task represents the scheduled visit to the depot to pickup or drop off goods for one or more activities.

Refer to [About Depot Runs](#) for more information.

NOTE: Logic to handle the business rules associated with locking an activity to the depot run it is currently scheduled to resides in a "Lock Depot Run" algorithm on the activity's business object. Click [here](#) to see the algorithm types available for this system event.

Exclude Contractor

A dispatcher or a business process may exclude the contractor company from working an activity committed to them. When possible, the contractor company is automatically informed of the recall. Refer to [Forced Recall](#) for more information.

The excluded contractor company is added to the activity's excluded contractors list and the scheduler would no longer consider any shift (or contractor capacity) that belongs to that contractor company.

Refer to [Excluded Contractors](#) for more information.

NOTE: Logic to handle the business rules associated with excluding the contractor the activity is currently committed to resides in a "Exclude Contractor" algorithm on the activity's business object. Click [here](#) to see the algorithm types available for this system event.

Revert To Company

This action marks an activity as allowed to be performed only by company resources. It may be used when an activity is eligible for contracting but for some reason the dispatcher needs to revert it back to company only crews.

If the activity is currently dispatched, the dispatcher is notified of the situation as it involves a recall of the activity from the currently assigned resource. Refer to [Forced Recall](#) for more information.

Forced Recall

If the activity has already been dispatched and the requested change should recall it from the currently assigned crew (or contractor company), a warning message notifies the dispatcher that the activity needs to be recalled. The dispatcher may either confirm and proceed with the action or cancel the action.

If confirmed, the system completes the requested action right away. In addition, the system attempts to recall the activity from the currently assigned crew in an effort to avoid as much as possible a situation where more than one crew is sent to perform the same activity. Refer to [Concurrent Assignments](#) for more information.

In the same way, the system sends a message to the external contractor company notifying them of the recall. This is only applicable if the contractor company's system is integrated with your system.

NOTE: The system does not wait for the activity to be actually be recalled to perform the requested action. The recall request may or may not succeed with returning the activity, depending on the assigned resource's actual progress with

working the activity as well as their connectivity. It is the responsibility of the dispatcher to coordinate the recall with the currently assigned resource.

Refer to [Recall](#) for more information on how recalls are handled.

Dispatching An Activity

The following section describes key concepts related to the process of dispatching an activity.

The Dispatch Process

The process of dispatching an activity to the assigned crew (or communicating it to the contractor company) involves multiple steps.

- **Step 1 - Placed in queue.** The activity has to be placed in the queue for dispatch for the dispatch process to begin. This step is typically performed by the scheduler as it auto dispatches work, but other processes may do so as well.
- **Step 2 - Wait.** The activity remains queued until it is time to actually dispatch it.

Refer to the [Queued For Dispatch](#) for more information about the processes involved in these steps.

- **Step 3 - Dispatch.** Once all conditions are met for the activity to be dispatched, the system creates an assignment for activity and the assigned crew (or capacity contractor), and sets the activity to the "Assignment In Progress" state where it waits for the outcome of this assignment. A copy of the assignment is sent to the assigned crew (or capacity contractor).

Refer to the [Assignment In Progress](#) for more information about the processes involved in this step.

- **Step 4 - Acknowledgement.** The activity is considered dispatched if the assigned resource acknowledges receiving their copy of the activity assignment.
- If the activity is assigned to a crew, their assignment needs to be set to the "Dispatched" state. Refer to [Automatic vs. Manual Acknowledgment](#) for more information.
- If the activity is assigned to a contractor capacity, the contractor company has to approve taking on the activity by setting their assignment to the "Approved" state. Refer to [Managing Committed Work](#) for more information.

Refer to the [Assignment Workflow](#) for learn more about how the acknowledgement step.

Disabling Auto Dispatch

Certain types of activities may require an approval from a dispatcher before they can be automatically dispatched by the scheduler to a crew.

By default all activities are allowed to be automatically dispatched by the scheduler. If the activity type explicitly states that auto dispatch is not allowed, the scheduler would still lock the schedule of an activity of this type when it is time to queue the activity for dispatch, but the system would keep the activity in the being scheduled state, waiting for a dispatcher to [manually dispatch](#) it.

Once a dispatcher manually dispatches such activity, it set to queued for dispatch and the [dispatch process](#) progresses as usual.

Refer to [Controlling Auto-Dispatching](#) for more information about this process.

Allow Activities To Break Sequence Locking

When the crew's shift uses standard drip mode, they only have access to their next few tasks on the mobile device and are expected to work the tasks in their route sequence order.

By default, the scheduler locks the sequence of dispatched tasks for shifts configured this way so that the route communicated to the crew remains unchanged. The scheduler would not dispatch a task to be worked prior to another already dispatched task, nor would it switch the order between already dispatched tasks.

You can configure the shift to deviate from this default and allow the scheduler to make such changes to the dispatched portion of the route.

Refer to [Controlling Auto-Dispatching](#) for more information about the various options available on the shift to control dispatched route changes.

There are situations though where even if the shift is configured to not allow route changes, certain activities should be allowed to be dispatched and injected before already dispatched tasks. The indication to ignore sequence locking is set on the activity and is defaulted from the activity type, if not explicitly specified. A dispatcher may directly update an activity to allow this as needed.

If the activity is marked as allowed to ignore sequence locking, the scheduler may dispatch the activity to be worked before another dispatched task as needed.

Completion Information

When an activity is completed by the assigned resource a completion message is sent to the host with all completion information, including any [attachments](#) captured in the field.

NOTE: Logic to compose the completion message for an activity resides in a "Prepare Completion Data" algorithm on the activity's business object. Click [here](#) to see the algorithm types available for this system event.

Activity Cancellation By Host

The host system may send a request to cancel an activity they have previously asked to perform with the intention to prevent the work from being performed.

If the request reaches the system when the activity is already completed then there is no way to cancel the activity and the host is notified.

If the activity has not been dispatched yet, it is [canceled](#) right away.

If the activity is dispatched to a crew (or to an external contractor company) they may have already performed the work in the field and their completion update is yet to be reflected on the server. In this situation, the system cannot cancel the activity without ensuring the crew (or the external contractor company) would indeed not perform the work, for which an alert is raised for a dispatcher to follow up. The immediate response to the cancellation request to the host system indicates that the activity cannot be canceled right away but the request is being forwarded to a dispatcher.

NOTE: The system creates the alert using the active alert type referenced as the **Activity Cancellation Alert Type** [global configuration](#) parameter.

Meanwhile:

- If the crew uses a mobile device, the system automatically sends a [recall](#) message to the crew's device requesting to return the assignment if still possible. Refer to [Recall](#) for more information on how this is processed. In this situation the alert to notify the dispatcher of the cancellation request does not immediately appears on the dispatch's alert queue to allow some time for the recall message to reach the mobile device and return the assignment when possible. If the activity is recalled within that delay then the alert automatically closes without unnecessary disruption to the dispatcher. The delay in minutes is configured as a parameter to the **Check For Automatic Recall (M1-CHKAUTRCL)** algorithm type. You may override this value as needed.

- If and when the activity is recalled, it automatically gets canceled instead of being rescheduled. A completion message is sent to the host indicating the activity was canceled. The dispatcher alert to follow up on this cancellation request gets automatically closed as well.
- If the activity was not automatically recalled in a timely manner, the dispatcher assigned to address the alert follows up with the crew.
- If the crew (or external contractor company) did not perform the work yet, they would decline the activity. This would automatically cause the activity to be canceled, similar to being recalled.
- If for some reason it is not possible to wait for the decline message to reach the server, the dispatcher can manually force the cancellation of the activity by choosing the [Forced Cancel](#) action on the alert raised to handle this situation.
- If the crew (or external contractor company) has performed the work already, the activity cannot be canceled. The host system would receive the completion information as usual.

Depot Related Activities

When an activity involves a transfer of goods between a specified location and a depot, the activity specifies one or more depots from which goods should be delivered or to which goods need to be returned.

The scheduler plans the route such that it also includes a trip to the depot to pick up goods before a delivery activity or to return goods picked up by a collection activity.

Refer to [About Transfer of Goods](#) for more information.

Handling Keys

If a key or card is needed to gain access to a work site or facility, the crew cannot start such activity without it. The crew needs to pick up the key before traveling to the work site. These activities need to specify the key identifier as well as the key location in order to provide information about where the crew can pick up the keys to gain access to the work site.

Keys are typically located at common locations such as service centers or depots. As such, key pick up is handled using depot functionality where the scheduler creates a depot task to pick up the key as part of scheduling such activity. A distribution depot must be set up for every location from which keys would be picked up.

If the crew has multiple activities that require them to pick up keys at the same location, the scheduler creates one task to pick up all of those keys when possible. Using the mobile application, when at the depot the crew is presented with a list of the activities that need keys and the description of each key to pick up.

Refer to [About Depot Runs](#) for more information.

Once the crew has the key or card, the activity should not be assigned to another crew or resource. If a dispatcher attempts to change the assigned crew, he or she receives a warning.

Activities that require keys may be scheduled to crew shifts as well contractor capacities. However, when scheduled to shifts, they are restricted to shifts that have the same logoff location as the location where keys are picked up to allow the crew to return the keys at the end of their shift. This is assumed to occur as part of the your business practice as the system does not validate whether or not a key was actually returned by the crew.

NOTE: Activities already associated with depot to pick up or drop goods cannot use keys since an activity can only be in a single depot run. This means that there cannot be depot activities that require "goods" to be picked up from depot A and require "keys" to be picked up from depot B. Then the activity would have to be in two depot runs to pick up both goods and keys.

About Complex Activities

Activities that require multiple shifts to complete are referred to as complex activities and are typical of, but not limited to, construction work. They are also sometimes referred to as long cycle work.

Complex activities are similar to regular activities in many aspects. It is important that you familiarize yourself with how regular activities are managed in order to better understand complex activities.

Refer to [About Activities](#) for more information.

NOTE: Transfer of goods, appointment booking, conditional booking, and site delay cannot be used with complex activities. They can also not be used for emergency work and are not eligible for capacity contracting. Refer to the [Capacity Model](#) for more information about capacity based contracted work.

This section describes concepts related to managing complex activities.

A Complex Activity Is Scheduled Differently

A complex activity is too long to complete during a single crew's shift. As such, it cannot be scheduled like a regular activity. The scheduler splits the complex activity into several crew visits to the work site and progress is tracked per visit. Each visit is a new task record created by the scheduler and scheduled to a crew's shift, very much like a regular activity.

When a crew completes a visit task, they can indicate whether this visit has completed the entire work or if more work is still needed they can adjust the estimated remaining time. A dispatcher may also update the complex activity in the same way.

With each visit completed, the scheduler re-optimizes the whole complex activity. If the entire activity is not completed yet, future planned visits are adjusted based on the remaining time indicated. If more time is needed to complete the activity, the scheduler may create more visits and/or extend the last visit accordingly. If less time is needed, the scheduler may delete future visits and/or decrease the duration of the last visit as needed. Once the complex activity is complete, remaining future visits are removed from the schedule.

Complex activities are considered mandatory work.

The following sections describe topics related to scheduling a complex activity.

Refer to [Use Case Example](#) for an example of how a complex activity is scheduled.

The Scheduler Creates Visit Tasks

Visit tasks are solely created by the scheduler when scheduling a complex activity.

Crews assigned to perform this type of complex work are assumed to be busy performing that work for their entire shift. They may occasionally be asked to take on emergency or other priority tasks, but for the most part, they are asked to perform only the complex activity for the day. Therefore, the scheduler assumes one visit would be needed per shift.

The complex activity type defines the average and minimum site time per visit. The scheduler splits the long complex activity into visits of average length and a last visit in the duration of the remaining time but no less than the minimum time requirement which may be used, for example, to allow the crew some time for setup and tear down.

As an example, assume a complex activity is set up with an estimated duration of 20 hours. The activity type defines a minimum visit time as 3 hours and average visit duration as 6 hours. The scheduler splits the complex activity into 4 visits, 3 visits of 6 hours each and uses 3 hours for the last visit because the minimum visit duration must be at least 3 hours. If the complex activity duration was 23 hours then the last visit would be of 5 hours.

The number of visit tasks created is determined based on the complex activity's remaining duration and the average and minimal durations defined on the task type. The scheduler reassess the number of visits needed as well as adjusts the last visit duration only when the complex activity's [remaining duration](#) changes. When the remaining duration changes, the

scheduler may add new visits, delete no longer needed visits, or adjust the last visit's duration. For performance reasons, the maximum number of free visits the scheduler creates at any point in time is defined by the **Max Free Visits for a Complex Activity** scheduler configuration parameter. Refer to this parameter for more information.

Unlike activities that have a predefined duration, the amount of time scheduled on a shift to work a visit task may differ from the duration used to create it. That is because a visit task is designed to consume the entire time available on the shift when possible. Nonetheless, only the visit's duration at creation time is considered by the scheduler, and not the visit's scheduled duration, when it comes to creating new visits or adjusting the last visit duration.

Refer to [Visit Task Duration](#) to better understand various aspects related to the duration of a visit task.

Visit tasks are created independently from any attempt to schedule them to a shift and may remain unscheduled if there are not enough compatible shifts to work them all.

Refer to [Scheduling Visits](#) for more information on how these visits are scheduled and adjusted over time.

Remaining Duration

The remaining duration of a complex activity refers to the time left to complete the entire work. The system recalculates a complex activity's remaining duration when

- A crew completes a visit task assignment and provides the actual time worked. The system relies on an accurate capture of how much time the crew has spent working the activity during each visit to properly account for the remaining time left to complete the entire activity. By default the time spent on site is assumed to be all work time but the crew may adjust the value as needed when they complete their assignment.
- A crew completes a visit task assignment and estimates the time remaining to complete the work based on their assessment in the field.

NOTE: The crew's estimated value is validated against a reasonable percentage extension limit defined on the complex activity type. If the value exceeds the limit, the crew receives a warning but is allowed to explicitly override the validation to proceed with their estimation update.

- The host system provides a correction to the overall duration of the complex work activity
- A dispatcher extends the overall duration of the complex work activity

The scheduler reassesses the need for new visit tasks, removes no longer needed tasks, or adjusts the last visit duration only on these events when the remaining duration of a complex activity changes.

Visit Task Duration

Unlike activities and other tasks, a visit does not have a fixed duration before it is dispatched.

A visit task is created with a target duration which may be viewed as a close estimate of how long the crew would work on the complex activity on a shift. All visits but the last one are created with a target duration set to the average duration defined on the complex activity type. The last visit is created with a target duration set to the remaining time but no less than the minimal duration also defined on the complex activity type.

Refer to [Creating Visit Tasks](#) to better understand when and how visit tasks are created.

A visit is anticipated to be the main task a crew is busy with on their shift. As such, it may consume more or less schedule time on the shift it is assigned to, but never less than the minimum duration specified on the activity type. Each visit task, except the last one, consumes the entire remaining time available to perform work on a shift. The last visit consumes at most its target duration.

This means that until a visit task is dispatched, its scheduled duration may differ from its target duration.

NOTE: The number of visits and the duration of the last visit is solely a product of target durations and never automatically adjusts to the optimized scheduled duration.

The scheduler is encouraged by cost control configuration to find a compatible shift with enough available time to accommodate the visit's target duration, when possible, but may schedule the visit to a shift with less time available. If a shift has more time available, then more time is scheduled to the visit.

As an example,

- Assume a complex activity is set up with an estimated duration of 20 hours. The activity type defines a minimum visit time as 3 hours and average visit duration as 6 hours.
- A visit was created based on the average duration with target duration of 6 hours.
- Assume it is scheduled to a shift of 7 hours planned duration and a planned break of half an hour. If travel to the activity location is an hour, the shift has 5:30 hours left to perform work. If no other tasks are scheduled to that shift, the visit is scheduled to consume 5:30 hours.
- In the same way, if that visit was scheduled to an 8 hour shift with same break setup, then 6:30 hours would be scheduled for that visit.

Target duration is a ballpark used to split the overall complex activity task to smaller visit tasks. As crews complete their visits and provide more accurate feedback as to the remaining time left to complete the work, then more or less visits are planned accordingly.

How close is the schedule duration compared to the target duration? If the average time specified on the activity type closely reflects average work time available on compatible crew shifts, the overall schedule of the complex activity may be insignificantly under scheduled or over scheduled. As crews complete visits and reduce the remaining time still needed to be worked, the overall scheduled duration becomes more accurate.

It is important to define realistic average duration time on the activity type so that the difference between the target duration and the scheduled duration is negligible and becomes more so over time. When setting up average duration, average travel and break time need to be taken into account and subtracted from typical shift planned time.

When the available time on a shift is less than the visit's target duration, the complex activity is under scheduled. It is preferred to set the average time slightly lower than anticipated and err on having more time scheduled on a shift for a visit than risk it being under scheduled as this situation may call for user intervention to ensure the entire work is properly planned.

Refer to [Under Scheduled Alert](#) for more information about handling an under scheduled complex activity.

Scheduling Visits

Visits are scheduled sequentially to non overlapping shifts, one visit per shift, unless a user manually allocates them differently.

Visits can overlap for a certain amount of time, for example, to allow time for one crew to pass on or debrief another crew that is taking over the activity. The maximum percentage of shift time that shifts are allowed to overlap is defined on the "Maximum Visit Overlap" scheduler configuration parameter. By default the scheduler considers shifts that overlap by more than that limit as concurrent and as such may only assign a visit of the same complex activity to one of these shifts.

A Visit Task Is Not A Business Entity

A regular activity is requested by a host system and captures details about the requested work. A visit task on the other hand is created by the scheduler as a way to divide a long running complex activity into multiple visits to the activity location until entire work is complete.

As such, a visit task bears no business information of its own. All information about the work to perform is provided on the complex activity it belongs to. A visit task may be viewed as a way to track the schedule and progress of a small portion of work time out of a long running complex activity.

A Visit Task Is Worked Like A Regular Activity

From a crew's perspective, performing the work associated with a complex activity on a specific visit is no different than working on a regular activity. A visit task is scheduled and queued for dispatch and eventually dispatched to a crew in the form of an assignment, exactly like a regular activity.

Refer to [About Assignments](#) for more information.

Actual Work Duration

When completing a visit task's assignment, the crew reports back how much time they actually worked which then impacts the remaining duration of the entire complex activity. The crew is presented with a calculated work time set to the entire time they were on site and they may adjust it as needed to provide a more accurate work time report.

Complex Activity Workflow

Complex Activity states are depicted in the following diagram:

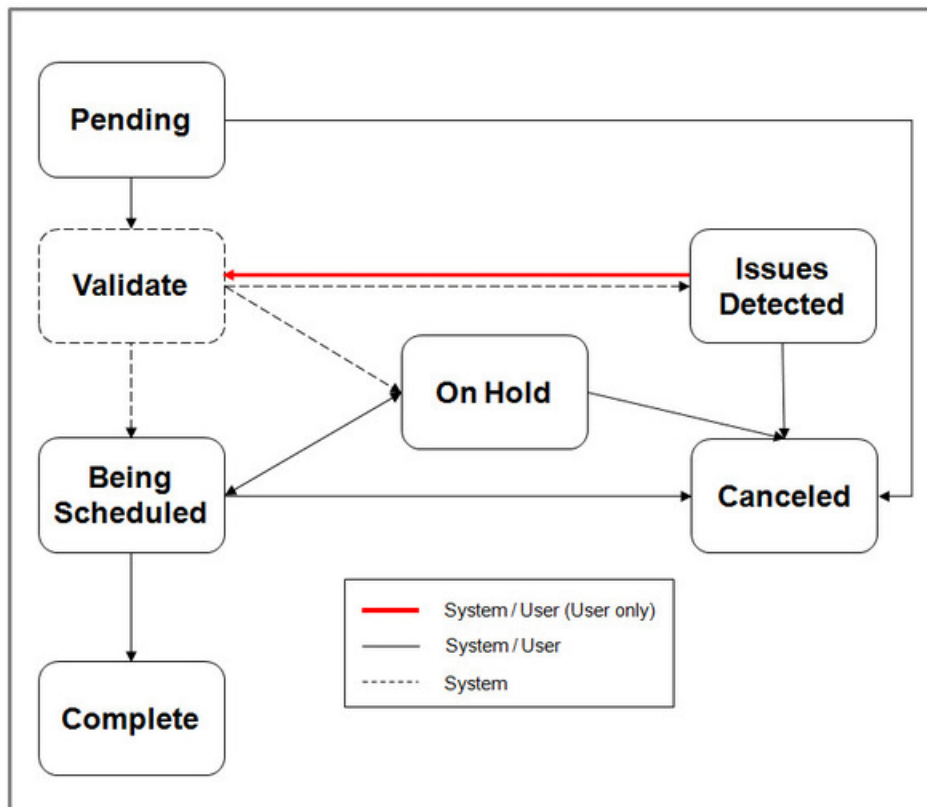


Figure 4: Complex Activity Lifecycle

The following describes key stages in the overall lifecycle of a complex activity.

Validation

A complex activity is validated very much like a regular activity.

Being Scheduled

Unlike a regular activity, a complex activity is not scheduled to a crew shift. The process of scheduling it involves splitting the activity into smaller visit tasks which are then scheduled to crew shifts.

Refer to [A Complex Activity Is Scheduled Differently](#) for more information.

While the complex activity is in the **Being Scheduled** state

- The scheduler [creates visit tasks](#) to cover the remaining duration of the complex activity as it gets updated and attempts to schedule them to crew shifts.
- The scheduler keeps improving on its scheduling decisions around the activity's visit tasks. During this continuous optimization phase, the activity's visits may frequently change their schedule until they are queued for dispatch, similar to a regular activity.
- Visit tasks are worked by crews. Refer to [Visit Task Workflow](#) for more information.
- The activity's time windows are checked and if the activity has expired, it may be automatically extended or [canceled](#) as per configuration on the activity type. Refer to [Expired Activities](#) for more information.
- The activity's time windows are recalculated to take into account new or canceled applicable override time window events, if any. Refer to [About Override Time Windows](#) for more information.
- Business rules may put the activity [on hold](#).

Completed

The activity is set to the **Completed** state when

- A crew completed a visit task assignment and indicated the entire complex activity work is complete. A completion message is sent to the host with all completion information, including any [attachments](#) captured in the field.
- The host notifies the system that the activity has been completed in the host system. This is not a typical situation but there are business practices where the crew may prefer to key in the completion information directly in the host system application. Since the host has the completion information already, there is no need to send them a completion message,

Regardless of how the complex activity is completed, any future visit tasks are deleted. If any are still dispatched to crews at that time, an alert is created to request a dispatcher to follow up with those crews similar to the process of [canceling a dispatched activity](#).

Refer to [Completion Information](#) for more information.

Canceled

The host system may send a request to cancel any further work on the complex activity. In this situation, any future visit tasks are deleted. If any are still dispatched to crews at that time, an alert is created to request a dispatcher to follow up with those crews similar to the process of [canceling a dispatched activity](#).

[Expired activities](#) may also be automatically canceled based on the activity type configuration.

When all visit tasks are finalized, a completion message is sent to the host with all completion information for work performed so far, including any [attachments](#) captured in the field.

Refer to [Completion Information](#) for more information.

On Hold

There are situations where the complex activity is sent in but it is not yet ready to be scheduled as it waits for some business event to occur or for proper timing, etc.

An existing complex activity that is already scheduled could also be placed on hold by the host system, a user, or any customer business rule.

While on hold, the complex activity is not scheduled. If it was being scheduled when placed on hold, it is removed from the scheduler and any future visit tasks are deleted. If any are still dispatched to crews at that time, an alert is created to request a dispatcher to follow up with those crews, similar to the process of [canceling a dispatched activity](#).

Visit Task Workflow

The workflow of a visit task resembles the workflow of a regular activity.

Visit states are depicted in the following diagram:

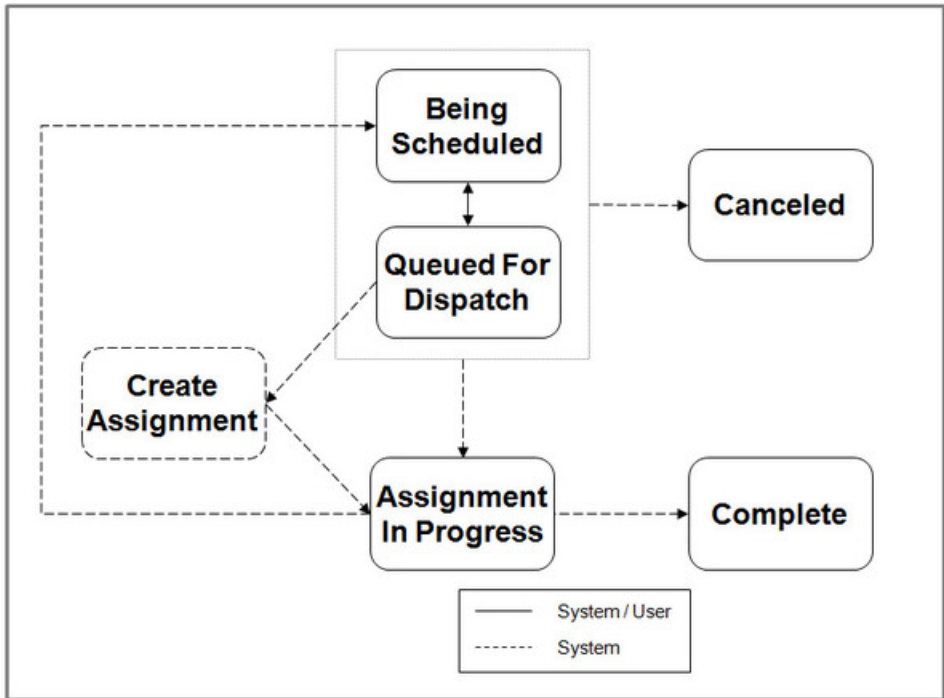


Figure 5: Visit Task Lifecycle

The following describes key stages in the overall lifecycle of a visit task.

Refer to [Activity Workflow](#) for more information.

Being Scheduled

A visit task is created by the scheduler in this state. Unlike a regular activity, a visit task does not need to go through a validation step as it contains no business information.

While the visit task is in the **Being Scheduled** state, the scheduler keeps improving on its scheduling decisions. During this continuous optimization phase, the visit task may be frequently assigned, reassigned, the order changed within its currently assigned shift's route, or become unassigned again.

Similar to the activity workflow, this goes on until it is time to queue the visit task for dispatch so it can be communicated and dispatched to the assigned crew.

Queued For Dispatch

A visit task is queued for dispatch exactly like a regular activity.

Assignment In Progress

A visit task is dispatched to a crew in the form of an assignment, exactly like a regular activity. Similar to a regular activity, the system waits for the resolution of crew assignments to assess the state of the visit task.

Completed

The visit task is set to the **Completed** state when at least one crew assignment for the visit is completed.

Unlike a regular activity, a complex activity typically requires multiple crews to complete the entire work and allows them to work concurrently if needed. As such, each crew provides their own completion information which is forwarded to the host system individually when they complete their assignment of a visit task and not when the visit task is completed. No completion is sent when the visit task itself is completed.

A final summary completion message is sent to the host when the entire complex activity is completed.

Refer to [Completion Information](#) for more information.

Canceled

The host system may send a request to cancel any further work on a complex activity or place it on hold. As part of this update to the complex activity, dispatched visit tasks may be canceled, similar to the process of [canceling a dispatched activity](#).

Use Case Example

Assume a complex activity is set up with an estimated duration of 20 hours.

Consider these parameters for this example:

- The activity type defines a minimum visit time of 3 hours and average visit time of 6 hours.
- Compatible shifts for this type of work are 8 hours shifts.
- Travel time to the activity location is 30 minutes each way, for a total 1 hour travel time.
- The crew is planned to take one 30 minute break on these shifts.

Initially the scheduler creates 4 visits: 3 visits of 6 hours average duration and the last visit of 3 hours minimal duration since the remaining duration of 2 hours is less than the minimum time required.

These visits are scheduled as follows

- Visit 1 scheduled to a shift on day 1.
- Shift available work time for the crew to perform the complex activity work during this visit is calculated as follows: 8 hours shift- 30 min break - 1 hour travel to activity location = 6:30 hours.
- Notice that even though the visit was created to consume 6 hours of work it actually consumes 6:30 hours of work because the crew is assumed to work mainly on this task for their entire shift.
- In the same way if only 5 hours were available on the shift because of other priority tasks scheduled to it then the visit would be scheduled 5 hours only. A visit would never be scheduled less than the minimum time required even if this means the crew is planned to work overtime.
- Visit 2-3 are scheduled to shifts in day 2 and 3 respectively.
- The shifts may be for same crew that started the work or another crew, depending on whether the complex activity requested that the [same crew](#) completes the entire work or not.
- Like visit 1 they are created with the target duration of 6 hours but are actually scheduled 6:30 hours.
- Visit 4, the last visit, is scheduled to a shift in day 4.
- Like visits 2 and 3 it may or may not be scheduled to the same crew.
- Unlike non last visits, the last visit cannot be scheduled more time than its target duration irrespective of the work time available on the shift. Even though 6:30 hours are available on the shift, the last visit is scheduled for 3 hours.

- Like all visits, it cannot be scheduled less than the minimal duration.

In the beginning of day 1, 3 visits are scheduled 6:30 hours each and the last visit is scheduled 3 hours which sum up to a total of 22:30 hours scheduled to work the entire complex activity.

With each visit task completion, the scheduler reassesses the remaining duration of the entire complex activity and adjusts the number of visits and/or adjusts the last visit accordingly.

- Day 1

- The crew assigned to perform visit 1 completed it and indicated they actually performed 9 hours of work as opposed to the 6:30 hours planned.
- The remaining work time for the complex activity is recalculated as 20 hours - 9 hours = 11 hours. The scheduler recalculates the number of visits needed to perform the remaining 11 hours as follows: 1 visit of 6 hours duration and a last visit of 5 hours duration.
- It now compares this with existing visits. There are 3 remaining visits out there, 2 of 6 hours assumed duration and the last visit of 3 hours assumed duration. Based on the new remaining duration there should now be only 2 visits: 1 of 6 hours and the last one of 5 hours. The scheduler therefore deletes one of the 6 hours visits and adjusts the last visit from 3 hours to 5 hours.
- When scheduled, visit 2 would be scheduled 6:30 hours and the last visit is scheduled 5 hours.
- Completion information captured by the crew on this visit is sent to the host system.

- Day 2

- The crew assigned to perform visit 2 completed it indicating they actually performed 5 hours of work as opposed to the 6:30 hours planned.
- The remaining work time is recalculated as 20 hours - 9 hours - 5 hours = 6 hours. The scheduler recalculates the number of visits needed to perform the remaining 6 hours and concludes it needs 1 last visit of 6 hours.
- It now compares this with existing visits. The last remaining visit is extended to 6 hours.
- When scheduled, the last visit would be scheduled 6 hours and not the 6:30 hours available because it is the last visit.
- Completion information captured by the crew on this visit is sent to the host system.

- Day 3

- The crew assigned to perform the last visit completed it indicating they worked 7 hours as opposed to the 6 hours planned, and that the entire work still requires 10 more hours.
- This update from the crew overrides the complex activity's remaining time to 10 hours. The scheduler recalculates the number of visits needed and creates 1 new visit of 6 hours duration and another last visit of 4 hours duration.
- The new visits are scheduled to days 4 and 5 respectively.
- Completion information captured by the crew on this visit is sent to the host system.

- Day 4

- The crew assigned to perform this visit completed it indicating they worked 9 hours as opposed to the 6 hours planned but now the entire complex work is complete.
- The scheduler deletes the unnecessary last visit.
- Completion information captured by the crew on this visit as well as a summary completion message for the entire complex activity is sent to the host system.

If the no crew indicates that the entire work is complete by the time the complex activity's time window expired, the scheduler disables the complex activity and it is no longer scheduled until the time windows are [extended](#).

Same Crew

Depending on the nature of work, either the same crew can perform the work from start to finish or multiple crews can perform the work. If same crew is to perform the work then the scheduler plans all visit tasks to be scheduled to shifts of the same crew.

This indication is defaulted from the complex activity type if not specified on the complex activity.

Under Scheduled Alert

An alert can be configured to notify a dispatcher when a complex activity is under scheduled beyond the configurable percentage threshold defined on the alert type.

A complex activity is deemed to be under scheduled if the percentage of the under scheduled amount of time as calculated by the scheduler and posted on its task schedule record relative to the remaining duration is greater than the threshold on the alert type.

Refer to [About Alerts](#) for more information.

Time Tracking

Typically the time a crew worked on site of an activity is automatically calculated from the moment they indicated they started the work until they marked it complete. This calculation may not correctly represent actual work time if the crew mistakenly started or completed the task too early or too late.

An accurate report of the actual time the crew spent working on site during each visit is key to an accurate calculation of the remaining time left to complete the entire complex activity work. For this reason, complex activities always allow the crew to review and adjust the time they spent traveling to and being on site for each visit.

Typically, crews stay on site for prolonged periods of time while they perform complex activity work. This means their breaks would most likely be taken while they are on site and therefore complex activity work allow breaks to be taken while on site.

Refer to [Crew Time](#) for more information.

Completion Information

Typically the crew enters completion information describing the work they have performed on a specific visit. For certain types of work, a single form or set of completion details may need to be filled out over the course of multiple visits. In this situation, you may configure the complex activity type to make the latest completion information available to the next crew visit. This information is only available to crews using the mobile application. The crew may choose to continue and work on the latest completion information provided or ignore it and provide new information.

The host system is updated upon completion of each crew's visit task assignment as well as when the entire complex activity is completed or canceled.

NOTE: Logic to compose the completion message for an activity resides in a "Prepare Completion Data" algorithm on the activity's business object. Click [here](#) to see the algorithm types available for this system event. You can turn off sending interim visit completion messages using an algorithm parameter.

About Activity Dependency

One activity may be dependent on another activity in various ways. A group of activities that have dependencies is referred to as "bound work".

This section describes the various types of bound work and the dependencies created when activities are tied to each other.

Job

A job is a group of activities that must be performed by a single crew during a single shift. For example, a job may consist of an activity to pick up goods from one location and deliver them to another location.

The following sections highlight concepts related to handling jobs.

Activities Are Grouped By A Job Identifier

A job is not a standalone entity but rather a set of activities referencing the same [job identifier](#).

Typically all activities of a job are interfaced together in a single service call but activities may be added individually as needed, as long as it references the same host job identifier. If a new activity is sent in later referencing an existing host job identifier, it automatically becomes associated with the corresponding internal job identifier.

Order

The order in which the activities in a job should be performed is established as usual by their time windows. A scheduling order number may optionally be specified to achieve a more explicit order. Notice that ordering should not contradict the time windows ordering.

NOTE: Only the host system can send in a set of activities linked as a job. The base product does not provide the means to compose a new job nor manipulate the set of activities included in an existing job within the application itself.

Loaded Goods Must Be Unloaded

The shift has to start and end with no load on board. A job does not necessarily have to involve transfer of goods. But if it does then all goods picked up by one activity in the job must be dropped off by another activity in the job.

Refer to [About Transfer of Goods](#) for more information.

All Scheduled or None

If there is any scheduling issue for one activity in a job, all of the activities are disabled from scheduling until the issue is resolved. Once the issue is resolved, all of the activities are rescheduled automatically by the scheduler.

Some Actions Affect Entire Job

Actions performed on an activity that is part of a job that cause the activity to be rescheduled from its currently assigned shift must affect the entire job so that all activities in the job are rescheduled together.

If such action is not performed on the entire job, the scheduler may disable the activities until they can be scheduled together again.

Therefore, performing any of the following actions on one activity in a job automatically performs the action to all other activities in the job.

- Allocation to a shift as well as removing such restriction
- Fix to crew as well as removing such restriction

- Unassigning an activity from its current shift
- Deferring the schedule of the activity to a later date
- Locking an activity to a depot run
- Excluding the contractor company currently assigned to perform the work, as well as removing such restriction
- Revert the schedule of the activity to company crews only

The CDI Gantt provides a visual cue to indicate that an activity is part of a job and therefore has dependencies. This helps the dispatcher to easily see how each activity is related and what would be the impact of certain actions, if taken.

Appointment Booking

Appointment booking and rebooking requests must include all activities in the job as the scheduler has to find availability for all of them on a shift.

Partial Completion

As explained earlier, all activities of a job must be scheduled to the same shift otherwise they are disabled by the scheduler.

There are times when the crew completes some activities of a job but does not manage to complete all of them for various reasons. To allow the scheduler to properly reschedule the remaining activities, the system automatically removes the completed activities from the job at end of shift. Those completed activities still reference their original job identifier for reporting purposes.

Viewing Activities In A Job

If an activity belongs to a job, the **Job Activities** zone on the activity maintenance portal lists all activities in the job in the order they should be performed, highlighting the current activity in that list.

In a partially completed job, only remaining non-worked activities remain in the job so that they can be rescheduled. The Job Activities zone displays all of the activities that were part of the original job.

Chain

A chain represents linked activities where each activity is dependent upon the start or completion of another activity, with dependency defined in terms of timing and assigned resource. The entire set of dependent activities represents a project like work request that is also sometimes referred to as long cycle work.

The following sections highlight concepts related to handling chains.

A Chain Dependency Can Span Multiple Shifts

Chains differ from jobs in that the activities are not bound to a specific shift, they are bound in time and can span multiple shifts. They are set up so that each activity depends upon the start or completion of the prior activity before it can start.

Dependency Rules

A dependent activity can start within an offset range relative to the start or finish time of the parent activity. An activity can also be restricted so that it can only be worked by a crew different than its immediate parent activity or by a crew other than any of those assigned to activities in the entire chain.

For example, if there is a project to dig a trench, lay pipes, and fill the trench, a chain dependency of tasks might be set up as follows:

- A surveyor is scheduled to lay out the plan before the start of work.
- Only after the surveyor is finished can the digging begin.
- The pipes cannot be laid out until the trench is complete.

- The trench cannot be filled until the pipe is laid.
- and so on

The following considerations apply for chains:

- An activity can only be immediately dependent on one other activity which could also be dependent on another activity and so on.
- A circular dependency is not supported, i.e. the activity cannot be dependent on itself even indirectly . (Example: A->B->C->A)
- Appointment booking, conditional booking, and chooser requests can only be used for an activity in a chain if the activity does not have any other activities dependent on it.
- The time windows for a dependent activity must be wide enough to support its timing dependency constraints, including offsets, with its parent activity. This has to take into account potential early or later start of the parent activity. For example, assume activity B is scheduled to start an hour after activity A starts. Since activity A may start earlier than planned by more than one hour, activity B should be defined with time windows wide enough to accommodate a reasonable early start of activity A. In other words that applying an offset towards the actual start of activity A would still be within the time window for activity B.
- Depot related activities may not be included in a chain.
- If the scheduler finds conflicts because these considerations have not been followed, the whole chain (non-started part of it) is disabled and not scheduled.

Assist Activities

There are times where a crew arrives on site and realizes they need assistance from another crew to help with the work of an activity. For example, a crew responds to an outage activity but upon arrival, realizes there is a tree that blocks access to the work location. They contact the dispatcher that organizes an assisting crew to help with the removal of the obstacle.

The assist activity needs to be scheduled at the same time as the original activity so that both crews are on site at the same time. In addition it must be scheduled to a crew other than the original activity's crew. The original and assist activities are an example of a small chain dependency.

The scheduler attempts to assign the assist activity to a suitable crew. However, the dispatcher can also manually allocate it to any specific crew as needed. To do that the dispatcher could take advantage of the shift chooser query option which consults the scheduler for a cost based recommended list of shifts to choose from.

Only specific activity types may be selected to create assist activities.

Refer to [Assist Activity Types](#) for more information.

Host Initiated Only

Only the host system can send in a set of activities linked as a chain. With the exception of [assist activities](#), the base product does not provide the means to compose a new chain, nor manipulate the set of activities included in an existing chain within the application itself.

Complex Activity Chains

Complex Activities can be part of a chain of activities, and can precede or follow regular or complex activities. However, any restriction dependency of a complex activity towards another activity and vice versa are not applicable because complex activities may be worked by more than one resource over time.

Viewing Activities In A Chain

The CDI provides a visual cue to indicate that an activity is part of a chain and therefore has dependencies.

If an activity is part of a chain, you can navigate from the activity portal, as well as from the CDI portal, to the [Activity Dependency](#) portal for an overview of the scheduling state of entire chain.

About Assignments

An assignment is a copy of an activity (or a complex activity's visit task) created by the system to track the assigned resource's (crew or a contractor company) progress in accomplishing the work.

Refer to [About Activities](#) and [About Complex Activities](#) for more information.

When the activity is ready to be worked by the assigned resource, the system creates their assignment and sets the activity to the **Assignment In Progress** state. The activity remains in this state until the assigned resource either completes their assignment or declines it.

Refer to the [Activity](#) and [Visit Task](#) workflows to better understand when an assignment is created and how it affects the activity when finalized.

Once the crew acknowledges their assignment, they may provide status updates informing the system of their progress. For example, when they start traveling to the location, when they have arrived there, when they start to perform the work as well as to provide all the required information when completed it.

When the activity is assigned to a [contracted capacity](#), work is outsourced to the contractor company and the system no longer has visibility into when and by whom the work is actually performed, nor does it know how many attempts the contractor company made to complete the work. Your company's expectation is to receive the completion information when the work is completed, otherwise a decline update if the work cannot be completed in a timely manner so it can be rescheduled.

This section describes concepts related to managing assignments.

Current Assignment

At any point in time an activity may only be scheduled to a single crew's shift (or [contracted capacity](#)) and if dispatched, there should only be one active assignment for it. The current assignment for an activity is the assignment created for its currently assigned resource.

Information about the work that needs to be performed is communicated to the crew (or contractor company) as part of the assignment dispatched to them. Any change made to this information afterwards is re-dispatched to the resource as long as they are still the currently assigned resource. For example, if work was reassigned to another crew then the previously assigned crew is not informed of the changes. Whether or not the crew (or contractor company) receives such updates in a timely manner depends on the nature of their communication method.

If the crew does not get to complete an activity assigned to them, their assignment for that activity is [returned](#) at end of their shift and the activity may be rescheduled to the same or other crew on a later day, repeatedly until the work is complete. In a similar way, if the contractor company determines they cannot perform the work in a timely manner, they return their assignment for the activity allowing it to be rescheduled as needed.

Over time there could be multiple assignments for the same activity, each representing an attempt to work it.

Refer to [Concurrent Assignments](#) to learn about situations where multiple active assignments may exist for the same activity.

Assignment Workflow

An assignment is created when the activity lifecycle reaches a point in which work should be dispatched to the assigned resource. The lifecycle of an assignment affects the lifecycle of its activity.

The assignment workflow differs when it tracks the progress of a crew working an activity as opposed to a contractor company's agreement to perform an activity.

The following sections describe key stages in each workflow.

When Assigned To A Crew

When the crew is ready to perform work and has started their shift, the system creates an assignment for each of their assigned activities to record their progress working that activity. If the crew uses the mobile application to manage their work, a copy of their assignment is sent to their mobile device.

Assignment states are depicted in the following diagram:

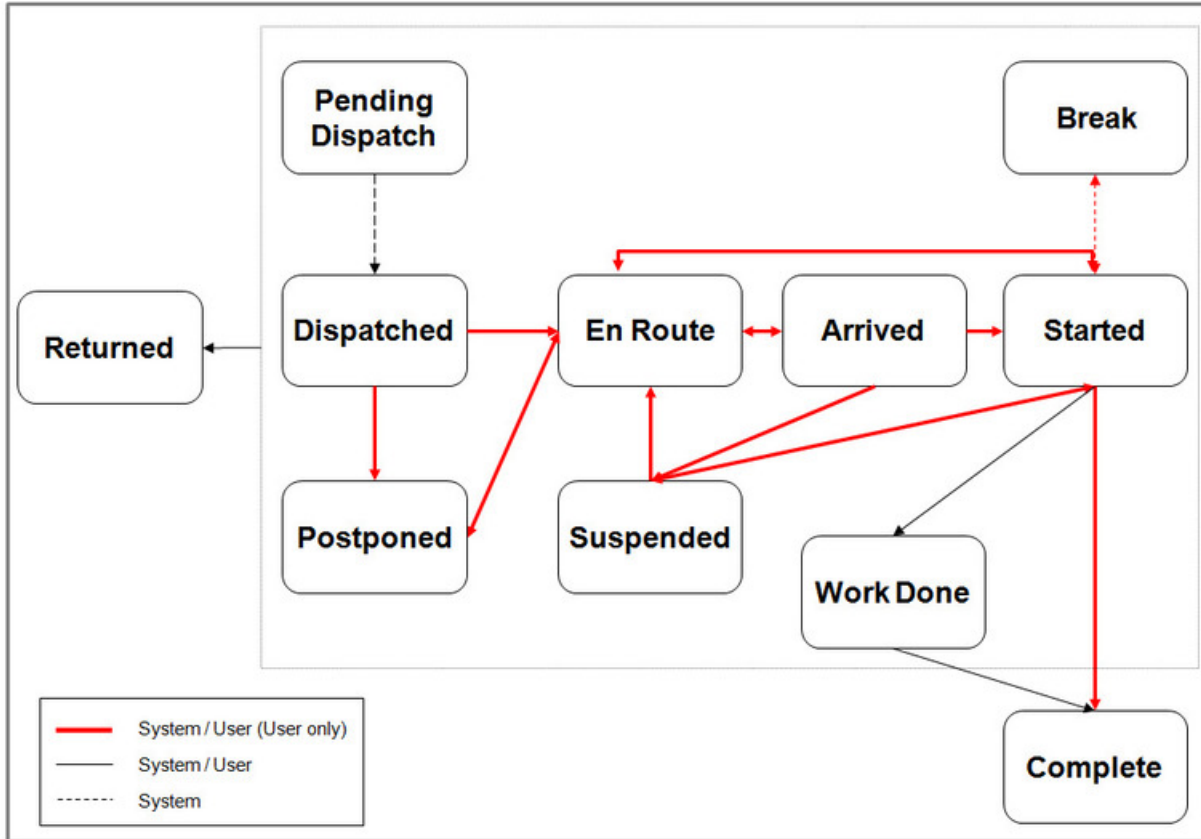


Figure 6: Crew Assignment Lifecycle

The following sections describe the crew's assignment overall lifecycle from the moment it is dispatched to the crew until the crew either completes the work or returns it back to be rescheduled.

Refer to [Crew Assignment](#) for more information.

Pending Dispatch

An assignment is initially created in this state waiting to be acknowledged as [dispatched](#).

If the crew uses a mobile device, a copy of the assignment is sent to the crew's device. If the crew uses SMS to work their tasks then a message with a brief description of the work is sent to their mobile phone.

Dispatched

An assignment is considered dispatched if the assigned crew acknowledges its receipt. This is the last step in the process of dispatching an activity. Refer to [The Dispatch Process](#) for more information

The mobile application acknowledges receipt of the assignment by sending a message to the server to update the assignment to the dispatched state. Refer to [Automatic vs. Manual Acknowledgment](#) for more information.

If the crew does not use a mobile device the assignment is automatically marked dispatched.

En route

This state indicates that the crew is currently on its way to the activity. The crew provides this [status update](#) when they are about to start traveling to the activity location. If the crew was previously out of service they are now considered back in service.

Time spent while in this state is considered travel time for [crew time](#) reporting.

Arrived

Typically the crew may start work immediately upon arrival to the activity location. This state indicates that the crew has arrived at the location but has not yet started work for some reason. For example, they may need to wait for the superintendent to provide access to the work site. The crew provides this [status update](#) when they arrive on site but cannot start the work right away.

Time spent while in this state is considered wait time for [crew time](#) reporting.

Started

This state indicates that the crew is on site and has started to perform the work requested by the activity. The crew provides this [status update](#) when they start the work.

Using the mobile application, the crew may also provide interim completion details as they work on the activity and send them right away to the server, even before fully completing the work. They may also report their estimation of how much time is still needed to complete the work when the activity takes much more or much less time than planned.

Time spent while in this state is considered work time for [crew time](#) reporting.

Break

This state indicates that the crew is taking a planned break while working on site. The system automatically sets the assignment to this state when the crew indicates they start their planned break. When the crew completes their break, the system automatically switches the assignment status back to started.

Refer to [Breaks](#) for more information.

Postponed

This state indicates the crew decided to postpone working on this activity any time prior to arriving on site, with the intention to come back to it later. Once the crew arrives on site the activity can no longer be postponed, but it can be [suspended](#) instead.

For example, the crew may have been on their way to the activity location and were instructed by a dispatcher to postpone this activity and handle an emergency situation or assist another crew nearby.

The system prompts for a reason for the postponement and the time to which the activity should be postponed, which should be within the current crew shift. A postponed activity remains dispatched to the shift but the scheduler re-optimizes its position within the route based on the estimated postponed time.

The crew provides this [status update](#) in that situation.

Suspended

This state indicates that the crew arrived on site and started work, but was unable to complete it for some reason. For example, the crew may be instructed by a dispatcher to handle an emergency situation or assist another crew nearby, etc.

The activity remains assigned to the crew, allowing them to return to it later on. At this time the system provides the crew with the option to also keep the activity with that crew for a future shift, if they are not able to complete it during the current shift.

The system prompts for a reason for the suspension, the time to which the activity should be postponed, which should be within the current crew's shift, as well as how much time is still required to complete the work once it is resumed. A suspended activity remains dispatched to the shift but the scheduler re-optimizes its position within the route based on the estimated postponed time.

The crew provide this [status update](#) in that situation.

If the crew had entered activity specific completion information while they were on site but never got to complete the activity the information they have entered is sent to the server when their assignment is returned. The system makes this partial completion information available to the next crew that is assigned to work the activity. This information is only available to crews using the mobile application. The crew may choose to continue and work on the latest completion information provided or ignore it and provide new information.

Completed

This state indicates that the crew has completed work and all completion details have been entered.

The crew may complete their assignment but provide a reason code that marks it incomplete when actual work is impossible to perform. For example when details on the activity are invalid or there was some problem denying the crew access to the work site.

Crews that use the mobile application are required to provide the necessary completion information when they complete their assignment. If they are connected or become connected later on, the mobile application forwards the entire completion information to the server.

If the crew uses SMS to work their schedule, they cannot provide the necessary completion information using text messages and therefore cannot complete their assignment. They can only mark their assignment as [work done](#). Later on when they are back at the service center, they may enter the completion information they have recorded on paper and complete their assignment.

In the same way, if the crew uses no device (non-MDT), they typically contact a dispatcher to report their [work done](#) and complete their assignment later when they are back at the service center.

Completing the assignment causes the activity to reevaluate its state.

Refer to the [Activity](#) and [Visit Task](#) workflow to better understand how finalizing an assignment reassesses the activity state.

Work Done

An assignment in this state indicates the work is performed but the crew did not yet enter the necessary completion information. This is typically used by crews that uses SMS text messages to report status updates as well as crews with no device (non-MDT). Later on at the office, the crew enters the completion details and completes the activity.

Returned

An assignment in this state indicates that the crew is no longer responsible to perform the activity and returns it back to be rescheduled if possible.

An activity assignment may be returned in the following situations

- The crew did not manage to work the activity before the end of their shift.
- The crew explicitly declined taking on the activity for various reasons. The crew has to specify the reason for declining the activity. If the reason is due to lack of worker or vehicle capability, a To Do Entry is created for a resource manager to follow up and ensure the capability information recorded for the crew's mobile workers and vehicles is accurate.
- The system automatically recalled the activity due to a host request to cancel the activity. Refer to [Activity Cancellation By Host](#) for more information about how such request is handled
- The dispatcher performed a schedule change action that [forced a recall](#) of the activity from the crew.

If the crew had entered activity specific completion information while they were on site but never got to complete the activity the information they have entered is sent to the server when their assignment is returned. The system makes this partial completion information available to the next crew that is assigned to work the activity. This information is only available to crews using the mobile application. The crew may choose to continue and work on the latest completion information provided or ignore it and provide new information.

Returning the assignment causes the activity to reevaluate its state.

If the shift is not completed, the shift is added to the activity's list of [excluded shifts](#) to prevent scheduling the activity back to the same shift.

Note that once the crew starts working on their assignment for a complex activity's visit the assignment can no longer be returned as it is needed to track the amount of time worked during this visit.

Refer to the [Activity](#) and [Visit Task](#) workflows to better understand how finalizing an assignment reassesses the activity state.

When Assigned To A Contractor Company

When an activity is committed to a capacity based contractor, the contractor company receives a copy of the activity information in the form of an assignment that they use to report their progress working that activity.

The following diagram depicts the lifecycle of a capacity contractor assignment.

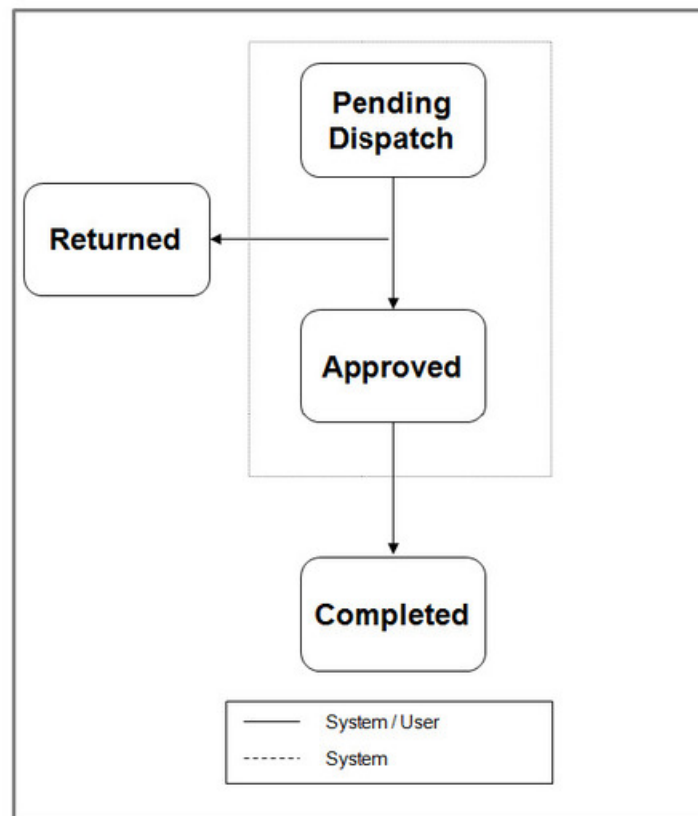


Figure 7: Contractor Assignment Lifecycle

The following sections describes the contractor company's assignment overall lifecycle from the moment it is dispatched to them until they complete the work or decline it.

Refer to [Managing Committed Work](#) for more information.

Pending Dispatch

An assignment is initially created in this state waiting to be [approved](#) or [declined](#) by the contractor company.

If the contractor company fails to approve taking on this activity within the agreed upon approval time, the system automatically declines their assignment and sets it to the Returned state.

Approved

This state indicates that the contractor company has agreed to take on this activity. From this moment on, the system has no visibility as to how the contractor company carries out the work and when.

Completed

This state indicates that the contractor has completed work and all completion details have been entered.

Completing the assignment causes the activity to reevaluate its state.

Refer to the [Activity](#) and [Visit Task](#) workflows to better understand how finalizing an assignment reassesses the activity state.

Returned

An assignment in this state indicates that the contractor company is no longer responsible for performing the activity and returns it back to be rescheduled.

An activity assignment may be returned in the following situations

- The contractor company has either declined taking on the activity committed to them or failed to approve taking the work in a timely manner. The contractor company is [excluded](#) from further working the activity.
- The dispatcher performed a schedule change action that [forced a recall](#) of the activity from the contractor company.

Returning the assignment causes the activity to reevaluate its state.

Refer to the [Activity](#) and [Visit Task](#) workflows to better understand how finalizing an assignment reassesses the activity state.

Crew Assignment

An assignment is created for each activity dispatched to a started shift. The following sections describe key aspects related to a crew's activity assignment.

Automatic vs. Manual Acknowledgment

Typically, it is enough for the crew's copy of the assignment record to reach the mobile device for the mobile application to automatically acknowledge the dispatch. The mobile application immediately sends a status update back to the server to mark the server copy of the assignment record as dispatched.

For certain activities though, typically though not limited to emergencies, mere receipt of the assignment on the device is not enough. The crew has to manually acknowledge taking on the work. If they cannot take on the work, it is important to know that as soon as possible so that the work is rescheduled quickly to another crew. Any activity, not just emergencies, may request manual acknowledgement. This may be configured on the activity type as well as requested on an individual activity.

When the activity explicitly states that a manual acknowledgement is required, the mobile application prompts the crew to accept or decline the new assignment sent to them. If the crew accepts the new assignment the mobile application sends a message to mark the server copy of the assignment as dispatched. If the crew declines taking on the assignment the assignment on the server is declined which causes the activity to be rescheduled, excluding this crew.

Note that until the assignment is acknowledged, manually or automatically, it appears to the dispatcher on the Gantt as not yet dispatched.

If the assignment is not delivered to the device in a timely manner, either automatically or manually, a synchronization alert may be raised for a dispatcher to follow up with the crew. Refer to [Synchronization Alert](#) for more information.

Synchronization Alert

When dispatching an assignment to a mobile device fails to complete in a timely manner (typically due to poor connectivity), the system creates a synchronization alert for a dispatcher to follow up with the crew.

NOTE: Initiating a synchronization alert is optional. An alert is only created if an active alert type referenced as the **Synchronization Alert Type** [global configuration](#) parameter.

Refer to [About Alerts](#) for more information.

A dispatcher assigned to handle such an alert may contact the crew and verbally communicate to them the key details of the activity. The crew can then create a local copy of the activity on their mobile device, capturing the key details provided and use it instead of the original assignment to record status updates and completion details.

This local copy of the activity has to be properly associated with the actual activity when the former is sent to the server. The activity identifier is too long to communicate by voice to the crew when they create their local copy so the system provides a shorter system generated alternate identifier to be used instead. The alternate identifier is captured on the original activity as well as provided to the crew as part of the details they capture on their local copy. The local copy of the activity is also known as a field-referenced activity (FRA) for that reason.

NOTE: Logic to handle the business rules associated with setting an alternate id for an activity resides in a "Set Alternate ID" algorithm on the activity's business object. Click [here](#) to see the algorithm types available for this system event.

Handling an assignment synchronization situation works as follows

- The dispatcher receives an alert when the synchronization of an assignment to the crew's mobile device takes too long or fails.
- The dispatcher clicks the **Voice** button on the Alert Queue. This action marks the assignment that failed to be sent as dispatched and cancels any further attempts to send it to the crew, so the assignment won't exist twice on the device once the FRA is created. Then, the system automatically navigates to the Activity portal and displays a message instructing the dispatcher to communicate the activity details to the crew. This message also includes the short alternate identifier of the activity.
- The dispatcher communicates the information to the crew (typically via phone), and the crew creates an FRA on the mobile device referencing the activity alternate identifier.
- The crew works the activity, updating its status as work progresses through completion, just like any other assignment. When the crew's connectivity is reestablished, the system automatically matches the FRA to the corresponding activity using the alternate identifier.

Keeping Mobile Copy Up To Date

If the crew is using the mobile application to manage their work they receive a copy of each of their dispatched assignments (and other tasks) sent to their mobile device from the moment they start their shift.

When key activity information changes after the assignment is dispatched, the system automatically synchronizes the mobile copy of the assignment and sends the crew a new copy of it with the latest activity information.

NOTE: Logic to initially send a copy of the assignment to the crew's mobile device as well as synchronize the mobile copy to keep it up to date when changes are made on the server resides in a "Synchronize Task" algorithm on the assignment's business object. Click [here](#) to see the algorithm types available for this system event.

Crew Status Updates

The mobile application allows a crew to advance the state of their assignment to reflect their current progress working the activity. They may do so even when they are not connected. The actual time of events is captured locally on the device and is forwarded to the server later on when they become connected again.

If the crew uses SMS to work their schedule, they send in a text message to update the state of their assignment on the server.

A crew that uses no device (non-MDT) typically calls in to report their status and a dispatcher makes the update to the assignment on their behalf.

Refer to [Working The Schedule](#) for more information.

Keep With Crew

Crews may keep certain activities with them for future shifts. Any work the crew did not manage to complete by the end of their shift may be rescheduled to a future shift for this or any other crew. A crew may decide to keep a certain activity they did not manage to complete by the end of their shift so that they can work on it the next day or during a future shift.

For example, a crew had to suspend working on an activity to respond to an emergency activity and never got the time to go back and complete the suspended activity. If the crew has keys or parts needed to perform the suspended activity, they may decide to retain the activity so that they can work on it the next day or any of their future shifts and not have it rescheduled to another crew.

The crew can make this request at the time they postpone, suspend, or decline an activity assigned to them. When the crew ends their shift, they also receive a reminder to review and mark activities they want to keep with them.

Refer to [Fix To Crew](#) for more information.

Crew Time

When the crew completes an activity, the system can be configured to present the crew with a summary of the time they spent on behalf of this activity in the following categories:

- **Travel.** This time is calculated from the moment the crew indicates they are en route to the activity location and until they indicate they are no longer traveling.
- **Wait.** This time is calculated from the moment the crew indicates they have arrived on site and until they start the work, capturing the amount of time the crew had to wait before they could start the work.
- **Work.** This time is calculated from the moment the crew start to work and until they indicate they are no longer working- either they complete the activity or suspend it for some reason.
- **Break.** This time is calculated for each break that the crew start and complete while being onsite. Refer to [Taking a Break While On Site](#) for more information.

If the crew has attempted to travel to the activity, then postponed it in order to work on another activity, and then resumed back to travel again to the previous activity, the time spent on these categories from all attempts are added up.

The crew may adjust the time information as needed to correct situations when they have provided status updates earlier or later than when they actually occurred. Both the calculated times and the adjusted times are sent back to the host system as part of the activity completion information.

NOTE: Crew time review is incorporated as part of the activity completion data entry, if enabled on the activity type, as it may be applicable to certain types of work only.

Recall

Recalling an activity refers to a system initiated request to finalize the [current assignment](#) of an activity for the currently assigned resource. When the recall is successful, the resource is no longer responsible for performing the activity.

The recall request may or may not succeed in recalling the activity. When it is successful, the assignment is set to the **Returned** state with a reason code stating it was returned due to a recall.

An assignment may be recalled when

- A dispatcher attempts to take certain actions on an activity that is already dispatched. Refer to [Forced Recall](#) for more information.
- The host system requests to cancel the activity. Refer to [Activity Cancellation By Host](#) for more information about how such request is handled.
- If allowed by shift configuration, the scheduler may make improvements to a crew's dispatched route that result in some activities being recalled and replaced with others. Refer to [Crew Undispatch and Sequence Lock](#) for more information

For crews that utilize the mobile application, the system attempts to recall the activity by sending a recall request to the crew's mobile device. When the recall message is received on the device, the mobile application determines whether the assignment is in a state that allows it to be recalled or not. If recall is possible, the assignment is automatically returned from the mobile device and removed from the crew's task list, notifying the crew of the recall. The recall is complete when the assignment status update reaches the server.

If the crew is only equipped with a SMS device, the system sends the crew an SMS message notifying them of the recall. The crew can either ignore the request if it's too late to recall or confirm the recall by sending a recall confirmation message back to the server. The recall confirmation message completes the recall by causing their assignment on the server to be returned.

The assignment remains with the assigned resource and the activity retains its schedule if the recall is not successful for any reason. If this was a forced recall, i.e. initiated by user in order to force a schedule change for a dispatched activity, then concurrent assignments may now exist for the activity.

Refer to [Concurrent Assignments](#) for more information.

For crews not using any device (non-MDT), there is no automated way of requesting and confirming a recall. The only recall allowed from these crews is that forced by a dispatcher action. In this situation the system assumes that the dispatcher has contacted the crew and confirmed the recall verbally before making the change and therefore their assignment is automatically recalled.

NOTE: The system is configured to consider all assignment statuses non recallable when the crew is already at the location.

Concurrent Assignments

The system would only schedule one crew (or contractor company) to work an activity at any point in time. As such, there should be only one active assignment for the activity, the one created for the [currently](#) assigned resource.

If that resource could not perform the work, their assignment is finalized indicating they have returned the activity to be rescheduled. When the activity is rescheduled to another resource, a new activity assignment is created for that resource which is now the currently active assignment for the activity.

However, due to a dispatcher initiated recall, it is possible for multiple active assignments for the same activity to exist concurrently. These are referred to as concurrent assignments.

Refer to [Forced Recall](#) for more information.

Typically the dispatcher coordinates the recall with the crews involved such that only one of them actually performs the work and completes their assignment where the others just return their assignment.

However, if there is some confusion in the field and more than one crew actually completed their assignment, the system cannot determine from which completed assignment to take the completion information to send back to the host. For example, a crew might arrive at the location, see that the work was already done, and then complete their assignment with no work details.

In this situation the system sets the activity to the [Collision Detected](#) state and raises a To Do Entry for a user to review the completed assignments and choose one assignment to complete the activity by clicking the "Complete Activity" action on that assignment.

About Activity Attachments

An attachment can be any file, such as a photo, a document with instructions or specifications, a spreadsheet, or any other supporting documentation. Several standard attachment file types are supported such as PDF, Word, Excel, or other document types.

Additional information about an activity, either sent from the host, added by the dispatcher, or captured by the crew, may be associated with an activity in the form of an attachment. This section describes functionality associated with activity related attachments.

Attachments Related To The Type Of Work

Attachments might be linked to specific activity types for the purpose of providing supporting information and guidelines to crews when performing this type of work. For example, an attachment may include instructions on how to perform the work or safety guidelines the crew should adhere to when performing such work.

Some attachments may be referenced on multiple activity types as applicable and are referred to as common or shared attachments.

Links to these attachments are available to the crew on their mobile device when they work activities of this type so they can download their content on demand, provided they are connected (to save bandwidth, time and storage space, etc.).

Common attachments the organization prefers to always be available on the device can be explicitly added to the mobile application deployment. This way the crew does not need to manually download them each time. Refer to [Managing Attachments](#) for more information.

These attachments may not be modified by the crew.

Attachments Related To A Specific Activity

Some attachments contain content related to a specific activity. The attachment may be provided by the host system along with the activity, captured by the crew in the field, or attached by a dispatcher to the activity.

This section describes functionality associated with activity specific attachments.

Attachments Sent In From The Host System

The host system may send supporting information in the form of attachments along with the activity.

For crews using the mobile application, the system includes links to these attachments when dispatching the activity to the crew's mobile device. The crew can download their content on demand provided they are connected (to save bandwidth, time and storage space, etc.).

These attachments may not be modified by the crew.

Capturing Content In The Field

Crews using the mobile application may capture content or attach documents related to a specific activity in the field. These attachments are later on synchronized and uploaded to the server.

NOTE: The types of attachments that can be captured depends on the mobile application and the mobile device used. In addition, attachments may be adjusted in size to not exceed a certain size restriction if specified on the mobile device type.

Sending Attachments To The Host System

By default, any attachments created in the field are always sent back to the host system as part of the activity completion information as it is the system of record for all types of completion information. Any other attachments added to the activity within the system must be marked explicitly in order to be sent to the host. Attachments sent in from the host system are never sent back.

For example, the host sends in an activity to investigate a possible electricity diversion by a customer. During the investigation, the mobile worker confirms the electricity diversion and takes a picture of the evidence. That picture is returned to the host system in the form of an attachment along with the investigation activity completion information. On a subsequent disconnect service activity resulting from the investigation, the host sends in the picture taken during the investigation to aid the mobile worker in the disconnection activity. The picture would not be returned to the host system again.

NOTE: If the host system is not configured to handle attachments you may change the default behavior to not send attachments back to the host system. Refer to the algorithm type **Attachment Audit Rules algorithm type** ([M1-SNDATUPDV](#)) for more information.

Dispatchers May Add Attachments

A dispatcher may attach files to a specific activity as needed. If the activity has already been dispatched, then for crews utilizing the mobile application the system sends the link to that attachment to their mobile device right away.

At the time of adding an attachment to an activity, the user can indicate whether it should be sent to the crew or not. For example, an attachment may not be required for work or of any interest to the crew in the field but rather added for the purpose of sending it to the host system when the activity is completed.

Maintaining Activities

This portal is used to display and maintain an activity.

Refer to [About Activities](#) for more information.

You can access the portal from the **Main Menu > Service Management > Activity**. You are brought to a query portal with options for searching for a specific activity. Once an activity has been selected you are brought to the maintenance portal to view and maintain the selected record.

The following zones may appear as part of the portal's **Main** tab page

- **Activity.** This zone provides information about the activity, including the various maintenance actions applicable to its current state.
- **Activity Time Windows.** This zone lists the various time windows available for the activity.
- **Activity Assignment Completion Details.** This zone displays the completion information provided by the assigned resource when they completed the work. The zone is only visible when the activity is completed.
- **Assignment History.** This zone lists all resource assignments for the activity over time.
- **Attachments.** This zone lists all attachments specific for the activity including those captured by crews in the field.

- **Job Activities.** This zone lists all activities that belong to the same job as the current activity. The zone is only visible if the activity is part of a job.
- **Activities Created By Crew.** This zone list activities created by the crew that are related to this activity.
- **Procedures.** This zone lists all procedures completed by the crew for this activity. The zone is visible only if such procedures exist for the activity.

The **Activity Schedule** zone on the portal's **Schedule** tab page displays the current scheduling information for the activity. Refer to [Scheduling An Activity](#) for more information.

NOTE: Depot transfer activities are not scheduled to crews and as such information related to crew compatibility and shift schedule is not applicable. Refer to [Depot Replenishment](#) for more information.

Request An Assisting Crew

Use the **Assist** action to create a new activity for the assisting crew of a specific activity type. The system creates the assist activity and navigates to the Activity portal so that you can add a note to the crew if needed. You can either allow the scheduler to assign it to an available crew or you can manually allocate the activity to a specific crew. If the original activity has already started, which is typical in assist situations, the scheduler attempts to dispatch it right away. You can also manually dispatch the activity as needed.

The current activity and the assisting crew's activity are bound as a chain so that the scheduler plans their arrival accordingly.

Refer to [Assist Activities](#) for more information on activity dependency.

Maintaining Current Crew Assignment

Use the current assignment link on the activity display zone to navigate to the activity's current assignment to a resource. This may be either an assignment for a crew on a specific shift or an external contractor assignment to perform the work as part of an agreed upon work capacity.

Refer to [About Assignments](#) for more information.

Crews that do not use any mobile device (non-MDT) to perform their work may rely on a dispatcher to act on their behalf and update the state of their assignment to reflect their progress. For example, if the crew cannot perform an activity and needs to decline it they may call in to report the situation. You can navigate to the current assignment and decline it providing the appropriate reason etc. This allows the activity to be rescheduled to another resource if possible.

However, crews using the mobile application maintain their copy of the assignment locally on the mobile device itself. When they make updates to their local copy the mobile application forwards the update the server when they become connected.

NOTE: It is important to let such crews fully control the state of their assignment and avoid manually updates to their assignment on the server. Manual updates may contradict status updates sent in from the crew and lead to unpredictable results. It is strongly recommended to confirm and coordinate such updates with the crew.

Completing Activities with Concurrent Assignments

If multiple completed assignments are received for the same activity, the system cannot automatically complete the associated activity. To resolve the situation, review each of the completed assignments for the activity and navigate to the assignment you want to use to complete this activity. Use the **Complete Activity** button to complete the activity with the completion information provided by this assignment.

Refer to [Concurrent Assignments](#) for more information.

Maintaining Complex Activities

This portal is used to display and maintain a complex activity.

Refer to [About Complex Activities](#) for more information.

You can access the portal from the **Main Menu > Service Management > Activity**. You are brought to a query portal with options for searching for a specific activity. Once an activity has been selected you are brought to the maintenance portal to view and maintain the selected record.

The following zones may appear as part of the portal's **Main** tab page

- **Activity.** This zone provides information about the activity, including the various maintenance actions applicable to its current state.
- **Activity Time Windows.** This zone lists the various time windows available for the activity.
- **Assignment History.** This zone lists all crew assignments for the activity from all visit tasks over time.
- **Attachments.** This zone lists all attachments specific for the activity including those captured by crews in the field from all visit tasks.
- **Activities Created By Crew.** This zone list activities created by crews at the activity location that are related to this activity.

The following zones may appear as part of the portal's **Schedule** tab page

- **Activity Schedule.** This zone displays the current scheduling information for the activity. Refer to [Scheduling An Activity](#) for more information.
- **Complex Activity Gantt.** This zone displays a graphical representation of the crews assigned to work the complex activity. The Gantt works in the same way as the dispatcher's scheduling Gantt on the CDI portal with the exception that filters and the ability to launch the geographic map are not applicable from this zone. Refer to [Using the Scheduling Gantt](#) for more information.
- **Crew Visits.** This zone lists visit tasks created for the activity.

Maintaining Assignments

This portal is used to display and maintain an activity assignment to a resource.

Refer to [About Assignments](#) for more information.

You can access an assignment from the [activity maintenance portal](#). You may navigate directly to the activity's [current assignment](#) or to any assignment listed in the "Assignments History" zone.

The **Assignment** zone on the portal's **Main** tab page provides information about the assignment, including the various maintenance actions applicable to its current state.

NOTE: Depot transfer activities are not scheduled to crews' daily shifts and as such differ from regular activities in how a work assignment is tracked. Refer to [About Depot Transfer Assignments](#) for more information.

Activity Dependency

The portal provides a schematic and a Gantt view of the chain dependency to which a specific activity belongs to. If an activity belongs to a chain dependency you can also navigate to this portal from the Dependency section on the activity details zone of the [activity maintenance portal](#).

Refer to [Chains](#) for more information.

The following zones appear as part of the portal's **Main** tab page

- **Dependency Gantt.** This zone displays a graphical representation of the crews assigned to work activities that belong to the current activity's chain dependency. The Gantt works in the same way as the dispatcher's scheduling Gantt on the CDI portal with the exception that filters and the ability to launch the geographic map are not applicable from this zone. Refer to [Using the Scheduling Gantt](#) for more information.
- **Dependency Hierarchy.** This zone shows a schematic view of the current activity's entire chain of dependencies.

Override Time Windows

About Override Time Windows

There are times when the hours of operation for certain types of work in all or some service areas may need to deviate from normal typically due to unplanned business events. Work that is normally allowed may become not allowed and vice versa, work that is normally not allowed may become allowed during that event.

This section describes concepts related to managing activity override time windows.

Prohibit Work When Otherwise Allowed

An override time window may be used to prohibit the schedule of certain type of work for the duration of an unplanned business event. For example, prohibit disconnects or deliveries from being performed during periods of extreme weather conditions.

Allow Work When Otherwise Not Allowed

Certain types of work may be not be scheduled on holidays specified on the [work calendar](#) associated with the activity type. They may also be restricted to business hours defined on the [work profile](#) associated with the activity type.

An override time window may allow such work to be performed if needed during those holidays or outside the specified business hours. For example, allow work on a holiday to manage a backlog of certain type of work.

Overrides Regenerate Arrival Time Windows

An override time window defines the period of time, the type of work and areas impacted as well as whether work is allowed or not allowed during that time. The override impacts activities for which at least one of their effective time windows (ETW) overlaps with the duration of the event.

When an activity is added or when its time window information changes, the system automatically calculates its arrival time windows(ATW) taking into account any overrides.

NOTE: Due to the potentially high number of existing activities that may be impacted by an override, it cannot be applied to them right away. The impact of introducing a new override or canceling an already applied override onto existing activities is delayed to the next time they are monitored by the Task monitor batch process. When next processed, the activity recalculates its arrival time windows(ATW) if it detects an applicable override or an override cancellation.

Refer to [When Can The Activity Be Actually Scheduled?](#) for more information about the impact of an override on an activity's arrival time windows(ATW).

Enforcing Overrides

The scheduler may schedule an activity to be started or completed after the end of an arrival time window (ATW) if this results in an overall cost benefit. In other words, lateness is cost driven and not a hard limit.

An override event to prevent certain work from being performed has a legal binding and as such, the system has to ensure work is not scheduled during that time at all costs. While the arrival time windows (ATW) of impacted activities are regenerated to exclude override periods, the scheduler also ensures that activities are not scheduled late, outside their time windows, and into the override event period.

NOTE: The scheduler honors such override as long as the activity has not been dispatched yet.

Once dispatched, activities that overlap with such override require manual intervention. The system provides a [KPI](#) to report these violations, allowing the dispatcher to contact the crew and ensure work is recalled.

Use Case Examples

For this example assume repair work is associated with a work profile that allows work to be scheduled from 8 AM to 5 PM Monday through Friday and is not allowed to be performed on the New Year's Day holiday.

Consider the following scenario:

- A repair activity is received from the host with an effective time window that starts on Monday 7 AM until Saturday 2 PM the week of the New Year's Day holiday.
- The activity is created with the following arrival time window (ATW) that week
- Monday 8 AM - 5 PM
- Tuesday 8 AM - 5 PM
- Wednesday 8 AM - 5 PM
- Thursday 8 AM - 5 PM
- No ATW is generated for Friday as that is the New Year's Day holiday.
- No ATW is generated for Saturday since it is not a business day based on the work profile.
- Due to a snowstorm expected to hit the area on that Tuesday, the company is required to suspend repair work in the area. To achieve that the company enters an override time window that does not allow repair work on that day and lists the specific service area.
- Due to a backlog in repair work, the company needs to allow repair work to be scheduled on that Friday from 9 AM to 6 PM in all areas even though it is the New Year's Day holiday. To achieve that the company enters an override time window that allows repair work on that Friday.
- Next time the Task monitor batch process processes the activity its arrival time windows (ATW) would be regenerated as follows
- Monday 8 AM - 5 PM
- No ATW for Tuesday because of the override to prevent work that day.
- Wednesday 8 AM - 5 PM.
- Thursday 8 AM - 5 PM

- Friday 9 AM - 6 PM because of the override to allow work during that time.
- No ATW is generated for Saturday since it is not a business day based on the work profile.
- New activities sent in from the host would have their ATWs reflect the overrides right away when created.
- On Monday evening the weather forecast cancels the snowstorm alert and the company cancels the override placed on Tuesday.
- The next time the Task monitor batch process processes the activity its arrival time windows (ATW) would be regenerated to add back the ATW for Tuesday.

Maintaining Override Time Windows

This portal is used to display and maintain an override time window event.

Refer to [About Override Time Windows](#) for more information.

You can access the portal from the **Main Menu > Service Management > Override Time Window**. You are brought to a query portal with options for searching for a specific override event. Once an override event has been selected you are brought to the maintenance portal to view and maintain the selected record.

The Override Time Window zone on the portal's **Main** tab page provides information about the override time window, including the various maintenance actions applicable to its current state.

NOTE: An override time window cannot be changed or deleted once added, it can only be canceled.

Crews

About Crews

A crew is a uniquely named group of one or more [mobile workers](#) and [vehicles](#) that perform work in [shifts](#).

The specific mobile workers and vehicles that make up a crew may be typically constant, like single person crews, or they can vary by day of week, from shift to shift. A crew might also perform different types of work and cover different service areas on different shifts.

This section describes concepts related to managing crews.

Work Is Scheduled To Crew Shifts

A crew is available to perform work during their shifts. The specific mobile worker(s) and vehicle(s) allocated to a crew may vary from one shift to another, and a crew may perform different types of work and cover different service areas on different shifts. Hence work is scheduled to crew shifts, and not to crews, as they represent the availability of that crew as well as what they should do and where they should work during that time.

[Mobile workers](#) and [vehicles](#) make up the crew's actual resources on a specific shift. A mobile worker or a vehicle may be part of only one crew shift at any given time.

Refer to [About Crew Shifts](#) for more information.

Mobile Workers

A mobile worker is an individual resource that performs work as part of a crew.

This section describes concepts related to managing mobile workers.

Worker Capabilities

An activity may require the mobile workers of the assigned crew to have certain work related [capabilities](#) in order to perform the work. These may be skills or activity specific requirements such as the need for a security clearance, building access, drug testing, permits, language proficiency, etc.

Some capabilities may have effective periods. For example, a skill might require a license that must be renewed every year. The mobile worker record lists the capabilities the individual has acquired over time.

A [mobile worker's type](#) defines worker capabilities that workers holding that job description are assumed to have. This list is used as the default capabilities list when setting up a new mobile worker, or when changing their mobile worker type, but can be further adjusted to remove or add capabilities as needed to describe the individual's actual capabilities.

When scheduling work, the system matches the worker capability requirements of each activity to the capabilities of the workers that make up the crew on a specific shift. It is important to properly describe a mobile worker's actual capabilities to avoid sending them work they are not able to or should not perform.

Base Areas

A mobile worker may be associated with one or more base service areas within which they typically work.

This information is intended to facilitate the search for suitable mobile workers to work in a certain area. It is not used to restrict the mobile worker to work solely in these areas.

Relative Efficiency

This specifies the mobile worker's efficiency relative to expected efficiency for mobile workers of this type.

The scheduler may plan for more or less time to work an activity based on the overall efficiency of the crew on that shift; the overall efficiency is the average of the relative efficiency of all mobile works allocated to that shift.

Home Address

A mobile worker's home address may be used as their logon or logoff location on their shift when the mobile worker is the only person on that crew. For this reason the home address needs to be geocoded so that the scheduler can plan the shift route to start and end from the home location.

Unless explicitly provided the system attempts to geocode the address using the [geocoding service](#) defined on the installation options expecting an exact match.

NOTE: Logic to geocode the home address resides in an algorithm associated with the mobile worker's business object. Review this algorithms and adjust its match quality parameter as needed.

Allow To Work Shift Using SMS Messaging

When a mobile worker uses SMS messaging to work their shift schedule, the system identifies the worker and the shift they work using their mobile phone number. The mobile worker has to be explicitly allowed to use their mobile phone to work their shift this way.

Refer to [Using SMS Messages](#) for more information.

My Calendar

As a best practice, you may want to set the home page for mobile workers to the [My Calendar](#) portal. This allows them to view their shift schedule and any planned leaves on a calendar view.

Vehicles

A vehicle may be needed for the crew to perform their work. The vehicle is mainly used for travel but may also carry pieces of equipment or have certain features that are required to perform certain types of work.

This section describes concepts related to managing vehicles.

Vehicle Capabilities

An activity may require the assigned crew's vehicle to have certain pieces of [equipment](#) in order to perform the work. For example, the crew may need a bucket truck to perform line pole work.

Some equipment may be available on the vehicle for limited time only. The vehicle record lists equipment associated with this vehicle over time.

An activity may also require the assigned crew's vehicle support the transfer of certain [categories of load](#). Some types of goods may require the crew's vehicle to meet certain business conditions before it can be used to transfer such goods. For example, assume an activity to transfer a certain type of fuel. A load category for that type of fuel may be specified on the activity requiring the vehicle to be able to handle that type of fuel. To avoid product contamination, a vehicle that was used to carry another type of fuel in its container cannot be used to carry this activity until its container is rinsed.

A vehicle may be associated with a list of load categories. This information is defined on the vehicle type but may be specified for a specific vehicle as needed.

A [vehicle's type](#) defines equipment assumed to be available on vehicles of that type. It may also define a list of load categories. These are used as the default equipment and load category lists when setting up a new vehicle but can be further adjusted to remove or add pieces of equipment and load categories to reflect those actually available on and supported by the vehicle.

When scheduling work, the system matches the equipment and load category requirements of each activity to the equipment and load categories actually available on and supported by the crew's vehicle on a specific shift. It is important to properly describe a vehicle's actual capabilities to avoid sending the crew work they are not able to perform.

Relative Speed

This defines the vehicle's speed relative to the expected speed for vehicles of this type.

The scheduler may plan for more or less travel time based on the overall speed of the vehicles utilized on that shift; overall speed is calculated as the slowest relative speed of all vehicles in use on the shift.

Fixed Crew Resources

Based on the [crew type](#) configuration, if the resources that make a crew are typically the same on every shift, they may be specified as the default resources on the crew record.

For example, the base product supports a setup for a single person crew where the mobile worker and their vehicle are defined as part of the crew record. Your implementation may support other fixed resource configuration. These resources are then defaulted onto any new shift created for the crew, whether added online or generated by a template.

Note that this does not prevent the crew from having different resources on specific shifts as needed.

Resource Capabilities May Change

When scheduling work, the system matches the requirements of each activity to the capabilities of available crew shifts. The scheduler considers the capabilities of all mobile workers and vehicles allocated to a crew shift to determine which activities the crew can work on a particular shift.

Adding or removing worker capabilities from a mobile worker and/or pieces of equipment from a vehicle impacts all non-finalized crew shifts to which the mobile worker and the vehicle are currently allocated. The scheduler then reevaluates the shift schedules accordingly, and may add or remove activities as a result.

However, if an activity has already been dispatched or marked for dispatch, it remains so. It is the responsibility of the crew to decline an activity they cannot perform. Once the activity is declined it may be rescheduled to another qualified crew.

Crew Locations

A crew may be affiliated with specific locations for various purposes.

For example, if the crew typically starts their shift from the same service center or depot, it can be defined as their default logon location. The system then would use this information when creating shifts for the crew and a logon location is not explicitly specified.

Planned Time Off

A mobile worker may be on leave for a period of time.

A vehicle may be scheduled for some maintenance work causing it to not be available for a period of time.

A crew may need to attend work related meetings during which they are not available to perform work.

Periods of time when a resource is not available to perform work impacts their available work capacity. This information is recorded in the system using a period of unavailability record.

Refer to [About Periods of Unavailability \(POUs\)](#) for more information.

Crew Hierarchy

Depending on your business practices, crews might be organized within a reporting hierarchy.

Using crew hierarchy can help with reporting and also allows for advanced searching and/or filtering crews in the Gantt and from related records.

Refer to [Understanding Crew Hierarchies](#) for more information.

Depot Restrictions

If your organization uses depots, a crew may be restricted to specific depots they are allowed to visit. If depot restrictions are indicated on a crew, the crew's shift schedule is restricted to the specific depots indicated. If no depots are specified, the crew can travel to any depot as part of their schedule.

If for some reason a crew on a specific shift should not be allowed to visit a depot they are normally allowed to visit, this information can be manually overridden on the specific shift.

Refer to [About Depots](#) for more information.

Crew May Be Managed Externally

The product supports the following types of integration flows with a host system.

Crews Managed By Us

In this type of integration the system fully manages crews, schedules work to them, optionally communicates with the crews via a mobile application, and forwards the specific completion information they have provided for each activity back to the requesting host system.

As such, all the information needed to perform an activity as well the data needed to be collected in the field to complete it has to be available to the crew on the mobile application.

The system may be configured to send interim activity status updates to the host or the host may request that on demand.

This is the default configuration.

Crews Managed By The Host System

In this type of integration the system is used solely as a scheduling engine.

Crews are managed by the host system itself which sends in (through synchronization) enough details to generate shifts for scheduling purposes.

Crews communicate directly with the host system which then forwards activity status updates to the current system so as to provide real time feedback to the scheduling process. In the same way, activity completion information is sent directly to the host system and never passes through this system.

To indicate a crew is externally managed, reference the external system that manages it and with which the system integrates.

Refer to [Crews Managed By Host](#) for more information.

Maintaining Mobile Workers

This portal is used to display and maintain a mobile worker record.

Refer to [Mobile Workers](#) for more information about mobile workers and their relationship to crews.

You can access the portal from the **Main Menu > Resource Management > Mobile Worker**. You are brought to a query portal with options for searching for a specific mobile worker. Once a mobile worker has been selected you are brought to the maintenance portal to view and maintain the selected record.

The **Mobile Worker** zone on the portal's **Main** tab page provides information about the mobile worker.

The **Calendar** zone on the portal's **Calendar** tab page provides a calendar view of shifts and other items on the mobile worker's schedule for the current month.

Refer to the Calendar (F1-CALENDAR) zone type for more information.

Maintaining Vehicles

This portal is used to display and maintain a vehicle record.

Refer to [Vehicles](#) for more information about vehicles and their relationship to crews.

You can access the portal from the **Main Menu > Resource Management > Vehicle**. You are brought to a query portal with options for searching for a specific vehicle. Once a vehicle has been selected you are brought to the maintenance portal to view and maintain the selected record.

The **Vehicle** zone on the portal's **Main** tab page provides information about the vehicle.

The **Calendar** zone on the portal's **Calendar** tab page provides a calendar view of shifts and other items on the vehicle's schedule for the current month.

Refer to the Calendar (F1-CALENDAR) zone type for more information.

Maintaining Crews

This portal is used to display and maintain a crew record.

Refer to [About Crews](#) for more information.

You can access the portal from the **Main Menu > Resource Management > Crew**. You are brought to a query portal with options for searching for a specific crew. Once a crew has been selected you are brought to the maintenance portal to view and maintain the selected record.

The following zones may appear as part of the portal's **Main** tab page

- **Crew.** This zone provides information about the crew.
- **Shift Weekly Templates.** This zone lists shift weekly templates the crew is associated with over time.

The **Calendar** zone on the portal's **Calendar** tab page provides a calendar view of shifts and other items on the crew's schedule for the current month.

Refer to the Calendar (F1-CALENDAR) zone type for more information.

My Calendar

This portal provides a calendar view of your shifts and other items on your schedule for the current month. The portal is typically used by mobile workers, company or contractor users, as it provides them with an easy access their shifts and their schedules. If they do not use the mobile application to perform their work they may navigate to their daily shift from this portal in order to review their activities and enter completion information. This should also be set as the home page for these users.

You can access this portal from the **Main Menu > Resource Management > My Calendar**.

Resource Route Replay

This portal shows a replay of the actual route crews and vehicles during a specific period of time based on GPS data reported either from the crew's mobile device or directly from a vehicle's GPS system.

This functionality may be used to analyze route patterns and for auditing purposes. You can review routes for up to 12 crews and/or vehicles, with each route displayed in a different color.

You can access this portal from the **Main Menu > Resource Management > Resource Route Reply**.

The following zones may appear as part of the portal's **Main** tab page

- **Crews / Vehicles Search.** This zone is used to select up to 12 crews and/or vehicles to replay their routes. Broadcast each resource and when ready use to the Routes zone to complete the query.
- **Routes.** This zone shows the routes for the crews / vehicles selected using the search zone and for a specified period of time provided by the user.

Refer to [Shift Route Replay](#) view the route for a specific crew shift.

Crew Shifts

About Crew Shifts

A crew shift is a planned period of time in which a crew (one or more mobile workers and vehicles) is scheduled to perform work. For example, a crew shift might be defined as Crew A, consisting of Jack and his van, working on a specific date from 9 am - 5 pm.

Refer to [About Crews](#) for more information.

This section describes concepts related to managing crew shifts.

Crew Shift Type Controls Everything

Each crew shift references a [shift type](#) which defines attributes common to all crew shifts of that type, including the shift start time and duration of the shift etc.

For example, night shifts would have a different start time and duration than morning shifts, so they would have different shift types.

Creating Shifts

Typically, crew shifts are generated by templates that define the basic characteristics of the shifts. A crew's shifts may be generated by a single shift weekly template at any given point in time.

Refer to [About Shift Weekly Templates](#) for more information.

Shifts can also be manually created to address special situations on particular days.

Default Crew Shift Details

When creating a shift for a crew, either manually or via a shift template, certain details may be defaulted from the crew record. For example, a single person crew's record would state the mobile worker and potentially the vehicle that typically make up the crew each day. This information is used as to default the resource allocation for the crew when a shift is created for it online or by the shift generation process.

NOTE: Logic to default resource allocation and any other shift related information from the crew record onto a new shift upon creation resides in a "Prepare Shift Data" algorithm on the crew's business object. Click [here](#) to see the algorithm types available for this system event.

Crew Shift Workflow

The following describes key stages in the overall lifecycle of a crew shift.

Pending

A shift is initially created in this state. While in that state it is not sent to the scheduler and is considered a draft until marked as fully planned.

Planned

A shift in this state is no longer a draft and is available for scheduling as per its configuration.

If generated by a template, the shift is automatically set to the planned state. A user creating a shift online has to explicitly mark it as planned when set up is complete and the shift should be sent to the scheduler.

Started

When the crew starts the shift, their working state is either in service, i.e. performing work, or out of service, i.e. busy with non work related tasks such as taking a break, attend a meeting etc. Refer to [Planned Unavailability](#) for more information. You may prevent crews from starting their shift too early by setting up a time limit on the shift type.

By default, when the crew starts their shift they are assumed to be out of service until they start an activity at which time the system sets their shift to in service automatically.

Once the shift is started, the system automatically dispatches tasks that are queued for dispatch to the crew. Refer to [Locking The Schedule](#) for more information.

Out Of Service

This state indicates that the crew is not currently working or available for work.

Each time the crew starts a non work related task, such as attending a meeting, their shift is switched back to out of service, if it is not already in this state.

A crew may also be out of service for unplanned circumstances, for example, their vehicle breaks down. The crew may use the mobile application to report they are out of service and for how long. They can also do the same if they use SMS messaging to work their shift. Crews not using any device may choose to notify the dispatcher of their state so it can be reflected in the system. The scheduler is notified of this event and pushes the work on the crew's schedule to after the crew's estimated time they expect to be available again.

In Service

This state indicates that the crew is currently working (traveling to an activity location, working an activity, idle between activities, etc.). Each time the crew starts a work related task, such as an activity or a depot task, their shift is switched to in service, if it is not already in this state.

While in service, time between completing one task and before starting another is considered idle time.

Completed

This state indicates that the crew is done working. Activities the crew did not manage to complete are returned back to be rescheduled. Refer to [End of Shift](#) for more information on this process.

Canceled

This state indicates the shift is canceled and should no longer be used for scheduling. If any tasks are currently scheduled to the shift, they will be rescheduled.

Planned shifts that were not started before their planned end time are automatically canceled and removed from the scheduler.

A user or any process can cancel a shift that has not started yet. Once started, it can only be completed.

Controlling The Schedule

This section describes the various aspects of a shift used by the scheduler to control the crew's schedule of work on that shift.

Fixed Shifts

By default, the duration of a shift is fixed and defined as the time between its planned start and end times. The crew is assumed to start their shift at the planned start time and end it at the planned end time. While typically shifts are shorter than a day, there is not such limitation in the system; they can be set up to any length as the business requires.

A default duration is defined on the [shift type](#), allowing different types of shifts to be of different lengths; for example, night shifts could be shorter than day shifts. When setting up a shift online, the system automatically sets the planned end time based on the planned start time entered and the default duration, but allows the user to adjust the shift to be shorter or longer as needed.

By default, the scheduler starts to incur overtime cost for scheduling work beyond the earliest of the planned end time and the planned duration of time. If the crew start their shift late their schedule will be subject to overtime cost starting at the shift's planned end time. However, if the crew starts their shift early their schedule would be subject to overtime cost after the planned duration time which is earlier than the shift's planned end time. If your business rules require that overtime be applied at the planned end time even for crews starting their shift early you may indicate so the [shift type](#).

Variable Shifts

When crews work on appointments, there may be a lot of idle time between the time they start their shift and their first appointment. To help with more efficient scheduling, shifts can be set up as "variable shifts" to provide the crew with some flexibility on when they can start and end work.

A variable shift defines a broad time window within which a fixed duration shift is planned based on work load. The shift has to retain a fixed duration so the latest the shift can start is the broad time window end time minus the fixed duration. For these shifts, the scheduler continuously adjusts the estimated start work time as the start time of the first task scheduled to the shift, minus any travel time, but never later than the earliest shift start time.

The estimated start time is locked at some point prior to the broad time window start time so as to give the crew enough time to know when they should show up for work. Once locked, the shift is updated and sent to the scheduler as a [fixed shift](#) that is planned to start at the locked estimated start time and planned to end after the fixed shift duration specified. The estimated start time is rounded down based on the rounding option specified on the shift. The amount of time prior to start that the shift is locked is set by the advanced dispatch offsets shift configuration.

For example, a variable shift with a planned range of 8:00 AM to 8:00 PM and a fixed duration of 9 hours can start anytime between 8:00 AM to 11:00 AM (the start window is defined as the earliest shift start time to the latest end time, minus the duration - a three hour offset in this example). Depending on the shift start time, the shift will end 9 hours later for an end time between 5:00 PM and 8:00 PM. If the first activity is scheduled after the latest shift start time, the shift will start at 11:00 AM to ensure that the requirement for a 9 hour duration is met.

The default duration of the broader flexible period of time is also defined on the [shift type](#) along with the default fixed shift duration. When setting up a variable shift online, the system automatically sets the broad planned end time based on the planned start time entered and the default variable duration but allows the user to adjust the shift to be shorter or longer as needed.

Refer to [Advance Dispatching](#) for more information.

NOTE: [Non productive tasks](#) cannot be created for variable shifts.

Logon / Logoff Locations

The scheduler considers the crew's location at any point in time, including where they start and end their shift, to determine the most optimal route for them.

A crew shift must either reference a service center location or indicate the [mobile worker's home](#) as their logon and logoff locations. The latter is only applicable when the shift is for a single worker.

You may specify a maximum travel time the crew is expected to spend traveling from/to the logon/logoff location on their own time.

You can also reserve time at the beginning and/or end of the shift to perform review and preparation work. Note that the scheduler ignores logon time reservation when scheduling an emergency activity to this shift allowing the crew to address the emergency as soon as possible.

Crew Allocation

A crew on a shift has to include at least one mobile worker. A vehicle is not required unless it is needed by the crew to perform their work. Multiple mobile workers and vehicles may work together as a crew on a specific shift.

A crew and any of its allocated resources can only be part of one shift at any given point in time. This means shifts cannot be planned for the same crew or mobile worker or vehicle to overlap in time.

This section describes concepts related to allocation of resources to a shift.

Crew Allocation May Change At Logon

When a crew leader starts their shift, the additional resources that make up the crew that day may differ from those originally planned. For example, if a worker reported sick and is replaced by another worker that day, the crew leader can specify the new worker instead of the sick person.

Changes to the crew allocation are communicated to the scheduler, which may reconsider the planned work for the shift and rebuild the shift schedule as needed to reflect the [crew's current capabilities](#).

Resource allocation may also [change throughout the shift](#).

A crew and any of its allocated resources can only be part of one shift at any given point in time. Whether at shift start time or later, the system always validates changes to crew allocation to ensure neither crew nor resources are already part of another started or completed shift at this time.

Crew Allocation May Change Mid Shift

There are situations where the crew's resource allocation may change mid shift.

For example, a mobile worker may leave their crew mid shift for some reason or their vehicle breaks down and is replaced. In this situation, the crew's schedule for the remaining of their shift may be adjusted to meet the crew's [recent change of capabilities](#).

The scheduler reevaluates not yet dispatched work and may add or remove activities as a result. However, if an activity has already been dispatched or marked for dispatch, it remains with the crew. It is the responsibility of the crew to decline an activity they can no longer work. Once the activity is declined, it may be rescheduled to another qualified crew.

The system keeps track of all the resources that were part of the crew throughout the shift for reporting and timekeeping purposes.

Crew Capabilities

The scheduler assigns an activity to a crew that is capable of carrying it out. The crew's overall capabilities on a specific shift are those associated with the mobile workers and the vehicles [allocated](#) to that shift.

Refer to [Worker Capabilities](#) and [Vehicle Capabilities](#) for more information.

By default a crew is considered having a certain capability on a specific shift if at least one of the resources (worker or vehicle) allocated on that shift possess it. If a capability has to be possessed by all members of the crew, i.e. all workers for a worker capability or all vehicles for equipment, then the [capability](#) itself needs to be defined as required by all. In this situation, the crew shift is considered having the capability if all relevant resources have it.

Covered Areas

A crew may work in specific service areas on a shift. The scheduler would only assign the crew work within these service areas. A shift must cover at least one service area.

A shift may be configured to prefer work in some service areas over others. The scheduler is encouraged to assign work in the preferred service areas over other areas covered by the shift.

Primary Function

Types of work may be classified by service class. One of these service classes can be designated as the primary function of the crew on a specific shift. This encourages the scheduler to prefer work of that class, if any, over other classes of work.

Not Allowed Work

A crew shift may list specific service classes of work that the crew should not be assigned on that shift, regardless of whether they are capable of performing that work or not. The scheduler would not assign work that belongs to these service classes to that shift.

Reserve Capacity

A certain percentage of shift capacity may be reserved for up to 10 different categories of work.

Common reasons to reserve capacity on a shift may be as follows:

- To offer appointments up until the day before, your organization would like to reserve a certain percentage of a shift's time for appointments so that other non-appointment work does not take precedence over appointments.
- Your organization may have a lot of emergency work during certain months of the year for which they need to reserve a certain percentage of a shift's time for emergency work so that it does not get filled up with appointments that might be missed if pushed out to accommodate emergencies.
- Some time may be reserved for non appointment low priority work to handle backlogs.

A work category may be broadly defined for emergency work, appointments or non appointment work, and may be further restricted to specific activity types by referencing a service class.

Each reserve states the minimum and optionally, a maximum percentage of shift time to reserve for that category of work. Note that the maximum percentage is not applicable to emergency work as it needs to be addressed as soon as possible without any capacity restriction.

In addition, each reserve defines whether the reserved capacity should be released some time before or after the shift is planned to start, or not released at all.

For example, you might reserve 10% of a shift's time for emergency work that would never be released and 50% for appointment work that would be released 1 hour prior to planned shift start time.

Shift Cost Profiles

A shift is associated with a shift cost profile to define how the scheduler should cost travel time, overtime and other scheduler costs for this shift relative to other shifts.

Refer to [Understanding Shift Cost Profiles](#) for more information.

Depot Restrictions

If your organization uses depots, and for some reason a crew on a specific shift should not be allowed to visit a depot they are normally allowed to visit, this information can be overridden on the specific shift.

Refer to [About Depots](#) for more information.

Standby Shifts

Standby shifts are potential shifts that can be activated only when needed. They are not considered as available capacity, and as such are not scheduled until the crew actually starts the shift or the dispatcher [opens](#) it to be scheduled as usual.

For example, a crew may be planned to have a stand by shift on Sundays to handle emergencies. If the crew is indeed called on that Sunday and starts that shift, it becomes an active shift as usual and gets sent to the scheduler.

A shift is typically marked as "stand by" when created, but it can also be changed to standby while this shift is still in the pending or planned state.

Crew Own Travel Time

Depending on your business requirements, shifts may define a maximum period of time the crew is expected to spend traveling from the logon location to the first task of the day on their own time. If the time required to travel to the first task exceeds the time specified here, the delta of travel time is built into the schedule. In the same way, a maximum period of time the crew is expected to travel from the last task of the day to the logoff location on their own time may also be configured on the shift. If the time required to travel from the last task exceeds the time specified here, the delta of travel time is built into the schedule.

For example, assume a crew's shift is planned to start at 8 AM and configured such that the crew is expected to travel up to 30 minutes at the beginning of the shift on their own time. With this configuration, the scheduler is allowed to plan for the crew to start driving as early as 7:30 AM and up to 30 minutes of his driving time to the first scheduled task is considered their own time, i.e. not part of the paid shift period of time. Note that when there is no early task scheduled for this shift that requires an early travel, the scheduler plans their route to start at the logon location as usual at 8 AM (or even later for a variable shift). Regardless of whether travel is planned to start earlier than the shift plan start time or not, up to 30 first minutes of travel time is considered not paid.

Let's continue with the same example and assume the shift configuration is for a fixed planned start time shift, i.e. not a variable shift.

- When the time window for the first activity starts at 8 AM:
 - If travel to the activity's location is estimated to take 10 minutes the scheduler plans for the crew to start driving at 7:50 AM.
10 minutes of travel time are considered to be the crew's own time and shift paid time is calculated to begin at 8:00 AM.
 - If travel to the activity's location is estimated to take 30 minutes the scheduler plans for the crew to start driving at 7:30 AM.
30 minutes of travel time are considered to be the crew's own time and shift paid time is calculated to begin at 8:00 AM.
 - If travel to the activity's location is estimated to take 50 minutes the scheduler plans for the crew to start driving at 7:30 AM.
30 minutes of travel time are considered to be the crew's own time and shift paid time is calculated to begin at 8:00 AM.
- When the time window for the first activity starts at 9 AM:
 - If travel to the activity's location is estimated to take 10 minutes the scheduler plans for the crew to start driving at 8:00 AM.
10 minutes of travel time are considered to be the crew's own time and shift paid time is calculated to begin at 8:10 AM, i.e. 50 minutes of waiting idle at the door is considered paid time.
 - If travel to the activity's location is estimated to take 30 minutes the scheduler plans for the crew to start driving at 8:00 AM.
30 minutes of travel time are considered to be the crew's own time and shift paid time is calculated to begin at 8:30 AM, i.e. 30 minutes of waiting idle at the door is considered paid time.
 - If travel to the activity's location is estimated to take 50 minutes the scheduler plans for the crew to start driving at 8:00 AM.
30 minutes of travel time are considered to be the crew's own time and shift paid time is calculated to begin at 8:30 AM, i.e. 10 minutes of waiting idle at the door is considered paid time.
- When the time window for the first activity starts at 6 AM:

- If travel to the activity's location is estimated to take 10 minutes the scheduler plans for the crew to start driving at 7:50 AM.
10 minutes of travel time are considered to be the crew's own time and shift paid time is calculated to begin at 8:00 AM.
- If travel to the activity's location is estimated to take 30 minutes the scheduler plans for the crew to start driving at 7:30 AM.
30 minutes of travel time are considered to be the crew's own time and shift paid time is calculated to begin at 8:00 AM.
- If travel to the activity's location is estimated to take 50 minutes the scheduler plans for the crew to start driving at 7:30 AM.
30 minutes of travel time are considered to be the crew's own time and shift paid time is calculated to begin at 8:00 AM.

The crew is assumed to start their shift as usual before starting their travel from the logon location as well as complete their shift when they arrive to their logoff location.

Regardless of when the crew has actually started or completed their shift, the scheduler always applies the corresponding crew's own travel time to the actual start and completion times to calculate the paid shift start and end times respectively.

Time Zones

The time zone on a shift indicates the time zone in which the crew is planned to be working. All date and time information on the shift portal, including set up of breaks and non productive tasks for the shift, is entered and displayed in that time zone.

Crew Shift Actions

A dispatcher or a business process may perform various actions on a crew shift that may impact the shift's state and schedule.

The following section describes these actions and the business flows associated with them.

Allocate All

[Allocates](#) all of the activities currently scheduled to the shift that have not yet been dispatched.

NOTE: This action does not [close](#) the shift which means more work can be assigned to it by the scheduler.

Open

A user or any business process can open a [closed](#) or [standby](#) shift's schedule to allow the scheduler to assign more tasks to the shift, reassign work to another shift, and adjust the schedule as needed.

Close

By default and unless explicitly set up to be a [standby shift](#), the schedule on a shift is open for the scheduler to plan and optimize continuously. Closing a shift [allocates](#) all scheduled activities to it. The scheduler is prevented from removing any of the allocated activities or adding more activities to the shift.

When closing the shift you may also [manually dispatch](#) the allocated activities at the same time.

A shift can be closed in the following ways:

- A shift may be configured to automatically close (and queue for dispatch all tasks) some time prior to its planned start time. This is typically useful for crews that do not use a mobile device to work their schedule (non-MDT). Refer to [Advance Dispatching](#) for more information.
- A dispatcher can manually close a shift so as to finalize its schedule and prevent more work from being assigned to the crew. The dispatcher can re-open the shift at any time.
- If the crew goes offline, then all scheduled tasks are dispatched to the crew and the shift is closed. Refer to [Offline Mode](#) for more information.

Force Logoff

This action is only applicable to crews that utilize the mobile application to work their schedule.

When a crew uses the mobile application to start their shift, the system stamps the identifier of the device (MDT) they are using on the shift record. Message exchange between the crew and the system is assumed to be made to and from this device only.

This action forcibly logs off the association between the shift and the device currently being used to update it. This may be required if a logged on crew's mobile device is not working and they need to receive data on a new device. Once the shift is forced logged off, the crew can now access the shift from another device which then gets stamped on the shift as the current device being used by the crew.

NOTE: It is assumed that data on the old device is not recoverable. Any messages that may arrived stamped with the old device identifier would be discarded to avoid contradictory updates from multiple devices.

NOTE: This action solely disconnects a shift from a device and does not impact the status of the shift.

Planned Unavailability

There may be planned periods of time during a shift when the crew is not be available to perform work, for example, when they take their lunch break, stock their vehicle in the morning, or attend a meetings.

These non work related tasks consume time on a shift leaving less time to perform work. When planned in advance as part of the shift configuration, the system has a more accurate account of how much time remains available on a shift for work related tasks.

While performing these tasks, the crew is considered out of service and when they are completed, the crew is assumed in service again.

Refer to [Breaks](#), [Non-Productive Tasks \(NPTs\)](#) and [Meetings](#) for more information

Breaks

Crews may be entitled to take planned breaks after working some amount of time on their shift.

For example, a shift might include two 15 minute breaks, one in the morning and one in the afternoon. The morning break could be taken anytime from 1 to 2 hours after the crew has started their shift and the afternoon break could be taken anytime from 5 to 6 hours into the shift.

The following section describes key concepts related to planning and scheduling breaks.

Planned Relative To Shift Start Time

Breaks are planned to be scheduled during a time window starting at an offset relative to shift start time. Each break is a task associated with a break type. The duration of the break as well as the duration of the window during which the break

can be taken is defined on the break type, whereas the earliest time it can be taken, i.e. the offset from shift start time, is specified on the break task.

For example, if the offset is 2 hours and the shift starts at 8 AM, the break cannot start before 10 AM. If the shift starts at 9 AM with the same 2 hour offset, then the break task can be scheduled as early as 11 AM.

NOTE: Break tasks are only scheduled if enough work is scheduled to warrant a break from work.

No Location

Unlike other classes of tasks that take place at a specific location, a break does not have a defined location. The scheduler may schedule it to be taken anywhere on the way to a task location. A break may also be planned to be taken on site while working on a long duration activity, if permitted by the activity type.

If your company uses depots, the break may be scheduled to be taken at the depot while the crew is there to load or unload goods.

Creating Breaks

While you can manually add a break to an existing crew shift, breaks are typically incorporated into planned shifts at shift template design time.

Refer to [Planning Breaks](#) for more information.

Break Workflow

The following describes key stages in the overall lifecycle of a break.

Pending

A break is initially created in this state waiting to be scheduled and [queued for dispatch](#).

Queued For Dispatch

Depending on shift configuration, the crew may receive their entire schedule for the day or the first few tasks dispatched to them only when they start their shift (which they can preview in advance).

Refer to [Locking The Schedule](#) for more information.

Similar to an activity, when it is time to make the task available for the crew to preview or actually dispatch it to a crew's mobile device, the system queues it for dispatch so that its schedule is locked and changes can no longer be made once communicated to the crew.

Setting the break to the **Queued For Dispatch** state is the first step in the dispatch process.

The following business events queue breaks for dispatch

- If a crew's shift is configured to support [Advance Dispatching](#), the configured number of first tasks they may be previewed are queued for dispatch at a predefined time prior to the shift's planned start time.
- When a crew starts their shift, the scheduler automatically queues for dispatch (if not already queued in advance) their first few scheduled tasks, all of them or none, as per the shift's drip mode configuration. If a drip horizon is used, each time the crew starts or completes a task the scheduler proactively queues the next one for dispatch. Refer to [Controlling Auto-Dispatching](#) for more information

NOTE: Queuing a break task for dispatch should not be confused with the actual process of dispatching the task to the crew's mobile device. These are consecutive but separate steps.

Actual dispatching can only happen once the crew has started their shift and is ready to work. Moreover, the device the crew uses to work their schedule is only known once they use it to start their shift. Once known the device identifier is stamped on the started shift and the system directs all communications to the crew to that device. If the task was placed in the queue

in advance, prior to the crew starting their shift, then it can only be dispatched once the shift starts, otherwise it is dispatched right away.

NOTE: When conditions are met, the task would be automatically dispatched without relying on any background process to move it along. Even when the task is advance dispatched, it would automatically be dispatched when the shift starts.

When ready to be dispatched, if the crew uses a mobile device, a copy of the task is sent to the crew's device. If the crew uses SMS to work their tasks, a message with brief description of the task is sent to their mobile phone. Refer to [Keeping Mobile Copy Up To Date](#) for more information.

Refer to [Dispatching A Break](#) to better understand the big picture of the dispatch process.

Dispatched

The mobile application acknowledges receipt of the task by sending a message to the server to update the task to the dispatched state.

If the crew does not use a mobile device, the task is automatically marked dispatched.

Refer to [Dispatching A Break](#) for more information.

Started

This state indicates that the crew has started to take their break and are not working at the moment. While taking a break, the crew is considered out of service.

The crew provides this status update when they take this planned break. Task status updates by the crew are similar to assignment updates. Refer to the assignment [status update](#) for more information.

Time spent while in this state is considered break time for [crew time](#) reporting.

Completed

This state indicates that the crew is no longer at break and is back in service .

The crew provides this status update when they resume back to work. Task status updates by the crew are similar to assignment updates. Refer to the assignment [status update](#) for more information.

Canceled

A break in this state indicates it is no longer needed and should not be scheduled. Typically breaks are canceled before they are dispatched, but a dispatcher may decide to cancel a break at any time.

The crew may also decide to not take a planned break and cancel it. Task status updates by the crew are similar to assignment updates. Refer to the assignment [status update](#) for more information.

Once a shift is finalized all breaks that were not taken are automatically canceled.

Dispatching A Break

The process of dispatching a break to the crew involves multiple steps.

- **Step 1 - Placed in queue.** The break has to be placed in the queue for dispatch for the dispatch process to begin. This step is typically performed by the scheduler as it auto dispatches work, but other processes may do so as well.
- **Step 2 - Wait.** The break remains queued until it is actually dispatched and acknowledged (see step 4).
- **Step 3 - Dispatch.** Once all conditions are met for the break to be dispatched, the system sends a copy of the break task to the crew's mobile device, if any. Refer to [Keeping Mobile Copy Up To Date](#) for more information.

Refer to the [Queued For Dispatch](#) for more information about the processes involved in these steps.

- **Step 4 - Acknowledgement.** For crews using the mobile application the break is considered dispatched when the mobile application acknowledges receiving it. For all other crews a break is automatically acknowledged if the shift is started.

Refer to the [Dispatched](#) for more information about the processes involved in this step.

Keeping Mobile Copy Up To Date

If the crew is using the mobile application to manage their work, they receive a copy of each of their tasks sent to their mobile device from the moment they start their shift.

When key activity information changes about the task after it was dispatched, the system automatically synchronizes the mobile copy of the task and sends the crew a new copy of it with the latest information.

NOTE: Logic to initially send a copy of the task to the crew's mobile device as well as synchronize the mobile copy to keep it up to date when changes are made on the server resides in a "Synchronize Task" algorithm on the task's business object. Click [here](#) to see the algorithm types available for this system event.

Non-Productive Tasks (NPTs)

Crews may need to restock their vehicles or perform other administrative tasks at a specific location when they begin their shift or sometime after. These tasks are referred to as non productive tasks (NPT).

For example, a shift might start with a 30 minute task to replenish supplies on the truck from the service center in the morning. The task might take place anytime up to an hour at the beginning of the shift.

The following section describes key concepts related to planning and scheduling these tasks.

Planned Relative To Shift Start Time

Similar to breaks, NPTs are planned to be scheduled during a time window starting at an offset relative to shift start time. Each task is associated with a non productive task type. The duration of the task as well as the duration of the window during which it can be taken is defined on the task type, whereas the earliest time it can be taken, i.e. the offset from shift start time, is specified on the task itself.

For example, if the offset is 1 hour and the shift starts at 8 AM, the task cannot start before 9 AM. If the shift starts at 9 AM with the same 1 hour offset, then the break task can be scheduled as early as 10 AM.

Unlike breaks, NPTs are associated with a location where the task is performed. The location of an NPT may reference a common facility location or a specific address. A crew may be affiliated with a specific NPT [location](#) defined on the crew record. When creating an NPT, if its location is not explicitly specified, the system defaults it from the shift's crew record, if any is defined.

The scheduler needs the geocode of the task location in order to incorporate the task as part of the shift's route. If the task references an address and a geocode is not provided, the system attempts to geocode the address using the [geocoding service](#) defined on the installation options expecting an exact match.

NOTE: Logic to geocode the home address resides in an algorithm associated with the task's business object. Review this algorithms and adjust its match quality parameter as needed.

Creating NPTs

While you can manually add an NPT to an existing crew shift, these tasks are typically incorporated into planned shifts at shift template design time.

Refer to [Planning Non Productive Tasks](#) for more information.

NOTE: Non productive tasks cannot be created for [variable shifts](#).

NPT Workflow

The workflow of an NPT is similar to the workflow of a break, with the exception of handling the travel to the NPT location. As with other tasks associated with a location, the crew indicates they have started traveling to the NPT location and once there, they can start performing the task.

The following describes key stages in the overall lifecycle of an NPT.

Pending

An NPT is initially created in this state waiting to be scheduled and [queued for dispatch](#).

Queued For Dispatch

Setting the task to the **Queued For Dispatch** state is the first step in the dispatch process.

This is similar to the process of queuing breaks for dispatch. Refer to the break task [Queued For Dispatch](#) state for more information.

Dispatched

This is similar to the process of marking breaks as dispatched. Refer to the break task [Dispatched](#) state for more information.

En route

This state indicates that the crew is currently on its way to the task location. The crew provides this status update when they are about to start traveling to the task location. If the crew was previously in service, they are now considered out of service.

The crew provides this status update when they are about to start traveling to the location. Task status updates by the crew are similar to assignment updates. Refer to the assignment [status update](#) for more information.

Time spent while in this state is considered travel time for [crew time](#) reporting.

Started

This state indicates that the crew has started to perform the task associated with the NPT.

The crew provide this status update when they start the task. Task status updates by the crew are similar to assignment updates. Refer to the assignment [status update](#) for more information.

Time spent while in this state is considered NPT time for [crew time](#) reporting.

Completed

This is similar to the process of marking breaks as complete. Refer to the break task [Completed](#) state for more information.

Canceled

This is similar to the process of marking breaks as complete. Refer to the break task [Canceled](#) state for more information.

Meetings

Crews may need to attend work related events such as training sessions and meetings, also referred to as periods of unavailability (POUs). If the event overlaps with any of their planned shifts, the scheduler factors this event as part of the shift route.

For example, the crew can be out working activities, attend a training meeting and then return to work.

Refer to [Crew Meeting POUs](#) for more information

Locking The Schedule

Dispatching is an action that locks a task's sequence in the shift's schedule so no further optimization takes place, and marks it as ready to be sent to the crew.

This section describes shift configuration options related to dispatching.

Controlling Auto-Dispatching

Once a crew starts their shift, you can configure the scheduler to automatically dispatch to the crew all or a subset of their work without requiring dispatcher intervention.

This applies to any task, not just activities.

This is controlled by the shift's drip mode that supports the following options:

- **None.** Automatic dispatching is disabled. All scheduled tasks must be manually dispatched to the crew by the dispatcher or by another process.
- **All.** All scheduled tasks are automatically dispatched to the crew. Additional tasks scheduled throughout the day are also automatically dispatched. Activities for which automatic dispatch is **disabled** would still be marked as ready to dispatch by the scheduler, but they would rely on a dispatcher to manually dispatch them and place them in the dispatch queue.
- **Standard:** Only a fixed number of scheduled tasks are automatically dispatched to the crew at any given point of time. The number of tasks to dispatch at any point in time is referred to as the drip horizon. Activities for which automatic dispatch is **disabled** will not be assigned to shifts using this drip mode. Depot tasks do not count against the drip horizon. For example, if the shift is scheduled with a distribution depot task followed by an activity and the drip horizon set to 1 then both the depot task and the activity would be dispatched as part of the horizon.

The intent behind the more progressive standard drip mode is to dispatch just enough to keep the crew busy and allow the not yet dispatched portion of their route to be flexible enough to be rescheduled to other crews as exceptions arise throughout the day.

Depending on scheduler configurations, non work tasks, i.e. breaks, NPT and meeting POU tasks, may be automatically dispatched when a crew starts their shift, even if they are not yet within the dispatch horizon.

Refer to [Auto-Dispatching](#) for more information.

Advance Dispatching

Regardless of the automatic dispatching setting, you can also configure the system to queue for dispatch a number of tasks in advance of the shift start. This would ensure the schedule of these tasks is locked for various business needs. For example, it allows the crew to preview the first few tasks of their schedule before starting their shift.

It is also at this time that variable shifts become fixed shifts. Refer to [Variable Shifts](#) for more information.

The amount of time prior to shift start that the shift's schedule is locked is set by the advanced dispatch offset specified on the shift. Using the advance dispatch mode, you can also indicate that at the same time the shift should be closed and all the tasks scheduled to it should be queued for dispatch.

Closing the shift and queuing for dispatch all its scheduled tasks is also referred to as the "shift cutoff".

NOTE: This functionality relies on the **Crew Shift Monitor** batch process to be scheduled in appropriate frequency. Refer to [MI-CRSHF](#) batch control for more information.

Crew Undispatch and Sequence Lock

Some business practices require the dispatched route to be completely locked, while others may allow it to change for a significantly more optimal route and respond to real-time exceptions.

The system supports the ability to configure a shift to allow or disallow route changes in the form of undischarging tasks or changing the order of the dispatched tasks with the following options:

- Not allow undispach at all, allow it only for excess tasks dispatched to the shift beyond the horizon, or allow it always.
- Not allow sequence change at all or allow it always.

NOTE: When allowed, the scheduler would only consider making these changes if the cost of the route change does not exceed the **Mobile Update Cost** cost control threshold. The scheduler does not undispach activities that are manually allocated to the shift or advance dispatched.

Offline Mode

Some crews may use no mobile device or may operate in areas with poor or no connectivity coverage. These crews don't have the means to automatically receive new work nor changes to their route.

In this situation, there is no point in keeping their schedule open. These shifts needs to be closed with their entire schedule locked before the crew is out in the field with no means of connection.

Closing the shift and locking its schedule ensures the route remains unchanged. But in this situation, since the crew cannot provide status updates as to their progress, the scheduler does not have a way to validate whether or not the crew is on schedule. This in turn may result in false alerts being raised, for example, reporting that the crew is late for an appointment when they actually made it on time but could not or did not report it promptly.

To prevent scheduling discrepancies, the crew's shift has to not only be closed but also set to offline mode. In this mode, the scheduler assumes the crew is on track and does not push back the schedule. Also, base product alerts are designed to ignore shifts in offline mode.

NOTE: A shift has to be closed before it can switch to offline mode.

Shifts for crews that do not use a device are typically configured to close some time before the shift's planned start time. When the crew starts the shift, or a dispatcher does it on their behalf, the system automatically switches the shift to offline mode.

Refer to [Using No Device](#) for more information.

Crews using the mobile application may explicitly request to switch to offline mode before they travel to an area with poor connectivity for a long period of time. The request closes their shift, dispatches their entire schedule to the mobile device, and switches to offline mode. Note that closing the shift automatically allocates all scheduled activities to that shift and they remain allocated for the duration of the shift. Refer to the [Close](#) shift action for more information. The crew can request to switch back to online mode again, as needed.

Refer to [Switch To Offline Mode](#) for more information.

Working The Schedule

Crews may use mobile devices to carry out their scheduled work and provide real time updates as to their progress. Real time updates from the crew increases the accuracy of the information available to the dispatcher as well as the scheduler when making scheduling decisions and handling exceptions.

Crews may use the mobile application if installed on their device or provide simpler status updates using SMS messaging.

Crews that do not use a mobile device typically record their work on paper or other means, and enter the information via the main application at the end of their shift.

This section describes topics related to performing the work scheduled to a crew's shift.

Using The Mobile Application

The system supports dispatching work directly to mobile application installed on mobile devices including laptops, tablets and other devices.

The mobile application allows the crew to fully manage their work schedule, get more work as needed, handle exceptions as well as seamlessly report their whereabouts.

When the crew starts the shift from their mobile device key details about their shift are brought over and stored locally. If these details change later on the system automatically synchronizes the mobile copy of the shift and sends the crew a new copy of it with the latest information.

NOTE: Logic to initially send a copy of the shift to the crew's mobile device as well as synchronize the mobile copy to keep it up to date when changes are made on the server resides in a "Synchronize Shift" algorithm on the task's business object. Click [here](#) to see the algorithm types available for this system event.

Refer to [About The Crew's Application](#) for more information

Using SMS Messages

Crews may use SMS messaging to carry out their scheduled work. This method allows the crew to start their shift, get their tasks one at a time, provide status updates and complete their shift.

The following scenario illustrates how SMS messaging works

- A mobile worker sends a text message of "Logon". The system identifies the mobile worker via the [phone number](#) the text was sent from, starts their next shift, and sends details of their next task via SMS.
- As the crew works the task, they send messages to communicate their progress, such as "1" to indicate that they are en route to the task's location, and "2" to indicate arrival at the task location and starting to perform the work.
- When work is done, the crew sends a text message of "3". In response, the system sends the next task and so on. If this was an activity, this tells the system work is done but still waits for completion information to be keyed in. The crew would typically enter the completion details at the end of their shift when they return to the office. Refer to [Work Done](#) assignment state for more information.
- At the end of the shift, the crew sends a "Logoff " message.

The following are examples of the messages sent to crews when tasks are dispatched.

```
Activity Details MREAD at 12.00.00 for Swetha Paspuleti at 6284 dressler Road NW, North Canton
POU Details STATUS MEETING at 6284 dressler Road NW, North Canton from 14.00.00 to 14.30.00
Break Details COFFEE from 13.00.00 to 13.30.00
```

NOTE:

Logic to process incoming SMS messages from the crew and send messages back to the crew resides in algorithms on the installation options record.

Click [here](#) to see the SMS Receive algorithm types available for this system event.

Click [here](#) to see the SMS Send algorithm types available for this system event.

Using No Device

Crews that do not use any device to manage their work may be provided with a printout of their scheduled route at the beginning of their shift, or they may rely on the dispatcher to communicate the details of their next task verbally as they work them one at a time.

NOTE: Functionality around handling a shift's schedule printout is not provided by the base product. If needed, this is assumed to be a customization effort.

Since they don't have a way to provide real time updates from the field, crews typically rely on communicating them directly to the dispatcher who performs these changes on their behalf. For example, the dispatcher can start their shift, update the status of their next task to state they went en route to the task location, update it to note they have arrived, as well as mark the [work done](#), etc.

Throughout the day, the crew records their completion information on paper or other means and when back at the office, they use the main application to enter the data and complete the corresponding activities.

As the crew does not have a device to automatically receive new work nor changes to their route, there is no point in keeping their schedule open. These shifts therefore are typically configured to close and lock the entire schedule some time before they are planned to start the shift.

Refer to [Advance Dispatching](#) for more information.

If the shift is closed in advance, when the crew starts the shift or the dispatcher starts it on their behalf, the shift is automatically switched to offline mode.

Refer to [Offline Mode](#) for more information.

Auto-Completion

By default, the system assumes that work scheduled to a shift that was never started was not worked and therefore needs to be rescheduled. Expired shifts are automatically canceled and all the work scheduled to them is rescheduled.

However, some customers use the product purely for scheduling purposes where activities and shifts are never updated by the crew. No completion information is recorded in the system to be sent back to the host.

To accommodate this business use case, shifts can be configured to automatically complete as per their planned schedule assuming all scheduled tasks are worked as planned. If a shift is set up to automatically complete, the system starts it at the planned start time, progresses all its scheduled tasks as planned, and finally automatically completes it at the planned end time. Note that if by the planned start time all mobile workers allocated to the shift are on leave then the auto-completion process cancels the shift instead of starting it.

If there is an exception in working an activity in the field, it is assumed to be handled directly in the host system and resent as needed.

NOTE: The **Shift Auto Complete**[global configuration](#) parameter needs to be enabled before you can set up shifts to auto-complete.

NOTE: This functionality requires the shift record and the auto-completion sync request that manages its over time to be periodically monitored. The **Crew Shift Monitor** batch process initiates the auto-completion sync request when the shift is planned to start and **Sync Request Monitor** batch process advances the shift and its scheduled tasks as planned. It is important that these batch processes are ran in a timely manner. Refer to [M1-CRSHF](#) and [F1-SYNRQ](#) batch control for more information.

Working One Task At A Time

At any point in time during their shift, the crew is assumed to be working on a single task. Attempting to start a task when another is still in progress is not allowed. In the same way, a shift cannot be completed while a task is still in progress.

End of Shift

When the crew completes their shift the following steps are involved:

- Activity assignments the crew did not manage to complete are returned and the corresponding activities are rescheduled.
- All other types of tasks that were not completed are canceled.

- Allocated activities no longer remain allocated to the shift, allowing them to be rescheduled to other shifts.
- If timesheet functionality is enabled, the timesheet records are prepared for each period a mobile worker worked as part of the crew on the shift. Refer to [About Timesheets](#) for more information.

About Shift Weekly Templates

While crew shifts can always be manually created to address a specific situation, they are typically planned and generated ahead of time based on recurring templates. Templates provide a tool for quickly generating large numbers of shifts based on predefined parameters.

A shift weekly template defines a cyclical weekly pattern made up of weekday shift templates.

For example, you might define a weekly template in which the crew works its regular 9-5 shift Monday -Friday every week and also works an on-call shift on Saturday and Sunday every third week. Thus the shift weekly template would rotate every 3 weeks.

Resource planners can create shift weekly templates for specific crews, or they can create common templates and subscribe multiple crews to one template.

This section describes concepts related to setting up and maintaining shift weekly templates.

A Weekly Template Is Made Of Shift Templates

A shift weekly template is made of one or more shift templates that describe the recurring shift patterns. The details for each shift within a cycle are captured on a shift template record. For example, a weekly template might use a different shift template on Fridays than the one used on other days.

This section describes concepts related to setting up and maintaining shift templates.

Actual vs. Template Shifts

A shift template defines the details needed to generate actual shifts based on that template. Any detail you can define for an actual shift, you can also predefine on a shift template.

To enforce the close pairing between a shift template and the shifts it creates, both the actual and the template records are stored in the same table so they can share similar details. The same crew shift type defines both the business object used to create actual shifts of this type as well as the business object used to create shift templates of that type.

Refer to [Defining Shift Types](#) for more information.

A generated shift maintains a reference to the shift template that was used to generate it so that changes made to a template may be reapplied to those shifts.

Complete vs. Partial Shift Details

If the shift weekly template is dedicated to a [specific](#) crew, each shift template may include the complete shift information needed to generate a shift.

A shift weekly template may be defined as a common template to be [shared](#) by multiple crews. As such, details unique to each crew such as the crew and the resources allocated to it cannot be specified on the common template as they differ for each crew. The crew specific details are taken from the crew record at shift generation time.

Refer to [Shift Generation](#) for more information.

Planning Breaks

When you define shift templates, you can plan one or more break types for the shifts. When the shifts are generated, the system automatically creates breaks and incorporates them into the crew shift schedule.

Please refer to [Breaks](#) for more information.

Planning Non Productive Tasks

When you define crew shift templates, you can plan one or more non productive task types for the shifts. When the shifts are generated, the system automatically creates non productive tasks and incorporates them into the crew shift schedule.

The location of the non productive task is taken from the crew record at shift generation time.

Please refer to [Non Productive Tasks](#) for more information.

Crew Specific Weekly Template

Crew specific weekly templates are dedicated to a specific crew and cannot be shared with other crews. As such, each weekday shift template is fully defined with everything that the shift generator needs to generate a shift for this crew.

The initial date of the cyclical weekly pattern is the effective date of the template.

This type of weekly template references the crew it is designed for, as well as the resources allocated to the crew, on each shift template. This allows for more complex situations where different resources are allocated to the same crew on different days of the week.

This section describes concepts related to setting up and maintaining crew specific weekly templates.

Defining and Applying A Template

When you create a new shift weekly template, its status is **Pending**. No shifts are generated while the template is in this state allowing you some time to properly define the weekday shift templates and review them as needed.

After you have defined the shift templates, you must apply the weekly template so that actual shifts will be generated from it, based on the template's effective period. Applying the template changes its status to **Applied** and the system right away generates future shifts starting from the template's effective date (or from now if that date is in the past) and until the template's expiration date (if any).

NOTE: Generating shifts online is limited to a preconfigured number of days in order to avoid the apply action from taking too long and time out. The parameter is defined on the shift generation algorithm.

Once applied, the shift generation background process manages the generation of shifts going forward.

Refer to [Shift Generation](#) for more information about the generation algorithm as well as the overall process.

Making Changes To An Applied Template

Extending or shortening the effective period of an applied template adjusts future shifts accordingly as soon as you make the change.

If you attempt to add, delete, or modify a shift template for an applied weekly template, the system automatically creates a draft copy of the entire weekly template it belongs to where your change is made to the draft copy, leaving the original template unchanged. This allows you to make changes in a more controlled way.

You make changes to the draft copy and preview the shifts that would be generated should you accept the draft. If you decide to accept the draft, the original template is updated to match the draft and changes are immediately applied to already generated future shifts accordingly.

Refer to [Draft and Original Templates](#) for more information.

Common Weekly Templates

You can set up common weekly templates which can be shared across multiple crews. This is helpful for the efficient generation of shifts in organizations where there are sometimes hundreds or possibly even thousands of crews that follow the same weekly shift pattern.

This section describes concepts related to setting up and maintaining common weekly templates.

Subscription

The common template defines the initial date of the cyclical weekly pattern as well as the details for each shift template. These shift templates cannot include crew specific information, such as the resources allocated to the crew on that shift etc, as they need to be shared by multiple crews.

You may subscribe multiple crews to a common template as their template for generating shifts. Shifts generated from the common template are exactly the same except for crew specific details [defaulted](#) from the crew record such as allocated resources. Any crew that subscribes to a common template must have these types of details defined on the crew record.

Subscription to a common template is effective dated so that crews can subscribe to follow a common template as of a given date and this relationship can expire. Each crew can subscribe at different times, but their shifts are generated in the same rotation timing as all other subscribed resources, regardless of when they subscribed.

You may also unsubscribe a crew from the common template at any time or change effective dates to affect the creation of future shifts for those crews.

Activation

In a typical workflow, you would create a common weekly template and then subscribe multiple crews to the template. When template and subscriptions are ready to be used to generate shifts, you activate the template. Shifts are generated only when the common template is active.

Since a common template may involve a large number of subscriptions, actual shift generation is performed the next time the shift generation batch process runs. Note that you may still test your template and process it online for a subscribed crew without waiting for the next scheduled batch run.

Refer to [Shift Generation](#) for more information.

Making Changes To An Active Template

Changes made to an already active common template affect future shifts of all subscribed crews. Since a common template may involve a large number of subscriptions, such changes are not applied right away to all these shifts but rather they take effect the next time the shift generation batch process runs. This includes extending or shortening the effective period of an active template, subscribing or unsubscribing a crew etc.

If you want to test your recent changes without waiting for the next scheduled batch run you can do so by processing them online for one of the subscribed crews.

Similar to crew specific templates, if you attempt to add, delete, or modify a shift template for an activity weekly template, the system automatically creates a draft copy of the entire weekly template, leaving the original template unchanged.

This allows you to make changes in a more controlled way. You make changes to the draft copy and preview the shifts that would be generated should you accept the draft. If you decide to accept the draft, the original template is updated to match the draft. However, unlike crew specific templates, the impact on these changes on future shifts only takes effect the next time the shift generation batch process runs.

Refer to [Draft and Original Templates](#) for more information.

One Effective Shift Weekly Template For A Crew

A crew's specific weekly templates and subscriptions to common templates are either effective for a specified number of weeks starting from their effective date, or they may be perpetual, which means they have an effective date but no end date.

The crew can switch from a specific weekly template designed just for this crew or a subscribe to a common template, and vice versa, over time, but their shifts may be generated by a single weekly template at any given point in time. In other words, a crew's weekly templates must never overlap in time.

Shift Generation

The system generates future shifts for crews based on the weekly templates with which they are associated.

Common templates do not generate shifts as they are not defined for specific crews. Only a [crew specific](#) weekly template and a crew's [subscription](#) to a common template are used to generate shifts for the crew throughout their effective period. Unlike a crew specific weekly template that is self contained, a crew's subscription to a common template sets the shift generation effective period, whereas template details as well as the initial rotation date are defined on the common template.

Shifts are generated based on the cyclical pattern of the weekly template, starting from the weekly template's effective date (or from now if in the effective date is in the past) and up to a predefined number of days known as the planning horizon, or earlier, if the template expires before then.

The planning horizon is defined as a [global configuration](#) parameter and is used by the shift generation process to control the maximum number of days of shifts to generate and have available for planning purposes, at any point in time.

If the template references a work calendar, the process skips generating shifts that even partially overlap with a public holiday.

The first few shifts may be initially generated online when you apply a crew specific template or verify the setup for a crew's subscription to a common template, but the ongoing process of generating additional shifts to meet the planning horizon is handled by the weekly template monitor process.

NOTE: The shift generation process relies on the shift weekly template being periodically monitored by the **Shift Weekly Template Monitor** batch process as often as appropriate to meet your organization's needs. Refer to [M1-SHWKT](#) batch control for more information.

The background process is designed to add new shifts as needed on a daily basis. However, if any shift template was changed, the process would automatically regenerate future shifts that were earlier generated by that template to match its latest version. By default, even shifts that were manually overridden are regenerated to match the latest template version, but you may explicitly indicate to retain these changes when making a shift template change.

Refer to [Overridden Shift Details](#) for more information.

When a shift is about to be generated but happens to overlap with another shift for the crew that was not generated by this template, the generation of the shift is skipped and a user is notified by raising a to do entry that reports the situation.

NOTE: Logic used to generate shifts from a shift template resides in a "Generate Shifts" algorithm on the shift template's business object. Click [here](#) to see the algorithm types available for this system event.

Overridden Shift Details

In a situation where crew shift details differ from their template on a particular date, generate the shift initially from the template and then you can manually adjust it to override the template settings. This includes canceling a shift that is no longer needed. These modified shifts are called overridden shifts as they no longer match the template from which they were generated.

By default, any changes made to a shift template are automatically reapplied to all future planned but not yet closed shifts that were generated from that template, and any manual overrides will be lost. If you wish to retain overridden shifts you may indicate so on the shift template being changed.

Draft and Original Templates

If you attempt to add, delete, or modify a shift template for a shift weekly template that has been applied (or activated for common templates), the system automatically creates a draft copy of the entire weekly template. This allows you to make your changes to the draft and preview the shifts that would be generated if you accepted the draft.

You can also see if any manual edits were made to previously generated shifts. By default, such overrides are lost when the shift template that generated them is changed when a draft is accepted. You may explicitly request to retained overrides when making the shift template change.

If you have reviewed this information and are sure you want to make the proposed changes, you may accept the draft. Otherwise you may delete the draft copy to discard the proposed changes.

When a draft template is accepted, the original template is updated to match the draft

- If a shift template was changed, future shifts generated by this template are regenerated to match the changed template.
- If a shift template was deleted by the draft, the corresponding original template is also deleted if not yet used to generate shifts, otherwise it is marked inactive. Future shifts generated by this template are deleted.
- New shifts are generated for any shift templates that were added by the draft.

By default, changes made to a shift template impacts all future planned but not yet closed shifts generated by it, including overridden templates.

NOTE: If the original weekly template is a crew specific template, changes are reapplied to future shifts right away when you accept the draft. If the original template is a common template, the impact is delayed to the next time the shift generation batch process runs.

Refer to [Overridden Shift Details](#) for more information.

Once a draft has been accepted, any further changes made to the original template will result in a new draft being created.

Maintaining Crew Shifts

This portal is used to display and maintain a crew's shift.

Refer to [About Crew Shifts](#) for more information.

You can access the portal from the **Main Menu > Resource Management > Crew Shift**. You are brought to a query portal with options for searching for a specific crew shift. Once a shift has been selected you are brought to the maintenance portal to view and maintain the selected record.

The following zones may appear as part of the portal's **Main** tab page

- **Crew Shift.** This zone provides information about the shift, including the various maintenance actions applicable to its current state.
- **Breaks / Non Productive Tasks.** This zone defines breaks and NPT tasks for the shift.
- **Timesheets.** This zone lists timesheet records for resources allocated to the shift.

The following zones may appear as part of the portal's **Schedule** tab page

- **Crew Shift Gantt.** This zone displays a graphical representation of the tasks scheduled to this shift. The Gantt works in the same way as the dispatcher's scheduling Gantt on the CDI portal with the exception that filters and the ability to launch the geographic map are not applicable from this zone. Refer to [Using the Scheduling Gantt](#) for more information.

- **Scheduled Tasks.** This zone displays the list of tasks scheduled to this shift in their work order.
- **Un-Scheduled Assignments.** This zone displays activities that were once dispatched to the shift but were since recalled.
- **Planned Procedures.** This zone lists procedure types that the crew is planned to fill out when they would work this shift. The zone is only visible when the shift has not started yet.
- **Procedures.** This zone lists procedures completed by the crew during this shift. The zone is visible when at least one procedure was completed by the crew on that shift.

The **Route** zone on the portal's **Route** tab page displays on a geographic map the planned route of the crew's schedule vs. the actual route the crew took on a specific shift.

The **Capabilities** zone on the portal's **Capabilities** tab page displays the crew's overall capabilities from all mobile workers and vehicles currently allocated to the shift.

Recommended Crew Allocation

This portal is used to assist in allocating resources when defining a shift.

Refer to [About Crew Shifts](#) for more information.

You can access the portal from the **Maintain Crew Allocation** link on the [crew shift maintenance](#) portal when the shift has not started yet.

The following zones may appear as part of the portal's **Main** tab page:

- **Crew Shift.** This zone provides information about the shift, including the various maintenance actions applicable to its current state.
- **Recommended Mobile Worker.** This zone allows you to search and allocate appropriate mobile workers to the shift.
- **Recommended Vehicle.** This zone allows you to search and allocate appropriate vehicles to the shift.

Maintaining Breaks / Non Productive Tasks

This portal is used to display and maintain a break task or a non productive task.

Refer to [Breaks](#) and [Non-Productive Tasks \(NPTs\)](#) for more information.

You can access the portal from the Breaks / Non Productive Tasks zone on the [crew shift maintenance](#) portal.

The **Break / Non Productive Task** zone on the portal's **Main** tab page provides information about the task, including the various maintenance actions applicable to its current state.

Maintaining Shift Weekly Templates

This portal is used to display and maintain a shift weekly template. You may create a common weekly template to which multiple crews may subscribe or you may create a weekly template for a specific crew.

Refer to [About Shift Weekly Templates](#) for more information.

You can access the portal from the **Main Menu > Resource Management > Shift Weekly Template**. You are brought to a query portal with options for searching for a specific shift weekly template. Once a shift weekly template has been selected, you are brought to the maintenance portal to view and maintain the selected record.

The following zones may appear as part of the portal's **Main** tab page:

- **Shift Weekly Template.** This zone provides information about the shift weekly template, including the various maintenance actions applicable to its current state.

- **Shift Templates.** This zone defines the individual shift templates that comprise the rotation cycle for this weekly template.
- **Subscriptions.** This zone is used to maintain crew subscriptions to the common weekly template. The zone is only visible if this is a common weekly template.
- **New Subscriptions.** This zone assists in searching for and subscribing crews to the common weekly template. The zone is only visible if this is a common weekly template.
- **Inactive Shift Templates.** This zone lists shift templates that are no longer active and is visible if inactive shift templates exist for the weekly template.
- **Overridden Shifts.** This zone lists shifts that were generated by this template but were since modified manually and no longer match the template from which they were generated. The zone is visible if such shifts exist for the weekly template.

The following zones may appear as part of the portal's **Review** tab page

- **Preview Calendar.** This zone allows you to review a month's worth of potential shifts that would be generated if you applied the recent changes to the template. You may broadcast a specific shift to display its details in the crew shift preview zone.
- **Crew Shift Preview.** This zone displays a shift broadcasted from the preview calendar zone.

Refer to the Calendar (F1-CALENDAR) zone type for more information.

Comparing Draft and Original Templates

When you make changes to a template that has already been applied, the system automatically creates a new draft version of the template. It puts your changes in the draft version and leaves the original applied version unchanged until you manually apply the draft. Before you apply a draft template, you should view the differences between the draft and the original to make sure you want all changes applied.

To view differences, click the **Show Differences** button of the Shift Weekly Template portal. The system displays a pop-up window listing each template that has been changed and the type of change (Deleted, Updated, or New).

You may also want to preview the shifts that will be generated if the draft is applied using the **Review** tab.

Navigating Between Original and Draft Templates

When the shift weekly template you are working with has an original version and a pending draft, you can navigate between them as follows:

- When viewing a draft template, click the **Original Template ID** link to navigate to the original template's maintenance portal.
- When viewing an original template, click the **Pending Draft Exists** link to navigate to the draft's maintenance portal.

Maintaining Shift Templates

This portal is used to view and maintain a shift template as part of setting up a shift weekly template.

You can access the portal from the shift templates zone on the [Maintaining Shift Weekly Templates](#) portal.

Refer to [About Shift Weekly Templates](#) for more information.

The following zones may appear as part of the portal's **Main** tab page:

- **Crew Shift Template.** This zone provides information about the shift, including the various maintenance actions applicable to its current state.

- **Overridden Shifts.** This zone shows shifts that were generated by this template but their setup has since been manually changed from the setup defined by the template.

The **Capabilities** zone on the portal's **Capabilities** tab page displays the crew's overall capabilities from all mobile workers and vehicles currently allocated to the shift based on this template.

Periods of Unavailability (POUs)

About Periods of Unavailability (POUs)

Resources may have planned times during which they are not available to perform work.

For example, mobile workers or dispatchers may take planned vacations, have personal appointments, attend meetings or training sessions, etc. Company vehicles may be unavailable for maintenance reasons. Crews may need to attend meetings or participate in training sessions and hence not be available to perform work during these events.

You may use a POU record to advise the system when resources are planned to be unavailable. Knowing upfront about these periods of unavailability allows for better planning and scheduling of work during those times.

This section describes concepts related to managing POUs.

Different Types of Events

Each POU references a **POU Type** which defines attributes common to all POUs of that type. For example, a status meeting type of event may only be applicable to crews while a vacation type of event is only applicable to mobile workers.

Leave POUs

You may use a POU to capture an individual's leave time such as when they report their planned vacation times or a vehicle's maintenance periods.

If a leave POU has been entered for a mobile worker or a vehicle, the system also takes into account whether or not the leave period overlaps with any of their planned shifts and factors this in when determining whether the resource is available to be scheduled for a shift.

If the leave overlaps with the shift for any amount of time, the resource is considered unavailable, even if the overlap is only for a very short duration of time. In this situation, the shift is resent to the scheduler with its capabilities adjusted to not include those possessed by the resource on leave.

Note that a shift requires at least one mobile worker for it to be properly scheduled. If all mobile workers on a shift are on leave, the shift is automatically disabled by the scheduler and no work is scheduled to it as there is no one to perform the work. If there is a change in the resources allocated to the shift or if a previously planned POU is canceled or moved to a different time, the shift would be automatically scheduled as usual.

Leave time periods are displayed on the resource's calendar so that it is easy for the resource planner to determine their availability.

Crew Meeting POUs

Crews may need to attend a meeting or a training session that takes place at a particular time during their shift hours. When they do, their schedule for the day is adjusted to take into account the time needed for them to participate in this event. In the same way, the time it would take them to travel to and from the event location is incorporated into their scheduled route.

You may use a POU record to capture the details of such event as well as the list of invited crews. This section describes concepts related to crew related POU.

Service Areas

In addition to the event time, duration, location, and the list of invited crews, the event may optionally be restricted to crews working in specific service areas.

This instructs the scheduler to incorporate the POU in the invited crews' shifts only if they are planned to work in any of the specified service areas that day. This is used to discourage the incorporation of the POU event into a crew's schedule if the crew happens to work too far from the event that day.

If no service areas are specified, there is no restriction around a crew's proximity to the event that day.

Tracking Crew Attendance

While a single POU record captures the meeting event, the system has to track the attendance of each of the invited crews individually.

The system automatically creates a POU task for every invited crew and then incorporates this task into the crew's shift schedule. The system uses the POU task type defined on the POU type record to create the POU task and if not specified a default task type is used from the global master configuration.

A POU task is dispatched to a crew, just like an activity, and its status is updated and tracked by the crew in the same way.

POU Task Workflow

Similar to other tasks associated with a location, the crew indicates they have started traveling to the event location, arrived at the location, as well as when they leave (complete) the meeting.

The following describes key stages in the overall lifecycle of a POU task.

Pending

A POU task is initially created in this state waiting to be scheduled and [queued for dispatch](#).

Queued For Dispatch

Setting the task to the **Queued For Dispatch** state is the first step in the dispatch process.

This is similar to the process of queuing breaks for dispatch. Refer to the break task [Queued For Dispatch](#) state for more information.

Dispatched

This is similar to the process of marking breaks as dispatched. Refer to the break task [Dispatched](#) state for more information.

En route

This state indicates that the crew is currently on its way to the event location. The crew provides this status update when they are about to start traveling to the location. If the crew was previously in service they are now considered out of service.

The crew provides this status update when they are about to start traveling to the location. Task status updates by the crew are similar to assignment updates. Refer to the assignment [status update](#) for more information.

Time spent while in this state is considered travel time for [crew time](#) reporting.

Started

This state indicates that the crew has arrived to the event location.

The crew provides this status update when they arrive at the location. Task status updates by the crew are similar to assignment updates. Refer to the assignment [status update](#) for more information.

Time spent while in this state is considered POU time for [crew time](#) reporting.

Completed

This state indicates that the crew has left the event.

The crew provides this status update when they leave the location (and thus complete the event). Task status updates by the crew are similar to assignment updates. Refer to the assignment [status update](#) for more information.

Canceled

This state indicates the crew is not going to attend the event and the POU is removed from the crew's schedule. A POU task may be canceled when the event itself is canceled or rescheduled to another time, or a dispatcher may have canceled the participation of the crew in the event. In either situations, if the crew is using a mobile device, the cancelation is communicated to the crew.

The crew may also explicitly decline their attendance of the event by canceling the task. Task status updates by the crew are similar to assignment updates. Refer to the assignment [status update](#) for more information.

Once a shift is finalized, tasks for events the crew did not attend are automatically canceled.

Time Zone

All date and time information on the POU portal is entered and displayed in the event location's time zone.

Creating POUs

Depending on the nature of the unavailability periods, they may be recurring or one time events. You may enter a POU to record a one time event or set up a POU template to handle a recurring event.

Refer to [About POU Templates](#) for more information.

About POU Templates

An unavailability event may either be recurring or a one time event. For a recurring event, you define a POU template from which the actual POUs are generated. For example, if you create a template for a regular weekly staff meeting, the system will automatically generate individual POUs for each meeting.

Reoccurrence may not be applicable to all types of events. For example, an individual's vacation time does not tend to repeat in a pattern.

Whether or not a POU can be recurring is configured on the [POU Type](#). If supported, then in addition to a one time POU event, you may also create a POU template to initiate a recurring event of this type.

In addition to the event details and invitees, a template includes information about the event's recurrence pattern and expiry date that the system uses to generate the recurring POU events.

This section describes concepts related to setting up and managing POU templates.

Actual vs. Template POUs

A POU template defines the details needed to generate actual POUs based on that template. Any detail you can define for an actual POU, you can also predefine on a POU template.

To enforce the close pairing between a POU template and the POUs it creates, both the actual and the template records are stored in the same table so they can share similar details. The same POU type defines both the business object used to create actual POUs of this type as well as the business object used to create POU templates of that type.

Refer to [Defining POU Types](#) for more information.

A generated POU maintains a reference to the template that generated it, so any changes made to a template are automatically applied to all future POUs generated from that template. If you modify a POU template after POUs have been generated from it, the system notifies you that all previously generated POUs will be deleted and regenerated, and that any manual overrides will be lost.

POU Generation

POUs are generated based on the recurring pattern defined on a POU template starting from the template's effective date (or from now if in the effective date is in the past) and up to a predefined number of days known as the planning horizon, or earlier if the template expires before then.

The planning horizon is defined as a [global configuration](#) parameter and is used by the POU generation process to control the maximum number of days of POUs it generates and have available for planning purposes, at any point in time.

The system supports the following patterns for a recurring POU event:

- **Daily.** The POU can recur every x days.
- **Weekly.** The POU can recur on a specific day of the week every x weeks. If you want the POU to occur on Tuesday and Thursday every week, create two POU templates: one that occurs on Tuesday every week and one that occurs on Thursday every week.
- **Monthly.** The POU can recur every x months on an exact date or a relative date (a day in a specific week). For example, the POU can occur on the 21st of every month or the last Friday every three months.
- **Yearly.** The POU can recur every x years on an exact date or a relative date (a day in a specific week in a specific month). For example, the POU can occur the 4th of July every year or the first Monday of January every two years.

If the template references a work calendar, the process skips generating POUs that fall on a non-workday or a holiday as defined by the calendar.

You may specify either an explicit expiration date or request a specific number of occurrences to be generated, after which the template would automatically expire. The number of occurrences defined on a POU template only counts actual POUs that are generated. For example, a POU template is defined as the first Monday of every month for 10 occurrences starting from the 1st of January. If two of those Mondays are holidays, the POUs will be created through December, skipping the two holiday days.

When a POU template is initially saved, the system generates POUs from the template automatically. The system only generates POUs to meet the planning horizon. The Period of Unavailability Monitor is responsible for generating more POUs going forward.

NOTE: The POU generation process relies on the POU template being periodically monitored by the **Period Of Unavailability Monitor** batch process as often as appropriate to meet your organization's needs. Refer to [M1-POU](#) batch control for more information.

On a daily basis, the background process is designed to add new POUs as needed. However, if any POU template was changed the process would automatically regenerate all future POUs that were generated by that template to match its latest version.

Refer to [Overridden POU Details](#) for more information.

NOTE: Logic used to generate POUs from a POU template resides in a "Generate POUs" algorithm on the POU template's business object. Click [here](#) to see the algorithm types available for this system event.

Overridden POU Details

In a situation where POU details differ from their template on a particular date, generate the POU initially from the template and then you can manually adjust it to override the template settings. These modified POUs are called overridden POUs, as they no longer match the template from which they were generated. The Overridden POUs zone on the POU Template portal displays a list of overridden POUs linked to the current template.

A generated POU maintains a reference to the template that generated it, so any changes made to a template are automatically applied to all future POUs that were generated from that template. If you modify a POU template after POUs have been generated from it, all previously generated POUs will be regenerated, and any manual overrides will be lost.

Maintaining Periods Of Unavailability

This portal is used to display and maintain a recurring or non recurring period of unavailability (POU). A recurring POU is also referred to as a POU template.

Refer to [About Periods of Unavailability \(POUs\)](#) for more information.

You can access the portal from the **Main Menu > Resource Management > Period Of Unavailability**. You are brought to a query portal with options for searching for a specific POU. Once a POU has been selected you are brought to the maintenance portal to view and maintain the selected record.

The following zones may appear as part of the portal's **Main** tab page

- **Period Of Unavailability**. This zone provides information about the POU, including the various maintenance actions applicable to its current state.
- **Invited Crews**. This zone maintains the list of invited crews for a meeting POU. The zone is only visible for a crew related POU.
- **Crew Search**. This zone allows you to search and invite crews to a meeting POU. The zone is only visible for a crew related POU.
- **Overridden Periods Of Unavailability**. This zone lists POUs that were generated by this template but were since modified manually and no longer matches the template from which they were generated. The zone is visible only for a POU template for which such POUs exist.

Maintaining Period Of Unavailability Tasks

This portal is used to display and maintain a POU task. Such task exists for every crew invited to attend a crews meeting event.

Refer to [Crew Meeting POUs](#) for more information.

You can access the portal from the Invited Crews zone on the [Maintaining Periods Of Unavailability](#) portal.

The **Period Of Unavailability Task** zone on the portal's **Main** tab page provides information about the task, including the various maintenance actions applicable to its current state.

Dispatchers

About Dispatchers

A dispatcher is a person within your organization who manages day-to-day field operations. Dispatchers are responsible for monitoring crews and activities within certain dispatch areas.

This section describes concepts related to managing dispatchers.

Managing Crews And Activities

Dispatchers are responsible for all the crews and work associated within the dispatch areas they manage. A dispatch area identifies a set of service areas and service classes representing the type of work being monitored in those service areas.

Refer to [Defining Dispatch Areas](#) for more information.

The dispatcher's record defines the dispatch areas that the dispatcher is allowed to monitor. When a dispatcher starts their shift, they may decide to monitor all or some of these areas as needed and, depending on their user privileges (also defined on their dispatcher record), they may also be able to set the range of dates to monitor.

These settings determine which crews and activities they are responsible for during their shift and they may change them as needed at any time.

Refer to [Dispatcher Shifts](#) for more information.

A crew is monitored by a dispatcher when at least one of their shifts covers a dispatch area managed by the dispatcher and falls within the dispatcher's monitoring time period. To be included in the dispatch area, the shift has to cover at least one of the service areas and allowed to perform one of the service classes defined by the dispatch area. Shifts that have not started yet need to have their planned time fall within the dispatcher's monitoring time period.

Crews that have started their shift and are still estimated to work during the dispatcher's shift, as well as those crews that have completed their shift during the dispatcher's shift are also considered irrespective of their original planned times.

NOTE: It is important to properly logoff your dispatcher shift and log back in each day to properly remove yesterday's completed shifts from the list of crews you monitor.

If your organization uses contractor management functionality, dispatchers can also monitor contracted work. Similar to crew shifts, any contractor capacity is considered within the dispatcher's responsibility when its service areas and service classes are covered by at least one of the dispatch areas currently being monitored, and it overlaps with the dispatcher's monitoring period.

Refer to [About Contractors](#) for more information.

Dispatchers monitor work within their dispatcher areas that is not yet assigned to any crew as well as all work already scheduled to crews (or contractor capacities) they monitor (even if these activities are not within their dispatch areas). Activities where the earliest time they can be scheduled falls outside the dispatcher's monitoring period are excluded.

An activity is considered to be within a dispatch area if its service area and at least one of the service classes it is associated with are covered by the dispatch area.

The dispatcher's scope of crews and activities they watch determines the data displayed on their CDI portal.

Refer to [About CDI](#) for more information.

If the dispatcher is set up to [handle alerts](#), the system may direct to them alerts related to the crews and activities they monitor, representing situations that require attention or manual intervention.

Refer to [About Alerts](#) for more information.

Dispatcher Shifts

There is no need to create a dispatcher shift in advance. The system creates a new dispatcher shift automatically when an authorized dispatcher logs on to start their shift from the CDI portal.

When logging in, the dispatcher can select which dispatch areas (of those specified in the dispatcher record) they want to monitor for the shift. They can also set the date range to monitor if they are allowed to do so based on their dispatcher record. These settings may be changed as needed throughout their shift.

The list of dispatcher areas and the period of time being monitored control the [scope](#) of crews and activities the dispatcher is responsible for during their particular shift.

While they are logged on, the dispatcher may use the CDI portal to manage their crews and activities as well as receive alerts related to these crews and activities. When the dispatcher logs off, they are officially no longer responsible for any crews or activities and alerts that may have been assigned to them are removed from their queue and reassigned to other logged on dispatchers as needed.

NOTE: For a dispatcher to be considered "logged on," they must have a logged on dispatcher shift. This is not achieved by just logging in to the application, but rather explicitly using the Logon button in the CDI portal.

Refer to [Handling Alerts](#) for more information.

A dispatcher can log off and log on again at any time, and each time they log on, the system creates a new dispatcher shift record. Each new dispatcher shift is automatically defaulted with the previous shift's settings that the dispatcher may adjust as needed.

From business perspective, a dispatcher's shift may be an 8-hour period from 9:00 to 5:00, but from a system standpoint, the shift lasts from logon to logoff. It is strongly recommended that the dispatcher logs off when they are done for the day so that alerts are properly reassigned to other logged on dispatchers and thus get addressed in a timely manner.

The dispatcher type record (and optionally also the dispatcher record) may be used to define an estimated duration for a dispatcher's shift. The system informs the dispatcher at logon time of their planned logoff time which they can adjust any time throughout their shift. If the dispatcher does not properly logoff at their estimated logoff time, the system would automatically do so some time after their planned time.

NOTE: Logic to handle dispatcher logoff is handled by a monitor algorithm defined on the logged on state of the dispatcher shift business object. Refer to this algorithm for more information.

NOTE:

NOTE: Additionally, the **Dispatcher Shift Monitor** batch process needs to run frequently to properly keep the list of crews being monitored by a dispatcher up to date by taking into account new shifts that may have been added after the dispatcher has logged on to their shift. It is also responsible for handling the logoff of dispatcher shifts that were not explicitly logged off in a timely manner. Refer to [M1-DSMTR](#) batch control for more information.

Handling Alerts

When a situation related to a crew or an activity warrants a dispatcher's attention, the system creates an alert for it and automatically assigns the alert to a logged on dispatcher to address it in a timely manner. Based on alert type configuration, these alerts may be reassigned to another logged on dispatcher if not addressed in a timely manner.

A dispatcher is considered logged on and hence available to receive alerts while their [dispatcher shift](#) is still in the logged on state. When they logoff their dispatcher shift, any remaining alerts still assigned to them are removed from their queue to be reassigned as needed.

NOTE: It is important for dispatchers to properly logoff their dispatcher shift in the CDI portal in a timely manner so as not to receive new alerts while they are not actually there to handle them.

Alerts assigned to the dispatcher and awaiting their action are listed in the alert queue zone on the dashboard. A dispatcher may indicate on their dispatcher shift that they back up another dispatcher while that person is temporarily away from their desk. When they do so the alerts assigned to the dispatcher they back up appear on their backup alert queue zone.

Dispatchers that are focused on reviewing future schedules may be set up to opt out of handling alerts on their dispatcher record. This allows them to concentrate on reviewing the next few days' schedule as opposed to handling today's exceptions.

Refer to [About Alerts](#) for more information.

KPIs

Key performance indicators (KPIs) are quantifiable measurements related to crews and activities dispatchers monitor during their dispatcher shift. These measurements are presented on the CDI portal to support dispatcher decision making and in-day exception handling.

KPIs are associated with either activities or crew shifts. For example, an activity-related KPI might indicate the number of appointments in jeopardy of being missed. A shift-related KPI might indicate the number of idle crews.

Refer to [Understanding Key Performance Indicators \(KPIs\)](#) for more information.

The dispatcher record lists the key performance indicators (KPIs) that the dispatcher should monitor and only those may appear in the CDI portal's "KPI Summary" zone.

Refer to [Monitoring KPIs](#) for more information.

Managing Depots

The CDI portal allows dispatchers to toggle their Gantt view to display the work scheduled to bays of a specific depot they are responsible for. The dispatcher record defines the list of depots the dispatcher monitors and only those depots are available for viewing on the Gantt.

Only depots that use bays may be monitored this way. Refer to [Bays](#) for more information.

NOTE: The depots a dispatcher manages do not impact the scope of crews and activities they are responsible for during their shift.

To Dos In Scope

Unlike an alert that represents a situation that requires a dispatcher's immediate attention, a To Do Entry may be raised to notify a user, not necessarily a dispatcher, of a situation that typically does not call for an immediate attention. To Do Entries associated with crews and activities that the dispatcher is currently monitoring are summarized in the "To Do In Scope" dashboard zone.

About CDI

The Common Dispatching Interface (CDI) portal helps dispatchers effectively manage day-to-day operations and facilitates communication with crews in the field. The CDI supports context-based decision making for handling exceptions that arise during the day. Although the system automatically schedules and dispatches tasks to available and qualified crews, a dispatcher can override virtually any decision made by the system to handle special situations.

Most dispatcher functions are performed from the CDI portal, which provides a single point of interaction for monitoring and managing tasks, crew shifts, and KPIs. You can also view and manage alerts and To Do entries from dashboard zones without ever leaving the portal.

Refer to [About Dispatchers](#) for more information.

This section describes key aspects related to functionality supported by the portal.

CDI is for Dispatchers Only

You have to be defined in the system as a dispatcher to take advantage of the functionality supported by the CDI portal. The crews and activities you are responsible for as a dispatcher and that are managed by the portal are determined by settings provided by you when you log on to start your [dispatcher shift](#). Until you log on to your shift, no information is displayed.

Refer to [Dispatcher Shifts](#) for more information.

CDI Data

The dispatcher's scope of crews and activities they watch during their shift determines the data displayed across all zones of the CDI portal and the CDI geographic map.

Only crew shifts, contractor capacities, and activities managed by the dispatcher during their current shift may be accessed and displayed by the CDI portal. Use corresponding general search portals, outside of the CDI portal, to access entities not currently managed by the dispatcher.

In the same way, if the dispatcher handles alerts, alerts are assigned to them based on the crews and activities they currently watch.

Refer to [Managing Crews And Activities](#) for more information.

NOTE: The **Dispatcher Shift Monitor** batch process needs to run frequently to properly keep the list of crews being monitored by a dispatcher up to date by taking into account new shifts that may have been added after the dispatcher has logged on to their shift. Refer to [MI-DSMTR](#) batch control for more information.

The scheduling Gantt zone supports the ability to apply filters that narrow down the list of crews displayed in the Gantt. It is important to note that filtering is a Gantt presentation feature only and as such does not impact the list of crews in the dispatcher's scope.

NOTE: While the Gantt list of crews is filtered, all other zones, CDI map, and alert functionality still access and rely on the dispatcher's entire scope of crews and activities.

CDI Map

The CDI map is launched from the toolbar on the Scheduling Gantt zone on either the main CDI or the KPI Specific KPI portals. The map's scope of data is the same as the scope of data of the CDI portal.

NOTE: Filtering the list of crews in the Gantt does not also filter the crews visible on the map.

The map supports the display of following geographic information:

- The location of any type of task, including activities, meeting places, etc
- The current location of a crew. The location of a crew is determined by its last known or estimated location. The estimated location is used when the last actual location is either not known or considered too old as per the **GPS Threshold Minutes**[global configuration](#) parameter.

- The location of a vehicle is determined based on its last known GPS location, if any.
- The locations where crews logon and logoff their shifts.
- The location of depots.
- The boundaries of service areas, if the information is available. Refer to [Understanding Service Areas](#) for more information.
- The boundaries of active alert areas. Refer to [About Alert Areas](#) for more information.
- The boundaries of custom geographic areas. Refer to [Defining Geographic Areas](#) for more information.
- Any information provided by custom MapViewer themes, for example, a network map.

Information is presented in layers that you can turn on and off as needed. Some of these layers may be configured to be turned on and displayed when the map initially opens.

NOTE: All CDI map layers are defined using the **CDI Map Layer (M1-CDIMapLayer)** extendable lookup which also allows you to customize which layer is to be displayed by default.

In addition to layered information, the map also highlights the following on demand from another portal:

- A specific task's location of interest
- A specific crew on the map
- One or more shift routes, each in a separate color. The route(s) presented by the map is the same route used by the scheduler to create the schedule. This allows the dispatcher to see an accurate depiction of the crew's actual work route.

CDI Is Customizable

Many aspects of the CDI portal are configurable and customizable. For example, you can configure custom task color rules, custom filters, control functionality used by the Gantt, add custom CDI map layers and more.

Refer to [Customizing The CDI Portal](#) for more information.

Maintaining Dispatchers

This portal is used to display and maintain a dispatcher record.

Refer to [About Dispatchers](#) for more information.

You can access the portal from the **Main Menu > Resource Management > Dispatcher**. You are brought to a query portal with options for searching for a specific dispatcher. Once a dispatcher has been selected you are brought to the maintenance portal to view and maintain the selected record.

The **Dispatcher** zone on the portal's **Main** tab page provides information about the dispatcher.

The **Calendar** zone on the portal's **Calendar** tab page provides a calendar view of items on the dispatcher's calendar for the current month.

Refer to the Calendar (F1-CALENDAR) zone type for more information.

NOTE: Typically the CDI portal should be defined as the Home page for dispatchers.

Maintaining Dispatcher Shifts

This portal is used to display and maintain a dispatcher shift record. Typically dispatchers create and maintain their dispatcher shift records directly on the [CDI](#) portal.

Refer to [Dispatcher Shifts](#) for more information.

You can also access the portal from the **Main Menu > Resource Management > Dispatcher Shift**. You are brought to a query portal with options for searching for a specific dispatcher shift. Once a dispatcher shift has been selected, you are brought to the maintenance portal to view and maintain the selected record.

The **Dispatcher Shift** zone on the portal's **Main** tab page provides information about the dispatcher, including the various maintenance actions applicable to its current state.

CDI

The CDI portal helps dispatchers effectively manage day-to-day operations and facilitates communication with crews in the field.

As a dispatcher, this portal should typically be set up as your home page. You can also access the CDI portal from the **Main Menu > Service Management > CDI**.

Refer to [About CDI](#) for more information.

The following zones may appear as part of the portal's **Main** tab page:

- **Logon**. Use this zone to logon and start your dispatcher shift. The zone is not visible once you have logged on to your shift. Refer to [Managing your Dispatcher Shift](#) for more information.
- **KPI Summary**. This zone provides a current summary of the key performance indicators (KPIs) related to crews and activities you monitor. Refer to [Monitoring KPIs](#) for more information.
- **Scheduling Gantt**. This zone displays a graphical representation of the crews you are monitoring and their scheduled tasks in a time line, allowing you to take action on those tasks. Refer to [Using the Scheduling Gantt](#) for more information.
- **Activity Search**. Use this zone to search, sort, and filter as well as perform standard dispatching functions on one or more selected activities you are responsible for during your shift. Refer to [Using the Activity List](#) for more information.

The following zones may appear as part of the **Contractor Capacity** portal tab page:

NOTE: This tab only appears if your company works with external contractor companies and you manage contracted capacities as part of your dispatcher shift responsibilities.

- **Capacity Search**. Use this zone to search, sort, and filter the contractor capacities you are responsible for during your shift.
- **Assigned Work**. This zone displays the activities assigned to a specific contractor capacity you have selected in the Capacity Search zone.

Refer to [Monitoring Contracted Work](#) for more information.

The following zones may appear as part of the **Shift** portal tab page

- **Logon**. The same zone as on the main tab is provided here for easier access.
- **Dispatcher Shift**. This zone displays your current dispatcher shift settings. Refer to [Managing your Dispatcher Shift](#) for more information.

Managing your Dispatcher Shift

The crews and activities you are responsible for as a dispatcher and that are managed by the portal are determined by settings provided by you when you log on to start your dispatcher shift, as well as the time you have started your shift.

Until you log on to your shift, no information is displayed.

To log on and start your dispatcher shift, click the **Logon** button either in the Main tab or the Shift tab. When you do so the system prompts you for your dispatcher shift settings, automatically defaulting them from your previous shift, if any.

Once logged on you can now begin monitoring your crews and work in your areas.

Use the Shift tab to view and change the settings on your dispatcher shift as well as to logoff the shift at the end of your day.

Occasionally as part of your shift, you may be responsible for backing up another dispatcher by monitoring their alerts when they step out for a while. To do so, add their user IDs to the list of users you back up on your dispatcher shift. Alerts assigned to these dispatchers would now appear on your backup alert queue.

NOTE: When done for the day, it is very important that you properly logoff your dispatcher shift so that you are no longer considered logged on and available to handle alerts.

Refer to [Dispatcher Shifts](#) for more information.

Using the Scheduling Gantt

The Scheduling Gantt zone displays a graphical representation of the crews you are monitoring and their scheduled tasks in a time line, allowing you to take action on those tasks.

The list of crews appears on the left pane and their schedules appear on the right pane. Each row in the schedule pane displays one or more shifts along with their scheduled tasks for the corresponding crew in the left pane. The schedule appears across a time line where a vertical line indicates the current date/time for reference.

This section of the user guide describes the various views you may use and actions you may take using this zone.

Gantt Views

The timeline view is the main Gantt view presented when you initially navigate to the CDI portal. The Gantt supports additional views, each providing a different perspective on the schedule.

The left and right panes, as well as the [toolbar](#), may show different content and options that are relevant to the type of view you are currently using.

The same information displayed using the main timeline view may be presented using an alternate "Plan View" where tasks in the schedule area are displayed in the order they are planned to be worked within the shift, irrespective of time. In this view the left pane lists individual shifts (as opposed to crews) and the schedule area only displays tasks, not shifts.

Refer to [Plan View](#) for more information.

You can also switch the Gantt view to display a single shift's list of scheduled tasks along with each task's available time windows in a time line. In this view, the left pane lists tasks and the schedule area portrays their schedule overlaying their time window information.

Refer to [Tasks View](#) for more information.

If you are tasked with monitoring the work scheduled to specific depots you may switch the Gantt view to display a single depot's schedule. In this view, the left pane lists the depot's bays and the schedule area portrays the tasks scheduled to each bay.

Refer to [Depot View](#) for more information.

Crews Pane

When using the main timeline view, the left pane shows a list of crews you monitor. Shift related columns in that pane pertain to the crew's currently active or earliest planned shift in the scope of data.

NOTE: Your organization may have configured the zone to show individual crew shifts on the left pane as opposed to crews. If the left pane shows individual shifts, shift related columns pertain to the specific shift. Refer to [Configuring The CDI Portal](#) for more information.

The **Crew** column on the left pane displays the crew name as well as their current working state, i.e. whether they are out of range, en route to an activity, working on site, out of service, etc.). The columns may also include the following icons:

- If the crew is using a mobile device to work their current shift, an icon appears to indicate whether they are connected or not.
- If your company works with contractor crews, a symbol next to the crew name indicates that this is a contractor crew.

You may include additional crew related columns that you prefer to work with and organize them as you like. Refer to [Customize Your Gantt Display](#) for more information.

You can temporarily further restrict the list of crews you see by using [filters](#).

Right-click anywhere on a row in this pane to choose from the various [actions](#) you can take on the corresponding shift.

The Schedule Area Pane

When using the main timeline view, each row in the schedule pane shows one or more shifts for the corresponding crew in the left pane along with their scheduled tasks.

When multiple shifts are displayed, large gaps in time in between shifts, may be collapsed by default to provide a more concise view of the schedule. You may turn this feature on or off to your liking. Refer to [Customize Your Gantt Display](#) for more information.

Icons, colors, symbols, and other graphical elements provide visual cues to information available regarding the crew's shifts and tasks.

Refer to [Visual Cues](#) for more information on how they may be customized.

This section describes the various visual cues used in this pane.

The Shift Bar

A shift is portrayed as a larger bar on top of which smaller bars are displayed representing tasks scheduled to the shift or locked to the shift but not yet scheduled to it.

The shift bar provides multiple layers of information. Using different **colors**, it shows the shift's originally planned start and end times, its broad variable scheduling period when applicable, any scheduled overtime, as well as the actual times the crew worked that shift.

- A shift's original planned time is displayed in the "planned time" color for the width of the shift's planned start to end times.
- For a shift defined with a variable scheduling period (and has not yet become fixed in duration), the planned time is displayed in the "planned time" color whereas the variable broader period is displayed in the "variable time" color as an underlay bar to the planned time bar. As the shift's start time approaches and it becomes fixed in duration, the variable underlay bar is no longer displayed.
- Once started, the actual shift time bar is displayed on top of the originally planned time bar for reference. The actual time bar starts from the actual time the crew started their shift and ends at their planned end time. The color of the actual time

bar reflects the shift state. While still in progress, if the shift enters overtime, the remainder of the shift is drawn using the "overtime" color. When the crew completes their shift, the actual duration bar ends at the completion time.

A thick black **frame** around the shift time bar represents a closed shift.

Right-click on the shift bar area to choose from the various [actions](#) you can take on that shift.

NOTE: Since the shift bar area is overlaid with task bar images, you are likely to right click on a task bar area. In this situation the task's list of options is presented from which you should choose the "Shift Actions" option to access the shift related actions.

The Task Bar

Each task bar uses visual cues to convey information about the task.

A task that is scheduled as part of a shift's route uses a preceding [travel time line](#) representing travel time for the crew to get to the task's location from the previous task's location (or from where they start their shift).

The length of the task bar represents the task's scheduled duration, from start to end, as well as any additional time spent on site. If the crew is planned to arrive or has arrived to the location before they can start the work, that waiting time is also included in the total length of the task bar.

The **frame** of the task bar distinguishes between emergency activities, appointments, and other tasks as well as the scheduling state of the task.

- Emergencies and appointments use a thicker frame with designated colors. Emergencies also have their frame blink to attract the dispatcher's attention to the situation.
- When a task is locked to the schedule but has not yet been incorporated into the shift's route, its frame uses a dotted line as opposed to the solid line used for scheduled tasks. Such a task may be a break planned for the shift but not yet used by the scheduler, or an activity that has just been manually allocated to the shift but the scheduler has not yet processed the request.
- While waiting to be scheduled, the task is stacked at the end of the new crew's schedule, with no travel to it (since it is not yet sequenced in the route and the previous task is not yet known).
- When an activity has just been manually deferred to a future time and the scheduler has not yet processed the deferral, a dashed lined frame is used. The activity remains with this visual cue until the scheduler provides a new schedule for it.

By default the **color** of the task bar indicates its status. Refer to [Task Color Rules](#) for more information on how to add custom color rules.

Icons are used to convey the following information about the task:

- This is a break task
- This is a depot task
- These are administrative tasks such as meetings and Non Productive Tasks (NPT). Some types of meetings may be configured with a unique icon designed to make it easier for you to distinguish them on the Gantt at a glance.
- Some activity types may be configured with a unique icon designed to make it easier for you to distinguish that work on the Gantt at a glance. If the activity type is configured with such an icon then it is displayed as part of the task bar for all activities of this type.

NOTE: Showing an icon on many of the tasks on the Gantt may be overwhelming for some and useful to others. Showing icons for activities is therefore optional. You may turn this feature on or off to your preference. Refer to [Customize Your Gantt Display](#) for more information.

- A dot at the top left corner of the task bar indicates that the activity is manually allocated to this shift.
- A dot at the bottom right corner of the task bar indicates that the activity is fixed to this crew.

- This activity is part of a dependency with other activities. This means the task is either part of a job or a chain, or is a complex activity visit task.

In addition, the task type code may also be included as part of the task bar. You may turn this feature on or off to your preference. Refer to [Customize Your Gantt Display](#) for more information.

Right-click on the task bar area (or the travel line if present) to choose from the various [actions](#) you can take on that task.

Travel Time Line

Until the crew notifies the system of the exact time they went en route to their next task, the system assumes the crew starts traveling right after completing their previous task (or starting their shift) without delay. Any idle time that may be incurred because the crew is planned to arrive (or arrived) at the next location before they can actually start their next task is included as part for the next task's time on site. This idle time is also referred to as schedule idle time.

Once the system knows the exact time the crew has started their travel to the next task, the travel line is adjusted to start at that actual time. The gap between the previous task (or the beginning of shift) to the start of the travel line is considered as crew idle time.

A travel line is also present from the last task in the shift route to the crew's logoff location.

The Toolbar

Use the toolbar to view and change the information that you see in the Gantt as described in this section.

Zoom In / Zoom Out. Adjusts the task schedule area to show a narrower/broader time frame.

Previous / Next. Scrolls the Gantt view to the previous and next day within your monitoring period.

Go to Now. Scrolls the Gantt view to the current date and time.

Show Map. Launches the [CDI Map](#) (or switches focus to it if already open).

Customize Display. Use this tool to change the look and organization of your Gantt view and save your preferences for future use. Refer to [Customize Your Gantt Display](#) for more information.

Filter. Use this tool to temporarily refine the list of crew shifts that are displayed in the Gantt. Refer to [Filter The List of Crews](#) for more information.

Plan View: Toggles the Gantt view between the default timeline view and [Plan View](#), and back.

Color Rule: By default task colors are determined based on their class and status. If your organization has defined other rules to set a task's color, you may choose such rule from the list.

Time Zone: By default the Gantt is displayed in your time zone as specified on your user record. If your company operates in multiple time zones, you may change the display to use another time zone from the dropdown list.

Customize Your Gantt Display

You can add additional columns to the left pane, change some look and organization aspects of the right pane, and save your preferences for future use.

To set up a new display preference:

- Click on the **Customize Display** option in the [toolbar](#) and select the **Change Display** to specify which columns you want to see on the crews pane as well as set options to control various aspects of the task schedule pane display.
- Back on the Gantt, rearrange the columns in the left pane as per your liking by dragging a column and dropping it before or after another column.
- Click on a column header to sort the list by that column. Hold down the **Ctrl** key before you click on the next column header to further sort by additional columns.

- Resize the left pane to best fit the columns selected.
- Set the Color Rule to be different than the default set for this display preference.
- After the display is changed, you can use the Customize Display tool again to save the display for easy access later as a one of your display preferences. Select **Save Display** and provide a name for your current display.

Access your listing of saved displays by selecting the **My Displays** option. You may select any display from this list to make it your current display at any time.

You may use the **Set as Default** option to mark the current display as the default view that opens whenever you open the portal.

Filter The List of Crews

Use this tool to further refine the list of crew shifts that are displayed in the Gantt according to various criteria such as by activity information, by capability, by coverage, by resource, by emergencies, or by custom filters that you save.

To set up a new filter:

- Click on the **Filters** option in the [toolbar](#) and select the type of filter you want to use.
- Provide corresponding criteria to filter by and test it before applying it into your Gantt.

Once applied, the list of crews on the left pane includes only those that have a shift you monitor that matches the filter criteria. Only one filter may be applied at any point in time and its description is displayed on the top bar along with any criteria that you may have specified.

It is important to note that filtering is a Gantt presentation feature only and as such, does not impact the list of crews in the dispatcher's scope.

NOTE: All other zones, including the CDI map, do not reflect the filtered list. In the same way, if you are set up to handle alerts, you are assigned alerts as usual and not based on the filtered list.

Click the close icon to remove filtering and resume the full view.

You can save filters that you frequently use to your list of preference filters which you can refer back to later:

- Click on the Filters option again
- Select **Save Current** to store your currently applied filter, along with the criteria you have specified, to your preference list and provide a name for the filter.

These filters can be saved to a list that you can easily select to toggle the display based on different filters and preferences that you have defined.

Access your list of saved filters by selecting the **My Filters** option from the Filters icon on the toolbar. You may apply any filter from this list at any time.

NOTE: Filters cannot be saved as your default display since it is a best practice to review all the information in your scope of crews and activities before applying filters.

Taking Actions

The Gantt allows you to perform actions on a shift as well as one or more tasks.

Shift Actions

Right-click on a row in the crews list pane or on a shift bar in the scheduler area pane to choose from the various actions you can perform on this shift. The available options may vary depending on the selected shift and your system's configuration.

The following highlights some of the more common actions you may perform for a shift:

- Actions that impact the schedule of the shift. Refer to [Crew Shift Actions](#) for more information.
- **Allocate.** This action works in conjunction with the [Activity Search](#) zone. Click the checkbox for one or more activities in the search zone, then use this action to allocate the selected activities to this shift.
- **Show Crew On Map.** Highlights the location of this crew in the open [map](#) window. This action is only applicable for crews that have started their shift and are still working. Before they start their shift and once they complete the shift, the crew's location is not known to the system.
- **Show / Hide Route On Map.** These actions draw the route for this shift on the open [map](#) window and remove it when no longer needed.
- **Tasks View.** Toggles the Gantt view to display the tasks scheduled to this shift only. Refer to [Tasks View](#) for more information.
- **View Details.** Allows you to navigate to the shift's maintenance portal to view more information about it or take further actions.
- **Send Mail.** Sends a message to the selected crew using the internal messaging system. If the crew uses a mobile device, the message is directed to their device, otherwise they may view the message when they are in the office.

Task Actions

Right-click on a task bar (as well as its travel line, if any) in the scheduler area pane to choose from the various actions you can perform on this task. The same menu is also available when you use the [Tasks View](#) and right-click on a task row in the left pane. The available options may vary depending on the type of task, its state as well as your system's configuration.

The following highlights some of the more common actions you may perform on a task:

- If this is an activity, you can perform actions that impact its current schedule. Refer to [Activity Actions](#) for more information.
- **Show On Map.** This action highlights the location of the task on the open [map](#) window.
- **View Details.** Allows you to navigate to the task's maintenance portal to view more information about it or take further actions.

Some of the actions may be taken on multiple tasks at once. To select multiple tasks, hold down the shift key and select the tasks. Right-click to choose from the applicable options.

You may also use **drag-and-drop** to unassign one or more activities from their currently assigned shift and allocate them to another shift.

Plan View

In addition to the time line view, the zone also supports a plan view that shows scheduled tasks in the order they are planned to be worked within the shift, irrespective of time.

When using this view:

- The left pane lists individual crew shifts (as opposed to crews), but other than that it works exactly like the [Crews](#) pane of the timeline view.
- The right pane shows the tasks scheduled to each shift, in their work sequence order. Each task is portrayed as a [Task Bar](#), exactly like the schedule area pane of the timeline view. There is no timeline, no shift bars, and no travel time information.

This view allows you to assess the overall schedule at a glance. As such, it may be even more useful when used in conjunction with [Task Color Rules](#) to highlight task related scheduling conditions in color.

Use the **Plan View** option on the [toolbar](#) to switch the timeline view to Plan View. Click on the icon again to toggle back to timeline view.

Tasks View

You can switch the Gantt view to display a single shift's list of scheduled tasks, along with their time window information, in a time line.

When using this view:

- The left pane lists individual tasks scheduled to the shift, but other than that, it works exactly like the [Crews](#) pane of the timeline view.
- The right pane displays time windows for each task on a time line, each type of time window in a different color. For reference, the task's schedule on the shift route is displayed as an overlay over the time windows. Each task is portrayed as a [Task Bar](#), along with a preceding [Travel Time Line](#). if any, exactly like the schedule area pane of the timeline view.

Time window information includes the activity's available time windows, as well as any override time windows that may have impacted them.

Refer to [When Can The Activity Be Actually Scheduled?](#) for more information.

This type of view allows you to easily visualize where each task is scheduled relative to its time windows, and assist you when you make scheduling decisions, monitor crew progress, and identify tasks that are in jeopardy.

To switch the Gantt to this view for a specific shift, right-click on a shift (left or right pane as applicable) and choose the **Tasks View** option. The view changes accordingly and the description of the selected shift appears in the Gantt bar to indicate you are in this view. Click the close icon to revert back to previous view.

Depot View

You can switch the Gantt view to display the schedule of a single depot you monitor.

When using this view:

- The left pane lists individual bays.
- The right pane shows the tasks scheduled to each bay in a time line. Each task is portrayed as a [Task Bar](#), exactly like the schedule area pane of the timeline view.

To switch the Gantt to this view for a specific depot, choose the depot from the dropdown list. Only depots listed on your dispatcher record as depots you are allowed to monitor are available for selection. Reset your depot selection to revert back to standard timeline view.

You can also switch the depot view to display a single bay's list of scheduled tasks along with each task's available time windows in a time line. In this view, the left pane lists tasks and the schedule area portrays their schedule overlaying their time window information.

Refer to [Bay View](#) for more information.

Bay View

While in Depot View you can switch the Gantt to display a single bay's list of scheduled tasks, along with their time window information, in a time line.

When using this view:

- The left pane lists individual tasks scheduled to the bay.
- The right pane displays time windows for each task on a time line, each type of time window in a different color. For reference, the task's schedule is displayed as an overlay over the time windows of the bay and task. Each task is portrayed as a [Task Bar](#), along with a preceding [Travel Time Line](#). if any, exactly like the schedule area pane of the timeline view.

Time window information includes the bay's time window and the activity's available time windows, as well as any override time windows that may have impacted them.

Refer to [When Can The Activity Be Actually Scheduled?](#) for more information.

This type of view allows you to easily visualize where each task is scheduled relative to its time windows, and assist you when you make scheduling decisions, monitor crew progress, and identify tasks that are in jeopardy.

To switch the Gantt to this view for a specific bay, right-click on a bay (left or right pane as applicable) and choose the **Bay View** option. The view changes accordingly and the description of the selected bay appears in the Gantt bar to indicate you are in this view. Click the close icon to revert back to previous view.

Using the Activity Search

The Activity Search zone allows you to search for activities and take action on them. Only activities you currently manage in your shift may be accessed using this zone. Use the general activity search portal to access activities you do not currently manage.

Use the action buttons to perform standard dispatching functions on one or more selected activities.

Refer to [Activity Actions](#) for more information.

You may also allocate selected activities to a specific shift on the Gantt:

- Select the activities to allocate as usual, but do not use the Allocate action button on this zone
- Instead, locate the shift you want to allocate the activities to on the Gantt and using the right-click shift actions menu, choose the Allocate option to complete the allocation update.

Click on the **Show On Map** icon to highlight the location of the activity on the open [map](#) window.

You can also locate a specific activity on the Gantt, if presented there, using the **Show On Gantt** icon.

Using the CDI Map

The CDI Map allows you to view on a geographic map the location of the activities you manage on your shift, your crews' current location, their shift routes, nearby company vehicles, and more.

Refer to [CDI Map](#) for more information.

The content of the map and the layers available correlates to the CDI portal you currently work with, the main CDI, or the KPI specific CDI. For example, the activities and crews layers show the entire scope of activities and crews you manage on your shift, if you are currently in the main CDI portal, or those associated with the selected KPI if you currently in the KPI Specific CDI.

NOTE: Gantt filtering does not impact the scope of data presented on the map.

To **launch** the map, click on the **Show Map** icon on the Scheduling Gantt's toolbar.

The map opens in a separate window, allowing you to use the CDI portal and the CDI map side by side.

In addition to toggling on and off the various layers of information on the map itself, you can also interact with the map from the CDI portal, requesting to:

- Highlight a specific task's location on the map. Only one task may be highlighted at any time.
- Highlight a specific crew's current location on the map. Only one crew may be highlighted at any time.
- Show or hide a crew's route on the map. Each new route to display is drawn in a different color.

NOTE: You must launch the map before you can display specific crews, activities, or routes on the map. After launching the map, leave the map window open and return to the **CDI** portal to use these options.

The **Clear All** button clears all route lines, if any are displayed, and also clears the latest marker of a highlighted crew or task shown on map.

The **Refresh** button updates the map with the latest available data.

The map interface works similar to other common mapping programs. You can pan and zoom using the icons in the upper left-hand corner of the map window. You can also click and drag any point on the map to move its position within the map window. The legend icon in the lower right-hand corner of the map window can be expanded to show the icons used to represent different objects on the map.

Monitoring KPIs

Key performance indicators (KPIs) are quantifiable measurements related to crews and activities you currently monitor. They highlight certain situations or provide additional information that may assist with your decision making and exception handling.

The KPI Summary zone uses colors and icons to easily identify KPIs based on their importance. Typically, green indicates normal KPI values; red is used for values that fall outside a pre-defined range and require immediate attention; yellow is used for warnings and to indicate a potential problem.

Some of the KPIs provide information only while others may require an action to resolve them. If action is required to resolve a KPI, click on that KPI to navigate to the KPI Specific CDI portal, which provides a narrower CDI view with information necessary to resolve the issue.

Refer to [KPI Specific CDI](#) for more information.

Monitoring Contracted Work

Contractor work capacities represent an external contracting company's agreement to take on a certain amount of work of a certain type in specific service areas for a period of time. The tab is only accessible if your company uses contractor functionality and the dispatcher's data scope includes contractor capacities.

The zones in this tab allow you to search for contractor capacities you monitor in your dispatcher shift and review their state and progress. You can also display the activities assigned to a specific capacity and take actions on these activities as needed.

Refer to [About Contractors](#) for more information.

My Alert Queue

The "Alert Queue" appears on the dashboard portal when you have alerts assigned to you.

Similarly the "Backup Alert Queue" zone appears for as long as you are backing up other users and they have alerts assigned to them.

These zones are there for you irrespective of whether you are currently working with the CDI portal or not.

Refer to [Handling Alerts](#) for more information.

Managing To Do Entries

The "To Do In Scope" zone appears on the dashboard portal when any of the crews or activities you currently monitor have To Do Entries associated with them.

The zone summarizes open To Do entries sorted by their type. Click on a specific To Do Type to navigate to the [Dispatcher To Do List](#) portal to review and address issues of the selected type.

KPI Specific CDI

You navigate to this portal when you click on an actionable KPI in the general [CDI](#) "KPI Summary" zone.

The "KPI Specific CDI" portal is similar in appearance and functionality to the general CDI portal, but only displays crews and activities needed to address a particular situation represented by the selected KPI. This is useful for troubleshooting the underlying conditions that contribute to the KPI, for example, when reviewing shifts going into overtime.

In the same way, the content of the [CDI map](#) and the layers available are adjusted to match the crews and activities related to the selected KPI.

NOTE: Using this portal does not narrow down your overall scope of crews and activities you are monitoring during your shift.

The following zones may appear as part of the portal's **Main** tab page:

- **Selected KPI.** This zone provides the current measurement for selected KPI. .
- **KPI Scheduling Gantt.** This zone displays a graphical representation of the crews associated with the selected KPI and their scheduled tasks in a time line, allowing you to take action on those tasks. Tasks associated with the selected KPI are distinguished from the rest of the tasks. The zone works similarly to the scheduling Gantt zone in the general CDI portal, with the exception of filtering. The latter is not supported by this zone as the list of crews is already filtered by the selected KPI.
- Refer to [Using the Scheduling Gantt](#) for more information.
- **KPI Activity Search.** Use this zone to search, sort, and filter, as well as perform standard dispatching functions on one or more selected activities associated with the selected KPI. Refer to [Using the Activity List](#) for more information.

Dispatcher To Do List

You navigate to this portal when you click on a specific To Do Type listed in the "To Dos In Scope" dashboard zone. The portal allows you to manage To Do entries related to activities or crews you are currently monitoring on your dispatcher shift.

Refer to [To Dos In Scope](#) for more information.

The **Dispatcher To Do List** zone on the portal's **Main** tab page lists open To Do entries for you to review and address as needed.

NOTE: You have to be defined in the system as a dispatcher to take advantage of this functionality.

Alerts

About Alerts

An alert represents a situation in which a predefined condition has been met that requires immediate user intervention. The alert may be related to a specific activity, a crew on a specific shift, or a company vehicle that caused the condition. For example, your system might generate an alert if an emergency activity is not acknowledged by a crew within 5 minutes.

NOTE: Alerts are not to be confused with To Do Entries. Alerts represent an issue that requires a dispatcher's immediate attention and as such, should be kept to exceptions of that nature.

The system automatically assigns an alert to a single logged on dispatcher, based on configurable business rules, and preferably to a dispatcher that currently monitors the entity in question.

Once assigned, the alert appears on the user's [Alert Queue](#) dashboard zone, waiting for the user to acknowledge they are aware of the situation. Depending on the configuration of each type of alert, acknowledgement may be all that is needed to close the alert and if not, the alert would reappear in the queue after some period of time until the issue is resolved or it is no longer relevant.

For alerts to be addressed in a timely manner it is important that dispatchers properly logoff their dispatcher shift at the end of their work day. This ensures alerts still assigned to them are properly recycled to other dispatchers that are still working, as well as prevent the system from assigning them alerts while off duty.

Refer to [Dispatcher Shifts](#) for more information.

This section describes concepts related to managing alerts.

Alert Type Controls Everything

Each alert references an alert type that represents a unique business condition to be checked and for which an alert should be raised when the condition is met. The alert type defines the attributes and controls the behavior associated with all alerts of that type, including:

- how the alert is raised
- how important the alert is relative to other types of alerts
- how the system determines if the underlying issue has been resolved and the alert can be closed
- how long before the alert reappears in the queue once acknowledged
- how long before the alert expires or recycles, i.e. directed to another dispatcher if not acknowledged after a period of time
- and more.

Refer to [Understanding Alert Types](#) for more information.

Alert Queue

The "Alert Queue" lists open alerts currently assigned to you. The zone appears on the dashboard portal when you have at least one alert assigned to you. As a dashboard zone it is guaranteed to always be present and available to you regardless of where you navigate to in the system.

The queue is sorted by priority order as defined on the alert type, and then by time so that more important types of alerts show before others, irrespective of time. Critical alerts may be configured to make a sound to ensure they are addressed as soon as possible.

Based on alert type configuration, alerts may support specific actions you may take to manage them.

Refer to [Alert Actions](#) for more information.

You should at least acknowledge viewing the alert, even if you don't have the time to address it right away. When you acknowledge an alert, you let the system know you are aware of the problem.

For some types of alerts, acknowledgement is all that is required from you and the alert automatically closes. For example, an alert that notifies you that an appointment has been missed.

When more intervention is needed, acknowledgement removes the alert from your queue for a predefined number of minutes, providing you some time to address the underlying situation, after which the alert shows up again in your queue, waiting for you to acknowledge it again. The alert may be snoozed in this manner until the underlying problem is resolved or the alert is recycled to another user, if not acknowledged in a timely manner.

From the alert queue, you can also navigate to the portal of the related subject of the alert, for example a crew's shift or activity, to better understand the problem and take actions accordingly.

Alerts are designed to automatically close when the underlying problem that caused them can be programmatically evaluated to no longer exists or is no longer relevant. This means that for most alerts you don't need to take any specific action to close them. For those where you need to address the underlying issue, addressing it causes the alert to close.

Refer to [Alert Evaluation](#) for more information.

Backup Alert Queue

The "Backup Alert Queue" lists alerts currently assigned to another user you are backing up. The zone appears on the dashboard portal when at least one alert is assigned to any of the users you back up.

As a dashboard zone it is guaranteed to always be present and available to you regardless of where you navigate to in the system.

Refer to [Handling Alerts](#) for more information.

Alert Actions

The available actions for an alert vary depending on the alert type and its current state. Common alert actions include the following:

Acknowledge

Indicates that you have viewed the alert and are aware of the problem. The system removes the alert from the alert queue for a specified number of minutes defined on the alert type and then reissues it. If an alert is not acknowledged within the time specified in the alert type, the system recycles it to another dispatcher.

Close

Closes the alert and removes it from your queue. This action is only available if the alert type supports manual closure.

Voice

For synchronization alerts only, this indicates that you want to voice-dispatch the activity.

Refer to [Synchronization Alert](#) for more information..

Force Cancel

You may respond to an activity cancellation alert assigned to you and use this action to recall the activity from the crew shift (or contractor capacity) it is currently dispatched to, and cancel it. This action also removes the alert from your queue. This is only supported for cancel activity alerts.

NOTE: It is important that you confirm the cancellation with the assigned crew (or contractor company) to ensure the work is not going to be performed. Failing to do so may result in a discrepancy of information provided to the host vs. the actual situation in the field.

Refer to [Activity Cancellation By Host](#) for more information..

NOTE: Logic to handle the business rules associated with forcing the cancelation of an activity resides in a "Cancel Activity" algorithm on the activity's business object. Click [here](#) to see the algorithm types available for this system event.

About Alert Areas

Alert areas are geographic areas that can be displayed on a map and that trigger an alert when resources enter or leave this area. Such alerts are sent to the dispatcher monitoring the resource or the closest service area to the alert location.

This section describes concepts related to managing alert areas.

Hazard Areas

Hazard areas establish a boundary or zone that, if entered, triggers an alert. If the most recent actual location of a crew or vehicle shows that they have entered the area, the system generates an alert. The hazard area supports an optional list of crew primary functions allowing you to limit the alert to crews performing certain types of work, as necessary.

An alert is generated when crews currently working these listed functions, if any, or any vehicles that are not part of any started shift enter the area. For example, supervisors or other company workers might drive around, but are not identified in the system associated to a shift. These would trigger an alert if they enter a hazard area. If no restrictions are indicated, any crew that enters the area will cause the system to generate an alert.

Vehicle Leave Areas

Vehicle leave areas establish a boundary or zone where vehicles are prohibited from leaving. This type of alert area supports an optional list of vehicle types, or vehicle primary functions, allowing you to limit the alert to vehicles of a certain type, as necessary.

An alert is generated when a vehicle of the listed types, if any, leave this area. This is typically used to track stolen vehicles or to make sure that crews are not using vehicles outside of designated work areas. If no specific vehicle types are defined, the alert applies for vehicles of any type.

Excluded Resources

Once presented with an alert for a resource, entering or exiting an alert area, a dispatcher may decide that the resource is allowed to be where it is and hence manually close that alert. Manually closing the alert excludes the resource from getting additional alerts for that area. However, if for some reason alerts should be created again for that resource, the dispatcher can manually take that resource off of the exclusion list.

Maintaining Alerts

This portal is used to view and maintain an alert. Dispatchers manage alerts assigned to them directly from the "Alert Queue" dashboard zone.

Refer to [About Alerts](#) for more information.

You can also access this portal from the **Main Menu > Service Management > Alert**. You are brought to a query portal with options for searching for a specific alert. Once an alert has been selected you are brought to the maintenance portal to view and maintain the selected record.

The **Alert** zone on the portal's **Main** tab page provides information about the alert, including the various maintenance actions applicable to its current state.

Maintaining Alert Areas

This portal is used to view and maintain an alert area.

Refer to [About Alert Areas](#) for more information.

You can access this portal from the **Main Menu > Service Management > Alert Area**. You are brought to a query portal with options for searching for a specific alert area. Once an alert area has been selected you are brought to the maintenance portal to view and maintain the selected record.

The following zones may appear as part of the portal's **Main** tab page

- **Alert Area.** This zone provides information about the alert area, including the various maintenance actions applicable to its current state.
- **Excluded Resources.** This zone lists all resources that have been manually excluded from alert processing for the specific alert area. Closing an alert for a resource automatically excludes the resource from getting additional alerts for that area. However, if for some reason you decide that alerts should be raised again for that resource, click the **Remove** button to take that resource off of the exclusion list.

Use the **Geographic Area Map** zone on the portal's **Map** tab page to define the map boundaries of the alert area. When ready activate the alert area.

Procedures

About Procedures

The system can be configured to include procedural instructions or questionnaires to be completed by mobile workers prior to taking on a specific primary function when starting a shift or changing their function mid shift, using a specific vehicle, or starting work on an activity. If the procedure fails, the corresponding business action cannot commence.

Usually procedures relate to safety precautions or regulatory requirements. For example, the mobile worker may need to complete a vehicle inspection at the beginning of a shift, or review safety instructions for a hazardous task.

Procedure enforcement is applicable to the mobile application only and is embedded as part of the following process flows:

- Request to start a shift
- Request to change a started shift's primary function
- Request to change a started shift's vehicle(s)
- Every time the crew starts an activity of a specific type. If they suspended and start it later, again they would be requested to complete the procedures.

Refer to [Enforcing Procedures](#) for more information.

Procedures cannot be manually created but rather initiated as part of the supported business events. The system tracks each procedure the crew has completed for audit purposes.

All procedures completed during a crew's shift are listed on the crew shift portal. Those related to a specific activity are also listed as part of the activity portal. You can also search for any procedure using the general procedure search portal.

This section describes concepts related to using procedures.

Procedure Type Controls Everything

Each procedure references a procedure type that controls the form presented to the crew as well as the business rules involved in completing it. The base product supports a configurable questionnaire type of procedure, but you can implement more sophisticated custom procedure forms to meet your business needs.

Refer to [Understanding Procedure Types](#) for more information on how to configure the system to enforce your procedures as part of the supported business events.

Shift Related Procedures

Certain procedures must be completed before the crew can start their shift as well as change their primary function or any of their vehicles mid shift. When this occurs, the crew is notified on their mobile device that they are required to complete one or more of these procedures.

Procedure types specified on a [Service Class](#) must be completed before the crew can start a shift referencing this service class as their primary function as well as each time they request to change their started shift to assume a new function.

Procedure types specified on a [Vehicle Type](#) must be completed before the crew is allowed to use a vehicle of this type. A separate procedure is required for each vehicle of the specific type, each time such vehicle is used.

Before the shift starts, you can view on the crew shift portal all the procedures the crew would be required to complete if they start the shift with its currently planned configuration. This includes any procedure types that are defined for the crew's primary function or related to any of the vehicles allocated to them on that shift.

At shift start time, the system determines again the list of procedures the crew is required to perform based on the actual primary function they chose and the vehicles they are about to use. Based on this list, the system creates corresponding procedure records for the crew to fill out and sends them to the crew's mobile device. The crew has to complete them and they have to all pass before the crew can start their shift.

In the same way, if a crew changes their primary function or vehicle allocation mid shift, corresponding procedures are yet again sent from the server and have to be completed before the changes take effect.

It is important to note that the list of procedures is determined and created on the server as part of processing the crew's synchronous request to start or update their shift. Procedures are sent back to the device as part of the response to this request.

NOTE: Your implementation may have different rules to determine which procedures are applicable for shift events. Logic to determine shift based procedures resides in a "Determine Shift Procedures" algorithm on the shift's business object. Click [here](#) to see the algorithm types available for this system event. Determining and creating procedures has to occur on the server so as to allow these algorithms to consult other entities and apply custom rules.

On the device, procedures are presented to the crew which then completes them one after the other. Each procedure is presented and validated on the mobile device.

A failed procedure may or may not be allowed to be manually overridden as passed based on your procedure type configuration. Once all procedures are passed or overridden, the system proceeds to perform the requested change.

Activity Related Procedures

Certain procedures must be completed on site before the crew starts to perform any activity work.

These procedure types are defined on the activity type. The system creates corresponding procedure records and sends them to the crew's mobile device when the activity is dispatched.

This is also true for an activity created in the field. When the activity is sent to the server it goes through the same process to handle procedures. This could be problematic in situations where the crew is currently working offline. To handle this, the activity type can be set to explicitly state how important it is for these procedures to be enforced when there is no connection coverage. If the crew decides to work the new activity and the activity type requires safety procedure clearance, the crew is required to be connected to receive the procedures and fill them prior to starting the work. Otherwise the crew may start the activity without completing the procedures in this situation.

NOTE: Your implementation may have different rules to determine which procedures are needed before the crew can start working on an activity. Logic to determine activity based procedures resides in the "Determine Task Procedures" algorithm on the activity's business object. Click [here](#) to see the algorithm types available for this system event. Determining and creating procedures has to occur on the server so as to allow these algorithms to consult other entities and apply custom rules.

On the device, when the crew attempts to start an activity that requires some procedures to be completed, the list of procedures is presented to the crew which then completes then one after the other. Each procedure is presented and validated on the mobile device.

A failed procedure may or may not be allowed to be manually overridden as passed based on your procedure type configuration. Once all procedures are passed or overridden the system proceeds to mark the activity started.

Maintaining Procedures

This portal is used to view a completed procedure filled out by the crew using the mobile application. As procedures are typically viewed in a context of a specific shift, you most likely navigate to this portal from the "Procedures" zone on the crew shift portal.

You can also access this portal from the **Main Menu > Resource Management > Procedure**. You are brought to a query portal with options for searching for a specific procedure. Once a procedure has been selected, you are brought to the maintenance portal to view and maintain the selected record.

Refer to [About Procedures](#) for more information.

The **Procedure** zone on the portal's **Main** tab page provides information about the procedure.

Timesheets

About Timesheets

A timesheet record summarizes the time a mobile worker spent while working a shift on performing work, traveling, attending meetings, and other broad categories of time usage.

At shift completion time, the system automatically creates a timesheet record for each period of time a mobile worker was part of the crew on that shift, allowing them to review, accept, and sometimes change their reported time information.

If the crew uses the mobile application, timesheet records for each mobile worker are prepared on the server and sent to the mobile device for review and approval. Crews that do not use the mobile application to work their shift may review and approve their timesheet records in the office when they complete their shift.

NOTE: Timesheet functionality is optional. [Global configuration](#) parameters must be specified to enable it.

This section describes concepts related to managing timesheets.

Timesheet Period

A crew on a specific shift may involve multiple mobile resources that may join and leave at different times. For example, one of the workers may leave the shift in mid day while another may join the shift later. In the same way, if the crew's vehicle breaks down in the middle of their shift they may switch to another vehicle. The system keeps track exactly when mobile workers join and leave a crew's shift, as well as tracks when vehicles are utilized by the crew on the shift. These period of times are recorded in the shift's crew allocation history information.

A timesheet record summarizes time usage for an individual mobile worker for each period of time they were part of a shift. Most of this would mean one timesheet record per mobile worker per shift but the system is designed to handle changes in resources throughout the shift.

NOTE: While base functionality mainly focuses on preparing timesheet records for mobile workers, the application can be extended to also manage vehicle utilization timesheet records if needed.

Time Usage Summary

Every minute during the timesheet period of time is summarized into a corresponding time utilization category, a "bucket", such as travel, work, etc., providing a breakdown of total time spent for each bucket.

For example, every time the crew indicated they went en route to a certain task location and until they indicated to are no longer traveling, that duration of time is added to the "travel" time bucket. In the same way, time they spent working on an activity on site is added to the "work" time bucket.

Time spent traveling to and working various types of tasks can be configured to be grouped into specific custom time "buckets" as part of each task type definition. In order to interface with your timekeeping system, these buckets should correspond to how the company wants to track time. For example, you may want to group time for all meter work in one bucket (both travel and work time), collections work in another bucket, and repair / construction work in a third bucket. Alternatively, you may want to group all travel into one travel bucket and all work in one or more work buckets.

NOTE: Time "bucket" categories are defined in **Timesheet Time Type (M1-TimeSheetTimeType)** extendable lookup along with a description of how they are used.

Generating Timesheets

When this functionality is enabled, the system initiates a process to generate timesheet records for the entire crew at the end of their shift.

Information about how time was spent during an entire shift is gathered from all the state transition log records recorded for the shift and those recorded for all the tasks the crew spent time on during the shift. When the crew uses a mobile device, actual event times are needed from logs captured on the device and sent to the server throughout the day until the shift is completed.

Timesheets can only be prepared once all this information is available on the server. Typically all information is available very close to ending the shift, but there could be exceptions that could some delays.

The process to generate timesheets for the entire crew works as follows:

- Wait for all shift and task information to be available on the server
- Once available, create a timesheet record for each mobile worker's period of time they worked as part of the crew on the shift. At this point, timesheets are ready for review.
- If the crew uses a mobile device and has indicated at end of shift that they are interested in reviewing their timesheets on the device, timesheets are sent to the crew's mobile device for review.

NOTE: Logic to request timesheets for the entire crew resides in an algorithm when the shift is completed. The request itself is handled by the "Crew Timesheet Process" business object. The calculation of an individual mobile worker's timesheet record is managed by the "Resource Timesheet" business object.

Timesheet types define attributes common to timesheet records of the same type. For example, you would need a timesheet type for managing the request to process timesheets for the crew and another timesheet type for the individual mobile worker timesheet records.

Refer to [Understanding Timesheet Types](#) for more information.

Review and Corrections

A timesheet record is initially created in the pending state, waiting for the mobile worker to mark it as reviewed.

Calculated time frames are strictly based on recorded system events. If for some reason the system times do not accurately reflect actual events, then depending on your business practices and settings, mobile workers may or may not be allowed to adjust their own calculated times as part of the review process.

NOTE: Use the [Timesheet Correction global configuration](#) option to allow or not allow mobile workers to correct the system calculated times.

Corrections to a system generated timesheet record are limited to adjusting time "buckets" and not to correcting the timesheet period itself.

If the mobile worker does not agree with their timesheet report or cannot correct the data, they can choose not to mark it reviewed and escalate the issue to their supervisor. The mobile worker or supervisor would have to make the correction in the office after returning from the field. If the problem is with the overall timesheet period, a new timesheet record would need to be created to reflect the correct period of time.

For example, if a mobile worker joined the shift at noon but the crew leader recorded the crew allocation change at 1 PM, then the system generated timesheet would show that the worker joined at 1 PM. Even if the mobile worker is allowed to adjust their time "buckets", they are not allowed to correct the overall time of their timesheet.

Crews using a mobile device are asked as part of the shift completion process whether they would like to review their timesheets on their device or later in the office.

Refer to [Reviewing Timesheets](#) for more information on how timesheets are handled by the mobile application.

Crews not using a mobile device can complete timesheets "in office" as part of their general shift completion workflow using the timesheet zone on the shift portal. The mobile worker may click on the **Process Crew Timesheets Now** button to generate the timesheet record for each individual resource that worked this shift or wait for the next timesheet batch process to do so.

Maintaining Timesheets

This portal is used to view and maintain a timesheet record. Timesheets are typically viewed in the context of a specific shift so you most likely navigate to this portal from the "Timesheets" zone on the crew shift portal.

Refer to [About Timesheets](#) for more information.

You can also access this portal from the **Main Menu > Resource Management > Timesheet**. You are brought to a query portal with options for searching for a specific timesheet. Once a timesheet has been selected, you are brought to the maintenance portal to view and maintain the selected record.

The **Timesheet** zone on the portal's **Main** tab page provides information about the timesheet, including the various maintenance actions applicable to its current state.

Mail

About Mail

The main purpose of the mail management system is to allow dispatchers to communicate with crews without leaving the application, as well as allow crews to respond to and communicate with dispatchers from their mobile device.

Messaging is context driven, providing dispatchers control over communication with crews based on their shift, availability, and whether or not they are logged on. They can choose to send messages as normal delivery, as emergency mail, or with other special handling such as setting the message to require manual acknowledgement or setting the message to expire after a certain amount of time without being acknowledged.

Any user can send mail to any other user or to a predefined mailing list. If the user is a mobile worker, the system knows to direct the mail message to their mobile device if they use the mobile application to perform their work.

Additional context driven functionality supports dispatchers and crews. Dispatchers can send a mail message to all of crews or only those they monitor. In the same way, crews can send mail to all logged on dispatchers or only those actively monitoring them. A crew can also send mail to a specific user, typically their supervisor.

Users have the option of selecting from a mail template to send preset mail content for specific situations, or they can create unique messages from scratch.

You can manage your mail using the [Mail Summary](#) dashboard zone.

This section describes concepts related to managing mail messages.

Emergency Messages

An emergency message is sent to a crew's mobile device in higher priority than regular messages. On the device, the message pops up over the current screen to grab the crew's attention. Messages can also be tracked to show whether or not they have been acknowledged.

Mail Summary

You receive notifications of new messages sent to you via the Mail Summary zone on dashboard.

From that zone you can also compose a new message as well as navigate to the My Mail portal to manage all of the mail messages.

Maintaining Mail

This portal is used to view and maintain one of your mail messages.

Refer to [About Mail](#) for more information.

You can access this portal by navigating to "My Mail" portal using the **Mail** link on the "Mail Summary" dashboard zone. You are brought to a query portal with options for searching for a specific mail message. Once a mail message has been selected, you are brought to the maintenance portal to view and maintain the selected record.

You can also access the "My Mail" portal from the **Main Menu > My Mail**.

The My Mail portal provides access to view mail you have received or sent. You can create new mail, reply, acknowledge mail you received, and archive your mail using this portal.

The **Mail** zone on the portal's **Main** tab page provides information about the mail message, including the various maintenance actions applicable to its current state.

In addition to the standard add menu option, you can start a new message in one of several ways:

- Dispatchers can right-click on a shift to send mail to the entire crew from the CDI portal.
- You can click the **New** button on the "Mail Summary" dashboard zone

Transfer of Goods

About Transfer of Goods

An activity may request to deliver goods to a certain location, for example to a customer. The reverse is also supported, a request to pick up goods from a certain location. Transporting goods from one location to another requires special scheduling and handling considerations around questions like:

- What are the dimensions of the delivered items? Large or heavy items may only fit in certain vehicles. Similarly, the system needs to make sure there is room available in the vehicle, considering other goods already on board the scheduled crew's vehicle.
- What can and cannot be shipped along with these goods? For example, items made of combustible materials should not be loaded together with matches.
- Where should the crew pick up the goods for a delivery activity? Where to unload goods for a pickup activity? If goods are delivered from or collected to a depot, what are the business hours of the depot?
- and more.

This section describes concepts related to managing activities that involve transfer of goods.

The Goods

Activities that handle goods specify information related to them such as the list of items being delivered, [load capacities](#) (weight, volume, size, etc) and restrictions relative to other goods on board. For example, an activity would indicate if hazardous load combinations should be avoided.

When applicable, the activity may indicate the following restrictions around transfers and goods on board:

- A list of other types of goods that **must not be** on board with this activity's goods.
- A list of other types of goods that **preferably should not be** on board with this activity's goods.
- No other goods must be on board, only this activity's goods.
- If a transfer is made using a pair of activities in a job or it involves a depot, the activities may indicate that no other activities may be scheduled in the route between the pickup and the delivery, i.e. make the transfer continuous without any other stops in between.
- Set a maximum transport time until goods reach their destination. For example, this would be important for items that require refrigeration or that are otherwise perishable or volatile.
- A broad category of the load may be specified to ensure the scheduler assigns the work to a crew whose vehicle is capable of handling this type of load. If goods are transferred from or to a depot this information is also used to ensure the depot can handle such load. Refer to [Load Compatibility](#) for more information.

If the [transfer](#) is from or to a depot, the activity also specifies:

- At least one valid depot where goods might be distributed from (for deliveries) or collected to after pickup

- An estimated time for how long it would take for the crew to load or unload the goods at the depot.

Refer to [About Depots](#) for more information.

The Transfer

A delivery activity specifies the location to where goods should be delivered, whereas a pickup activity specified the location from where goods are picked up.

The system supports the following transfer options:

- Use a pair of activities to define a transfer, one activity to pick up the goods from one place and another to deliver the goods to another place. These activities would have to be scheduled to the same crew shift in a specific order, pickup then delivery. Refer to [Job](#) for more information.
- If goods are to be loaded from a depot, the activity must specify one or more depots to choose from.
- If goods are to be collected to a depot, the activity must specify one or more depots to choose from.

As a rule, a crew's vehicle starts empty and should be empty again at the end of their shift. Goods must be delivered during the same transfer effort that picked them up. For example, assume a job with one activity to pick up an item from location A and another activity to deliver that item to location B. The total goods being transferred must be 0 at the end of the job, 1 picked up - 1 delivered = 0 on board at the end of that job.

Refer to [About Depots](#) for more information about depots.

Load Capacity

Specifying the overall capacity of the load on the activity is needed when the [maximum capacity](#) on the vehicles used to transfer the goods or [depot capacity](#) limits are to be considered.

Load capacity is not limited to any specific dimensions. The system supports up to 5 cumulative and 5 non cumulative types of capacities of your choice and applicable to your business.

Cumulative refers to dimensions like weight, volume, etc that can be added up for entire load to check against the vehicle's maximum capacity limit of this type if any. For example, if the maximum weight allowed on the vehicle is 10 tons, all of the activities' goods cannot weigh more than 10 tons.

Non cumulative may be for example length, width, etc. and for these, the maximum value of any of the goods is compared against the vehicle's maximum capacity limit, if any. If the vehicle's container size is 20 meters in length, than each of the activity's goods must be 20 meters or less in length.

The first cumulative capacity describing the load is also used as the quantity of goods being transferred for the purpose of calculating how long it would take to transfer the goods at the bay given the bay's service rate. Refer to [Handling Various Categories of Load](#) for more information.

NOTE: The types of load capacity limits to consider such as weight, volume, size, etc. must be configured before vehicle limits functionality can be used. Refer to [Defining Types Of Load Capacity](#) for more information.

Vehicle Maximum Capacity

For some types of goods being transported it is important to know how much of a load the transport vehicle can handle such as maximum weight, maximum volume, maximum size or any other [load capacities](#) as needed.

Maximum capacity of vehicles is defined on the vehicle type, not the crew shift. Any shift that references the vehicle type is bound by those capacity restrictions. When multiple vehicles are used on specific shift, the maximum capacity is set to the total of all vehicles' cumulative capacities, and the lowest limit for the non cumulative capacities.

However, this can be overridden on a specific shift to handle exceptions. For example, if a dispatcher is aware that a truck has low tire pressure, he might override the maximum capacity to indicate that the truck cannot carry such a heavy load.

Load Compatibility

Transferring some types of goods to and from depots may call for a vehicle that can carry such goods as well as scheduling the work to a bay at the depot that is equipped to handle it. This type of compatibility is enforced by the scheduler when the activity specifies the category of the transferred load.

When scheduling a depot related activity that involves the transfer of goods of a specific load category, the scheduler only considers crews whose vehicle can handle such load.

Refer to [Vehicle Capabilities](#) for more information.

In the same way, when scheduling a visit to a depot the scheduler would only consider depots that can handle the load category associated with a depot related activity or a depot transfer activity.

Refer to [Handling Various Categories of Load](#) for more information.

About Depots

If your business handles delivery and pickup activities where goods are loaded or unloaded at particular centers, your organization can set up depots to capture the information related to these centers, such as their physical address, the days and times the depot is open, and more. This section describes concepts related to managing depots.

Location

A depot record defines the physical address of the center by referencing a common [location](#). The scheduler uses this information to plan a trip to the depot's location to load or unload goods as needed when scheduling delivery and pickup activities respectively.

Depots may serve distribution and collection operations. A depot where crews pick up goods for delivery is referred to as a distribution depot whereas a depot where collected goods are dropped off is called a collection depot.

NOTE: If an actual depot participates in both collection and distribution operations then two depots need to be set up for the same location.

Depot Time Windows

Access time windows define the business hours when the depot is open. Each access time window also has an associated window cost and late cost. The window cost indicates the cost of using the time window, whereas the late cost indicates the cost of being late and arriving at the depot outside of the allotted time frame.

Service time windows define the preferred business hours within the normal operational hours. Lower cost is associated with scheduling work during preferred service times.

A work calendar may be specified to define public holidays when the depot is closed.

The scheduler uses this information to calculate the actual time windows during which depot tasks to visit the specific depot should be scheduled.

Multiple depots can share this information by referring to a [depot profile](#).

The system supports exceptions to time windows when the depot is open or closed.

Refer to [Depot Time Window Overrides](#) for more information.

Depot Capacity

A depot's scheduling capacity for handling goods may be controlled either in the form of the number of bays it has and their business hours or as an explicit maximum capacity limit associated with these business hours. These two mechanisms are mutually exclusive, a depot either uses bays and their time windows to limit capacity or it defines maximum capacity for specific time windows.

Refer to [Bays](#) for more information.

Maximum capacities may be specified for an access time window to indicate total capacity that can be handled during that time window at the depot, for example maximum weight, maximum volume, or any other type or capacity as configured by your implementation.

The scheduler uses maximum capacity information to ensure that the total goods associated with all depot tasks scheduled to a depot time window does not exceed this capacity limit.

NOTE: The types of maximum capacity limits to consider such as weight, volume, size, etc. must be configured before limits can be specified for a depot time window. Refer to [Defining Types Of Load Capacity](#) for more information.

Depot Site Delay

There may be a fixed delay time that is not related to the actual loading or unloading of the goods when visiting a specific depot. For example, the amount of time needed to access the depot center on each visit, regardless of specific goods being transferred. The scheduler uses this information when planning the [duration](#) of the crew's visit to the depot.

If handling of goods takes place at designated bays, a site delay may also include a fixed delay time during which the crew occupies the bay prior to loading / unloading. For example, it may be used to reflect quality control delay prior to unloading goods at bay.

Multiple depots can share this information by referring to a [depot profile](#).

Depot Task

When scheduling a delivery activity, the scheduler plans a task for the crew to visit one of the depots mentioned on the activity and pick up the goods prior to the delivery activity's scheduled time.

In the same way, when scheduling a pick up activity to collect goods from some location the scheduler also plans a task for the crew to drop off the collected goods at one of the depots specified on the activity.

Tasks to visit depots for loading and unloading purposes are generated by the scheduler and incorporated into the crew's shift route, like any other task. As with any task they perform, crews can provide status updates related to a depot task, allowing the system to track the crew's travel, arrival, and completion of the depot visit.

The scheduler creates a depot task using the task type associated with the depot record, its [depot profile](#) if any, and otherwise if not specified, it is defaulted from the [global configuration](#) record.

For example, you might have a task type to pick up keys, one to pick up hazardous goods, and one to drop off goods. Using a designated task type allows the description of the depot task everywhere in the system it is displayed to more clearly convey the purpose of the visit to the depot.

The complete transfer operation that includes the depot related activity (delivery or pickup) and its corresponding depot task (to load or unload the goods) is referred to as a depot run.

Refer to [About Depot Runs](#) for more information about depot tasks and the entire transfer run.

Depot Profile

A depot may refer to a depot profile for shared settings of certain details.

Depot profiles can be set up to establish common settings for depot task type, site delay, access time windows, and service time windows across multiple depots. All other details are defined on each individual depot.

If a depot refers to a profile, the settings established by the profile cannot be overwritten. However, anything that is changed on the depot profile affects all depots that references the profile.

Depot Time Window Overrides

An override time window is used to redefine a depot's access or service time window as originally defined on the depot. For example, if the dispatcher knows that the depot will be open later than usual on a particular day, he can override the default time window and adjust for the later time. The override will apply to a specific day and time for the depot rather than the generalized weekly schedule defined by the default time windows.

An override can also be used to completely remove a specific time window or to adjust the maximum capacities defined on the time window, such as the maximum size or weight of distribution that the depot can handle.

The host system may send in a depot time window override. Refer to [Depot Management](#) for more information.

Crews Can be Restricted to Specific Depots

Typically in depot based operations, crews logon and logoff from their base depots. Crews can be configured so that they are restricted to their base depots only. Restricting crews to specific depots may be necessary, for example, if the crew's vehicle is too large to visit other depots.

The scheduler is limited to schedule depot related work to shifts of crew that are allowed to visit these depots. If a crew is not limited to specific depots, it is assumed that they are allowed to visit any depot. A dispatcher can override the list of depots a crew is restricted to on a specific shift to handle exceptions that day.

Depot Breaks

Certain break types can be configured to be taken at a depot. If the shift is not scheduled for depot work, the break is scheduled as a standard break.

Bays

Depending on the nature of the operation at a depot, goods may only be handled in single queues at designated bays. The scheduling capacity of such depot is then dependent on the availability of these bays.

A depot's record indicates whether it uses bays, and when it does a record for each bay must also be defined to describe the bay's scheduling information.

NOTE: Changing the record of an existing depot to use bays requires the scheduler to be restarted if that depot is already referenced on tasks and other entities.

This section describes topics related to bay scheduling.

Bay Usage

All bays for a distribution depot may only be used for loading goods. In the same way, all bays for a collection depot may only be used for unloading goods. If a depot actually supports both loading and unloading bays then it should be defined as two separate depot records.

A Bay Represents a Single Queue

When a depot uses bays the scheduler schedules depot visits to load or unload goods at one of the depot's compatible and available bays. Tasks are scheduled one at a time in the form of a queue. If a compatible bay is not available the task remains not scheduled.

Refer to [Scheduling A Task To A Bay](#) for more information.

Handling Various Categories of Load

Bays at a depot may vary in the types of goods they carry and are equipped to handle.

Moreover, the amount of time it takes to transfer goods at the bay may be dependent on the type of goods as well as the specific equipment or capabilities available at the bay. For example, fluid goods may flow at different rates at different bays based on their infrastructure. For those types of goods, you may specify a rate of units per second as applicable for a load category and a bay.

The quantity used for the purpose of calculating the transfer duration at the bay is specified as the first cumulative capacity describing the load. Refer to [Defining Load Capacities](#) for more information.

The depot record may be used to specify a default list of load categories and their service rate which are applicable to all or most of the bays. The default may be overridden at the bay record for specific bays.

NOTE: Service rate is currently supported only for loading goods.

Bay Scheduling Restrictions

Scheduling work to depot bays involves complex scheduling decisions and therefore imposes some limitations.

The following limitations apply when leveraging bay scheduling:

- A delivery activity may reference either a single depot that uses bays or multiple depots that do not use bays. A depot transfer activity, for loading and unloading goods, may only reference a single depot that uses bays.
- If a crew is associated with an explicit list of depots they are allowed to visit on a shift then all depots must either use bays or not. A scheduler configuration parameter **Shift Depot Mode** states whether a shift without an explicit depot restriction is allowed to visit only depots that use bays or those that do not.
- The scheduling horizon for schedulers handling depots with bays may not exceed 4 days.
- Chooser, Appointment Booking and Conditional Booking requests are not supported for activities associated with depots with bays. Note that dispatchers can still manually allocate activities to specific shifts.
- The base product's mobile application does not support transfer work. It also does not handle bay specific workflows for a depot task scheduled to a depot that uses bays.

Arrival Delay

There may be a fixed delay time involved in traveling from the depot main entrance location to a specific bay.

The scheduler takes this information into account when planning the start time of the crew's visit at the bay. When specified, the task at the bay is planned to start once the crew arrives at the depot's main entrance and after they travel to the bay.

A default value for all bays may be specified as part of the depot record as well as a value for a specific bay.

Clearance Delay

There may be a fixed delay time involved after the crew leaves the bay but not yet ready to leave the depot. For example, time needed for clearance, etc.

The scheduler takes this information into account when planning the completion time of the crew's overall visit at the depot. When specified, the crew is planned to leave the depot after they leave the bay and spend the specified amount of clearance delay at the depot.

A default value for all bays may be specified as part of the depot record as well as a value for a specific bay.

Scheduling a Task to a Bay

Bay scheduling refers to the process of incorporating a crew's visit to a depot to load or unload goods into the schedule of a depot's bay. The purpose of this process is to assign a task that involves a transfer of goods to a compatible depot bay at an appropriate time taking into account bays limitations such as time windows, the goods they can service and other tasks sharing the bays.

Depot tasks are scheduled to bays if the depot uses bays. Depot transfer activities are always scheduled to bays, i.e. the depots involved in the replenishment job are assumed to use bays.

Refer to [About Depot Runs](#) and [Depot Replenishment](#) for more information about these tasks.

To plan the best work schedule for each bay, the scheduler matches activity details such as duration, the time frame available and the category of load being handled with information about the capabilities and availability of each of the depot's bays.

When a task is scheduled to a bay it is assigned a specific time within one of the bay's time windows during which the transfer of goods work should start and complete. Depot tasks are scheduled both to a crew's shift as well as a bay whereas depot transfer activities are only scheduled to a bay.

Common reasons for why a task remains unscheduled may be

- No scheduler covers the task's location
- The task is placed on hold or has data issues preventing the system from sending it to the scheduler
- No compatible bay is available that matches the activity's scheduling requirements
- The activity has expired and the activity type does not allow further time window extensions
- The activity is optional and may be pushed out by higher priority work

Until the task is queued for dispatch, its schedule is constantly being optimized by the scheduler. The task may be scheduled to a bay for a while, rescheduled to another, become free, rescheduled again, etc.

When the task is queued for dispatch its scheduled time is locked but the bay may still change up until the crew's arrival to the depot, allowing the scheduler more time to optimize bay assignments. While the bay decision can be improved, the scheduler would always keep an activity in this state scheduled to a bay.

Once the crew arrives at the depot, the scheduled bay is no longer changed by the scheduler, but a dispatcher may override the scheduler decision and request to lock the task to a specific bay.

At any point in time if the scheduler cannot schedule an activity that otherwise should be scheduled, the task is disabled from scheduling and the scheduling error is reported in the scheduler's registry record for the task. The system provides information on the activity portal schedule tab to assist you in determining why an activity is not scheduled or disabled. Once the issue is resolved, the activity is automatically rescheduled.

For more persistent errors, the scheduler monitor batch process converts registry errors to To Do entries for a user to follow up. For dispatchers, the To Dos In Scope dashboard zone lists these entries for activities they are monitoring.

NOTE: The *Scheduler Monitor* batch process needs to run at least once a day to convert scheduling errors to To Do entries.

Bay Scheduling Restrictions

Scheduling work to depot bays involves complex scheduling decisions and therefore imposes some limitations.

The following limitations apply when leveraging bay scheduling:

- A delivery activity may reference either a single depot that uses bays or multiple depots that do not use bays. A depot transfer activity, for loading and unloading goods, may only reference a single depot that uses bays.
- If a crew is associated with an explicit list of depots they are allowed to visit on a shift then all depots must either use bays or not. A scheduler configuration parameter **Shift Depot Mode** states whether a shift without an explicit depot restriction is allowed to visit only depots that use bays or those that do not.
- The scheduling horizon for schedulers handling depots with bays may not exceed 4 days.
- Chooser, Appointment Booking and Conditional Booking requests are not supported for activities associated with depots with bays. Note that dispatchers can still manually allocate activities to specific shifts.
- The base product's mobile application does not support transfer work. It also does not handle bay specific workflows for a depot task scheduled to a depot that uses bays.

Depot Replenishment

Depot stock may be replenished by transferring goods from one depot to another.

This functionality is optional and depends on the nature of your business and how it handles stock inventory. When used, the host system sends in detailed transfer requests between depots based on its knowledge of inventory levels, timings and needs.

NOTE: The purpose of interfacing such work is so that it can be incorporated into a depot's bay schedule. As such this type of work cannot be scheduled to crews and is limited to depots using bays.

Refer to [About Depot Transfer Activities](#) for more information.

About Depot Runs

Depot related activities involve either picking goods from a distribution depot and delivering them to a customer location, or picking up goods from a customer location and bringing them to a collection depot.

[Depot tasks](#) represent a scheduled visit to a depot, with a task duration for loading or off-loading the goods. The scheduler automatically creates these depot tasks to handle activities as needed.

Depot runs represent the entire trip or travel to visit a depot, including the depot task and one or more depot activities. A single run always involves one depot and as such, never mixes distribution and collection operations.

Refer to [About Transfer of Goods](#) for more information.

This section describes concepts related to managing depot runs.

Depot Task Duration

If the crew is scheduled to pickup or drop off goods for multiple activities, they truly only visit the depot once. In this situation, the scheduler consolidates the individual depot tasks into a single depot task larger in duration to better represent the single depot visit. Only the scheduler is responsible for creating, deleting, and consolidating depot tasks as it optimizes the scheduler.

If the depot uses bays then the depot task is also incorporated into the schedule of one of the depot's bays. Refer to [Bays](#) for more information.

The duration of the depot task reflects the total time the crew spends at the depot from when they arrive at the main entrance until they leave. Duration calculation includes the following time components:

- A [site](#) delay for visiting the specific depot once, if applicable.
- [Time to arrive](#) to the scheduled bay, if applicable.
- The total estimated time for loading or unloading goods as specified on each of the activities in the depot run. If scheduled to a bay then depending on the type of load handled, the duration of the actual transfer may be a function of the quantity requested on each activity and the bay's service rate for that load category. Refer to [Handling Various Categories of Load](#) for more information.
- [Time for clearance](#) after leaving the bay and before leaving the depot, if applicable.

For example:

- There is a site delay of 10 minutes to arrive at the depot and pass through security.
- It takes 5 min to get to a depot bay.
- The depot run has 4 delivery activities that each specifies 6 minutes for loading. Alternatively, assume one delivery activity in the run for 7200 gallons of petrol and the service rate at the bay is 5 gallons a second, which means it would take 24 minutes to load each activity at the bay.
- The crew has to fill out clearance forms for 10 minutes.

The estimated duration of the depot task is $10 + 5 + 24 + 10 = 49$ minutes.

Depot Task Workflow

As with other tasks associated with a location, the crew indicates they have started traveling to the depot location and once there, they can start performing the depot task.

The following describes key stages in the overall lifecycle of a depot task.

Pending

A depot task is initially created in this state waiting to be scheduled and [queued for dispatch](#).

Queued For Dispatch

Setting the task to the **Queued For Dispatch** state is the first step in the dispatch process.

This is similar to the process of queuing breaks for dispatch. Refer to the break task [Queued For Dispatch](#) state for more information.

Dispatched

This is similar to the process of marking breaks as dispatched. Refer to the break task [Dispatched](#) state for more information.

En route

The crew provides this status update when they are about to start traveling to the location. Task status updates by the crew are similar to assignment updates. Refer to the assignment [status update](#) for more information. If the crew was previously out of service, they are now considered in service.

Time spent while in this state is considered travel time for [crew time](#) reporting.

Arrived

Typically, the crew may start work immediately upon arrival to the depot. This state indicates that the crew has arrived at the location but has not yet started work for some reason. The crew provides this [status update](#) when they arrive at the depot but cannot start the work right away.

For example, if the depot uses bays the crew may need to wait to be told which bay they should go to. The scheduler is notified of the crew's arrival and locks the scheduled bay for the activity. The host system polls for this information and once available, conveys it to the crew. If the bay is not available at this time they crew may be asked to wait shortly.

Time spent while in this state is considered wait time for [crew time](#) reporting.

Started

The crew provide this status update when start the loading or unloading work at the depot. Task status updates by the crew are similar to assignment updates. Refer to the assignment [status update](#) for more information.

Time spent while in this state is considered depot time for [crew time](#) reporting.

Completed

This state indicates that the crew has completed the work and left the depot. Task status updates by the crew are similar to assignment updates. Refer to the assignment [status update](#) for more information.

Canceled

A depot task in this state indicates it is no longer needed and should not be scheduled.

Once a shift is finalized, depot tasks that were not started are automatically canceled.

Depot Task Actions

A dispatcher or a business process may perform various actions on a depot task that may impact the task's state and schedule.

The following section describes these actions and the business flows associated with them.

Close

Closing a depot run prevents the scheduler from making further changes to the set of activities in the run.

A depot task can be closed in the following ways:

- When it is advance dispatched. Refer to [Advance Dispatching](#) for more information.
- A dispatcher can manually close the depot task to finalize its schedule and prevent more work from being assigned to the run. The dispatcher can re-open the depot task at any time.
- As part of a depot cutoff request. Refer to [About Depot Cutoff](#) for more information.

Open

Opening a depot run allows the scheduler to further optimize the run and make changes to the set of activities scheduled to it.

Lock To Bay

This action is only applicable if the depot task is scheduled to a depot bay and the crew has already arrived at the depot. It allows a user or a process to explicitly allocate the task to a specific bay.

About Depot Cutoff

A depot cutoff event closes all distribution depot runs that are planned to depart from the depot before a specified cutoff time, and instructs the scheduler to build new runs for the depot to depart after that cutoff time. The depot cutoff automatically locks all such runs from the distribution depot.

As an example, think of depot cutoff functionality applied for grocery delivery. All deliveries for a given day must be on the truck and ready to go by 10 AM to make the "early run". To allow enough time for the picking process to prepare the trucks, all orders included in the 10 AM run should be processed no later than 5 PM on the previous day. Any orders processed after 5 PM would be pushed to later runs.

To achieve that, a depot cutoff for tomorrow at 10 AM would be processed at 5 PM today. When processed, the depot cutoff closes all depot runs planned to departure before 10 AM so that the scheduler does not add or remove activities to and from each run. At this point, new activities can only be scheduled to depot runs planned to depart after 10 AM.

Once depot runs are closed for a specific cutoff time, the host system is provided with the scheduling manifest that includes the list of these runs and the activities scheduled to each in order to start with the picking process and preparation of the trucks.

NOTE: Cutoffs only apply for distribution operations hence only to distribution depots.

This section describes concepts related to depot cutoff functionality.

Processing A Cutoff Event

A depot cutoff event specifies a cutoff time as well as a trigger time, indicating the time at which the cutoff event should be processed. If no trigger time is provided, the cutoff is processed right away.

When a cutoff is processed, all depot tasks scheduled to depart on or before the cutoff time are automatically closed thus locking their schedule. The latest cutoff time for a depot is posted on the depot record notifying the scheduler that work can only be scheduled to later depot runs.

Only one cutoff event may be processed at any point in time for a depot and each cutoff time has to be later than a previously processed cutoff.

Regardless of how the cutoff event was [initiated](#), when the cutoff event processing is complete the corresponding [scheduling manifest](#) for all the depot run it has closed may be automatically sent to the host system based on depot configuration.

NOTE: Processing of depot cutoff events in a timely manner relies on the *Depot Time Window Monitor* batch process to run frequently. Refer to [M1-DPTTW](#) batch control for more information.

Initiating Cutoff Events

Depending on the integration with the host system, cutoff events for a depot can be configured to be generated in advance from within the application or requested on demand from the host system.

Generation of cutoff events is based on a weekly schedule defined on the depot or the depot profile it uses. The schedule sets predefined cutoff times per weekday, as well as a trigger offset prior to each cutoff time to state when the event should be processed. In addition, the schedule defines how many days in advance to generate cutoff events. A cutoff event record is generated for each cutoff time in the schedule for the cutoff horizon specified.

NOTE: Logic to generate depot cutoff events relies on the depot being periodically monitored by the **Depot Monitor** batch process. Refer to [MI-DEPOT](#) batch control for more information.

Cutoff events may also be initiated by the host system on demand using the web service. Refer to [Depot Management](#) for more information.

A user may also manually initiate a cutoff event as needed by using the **Cutoff** action on the depot portal.

Scheduling Manifest

A scheduling manifest for a depot cutoff includes information about all shifts that are associated with depot runs closed by the depot cutoff request, as well as all the activities scheduled to each depot run.

Depot configuration can indicate whether or not a scheduling manifest should be automatically sent to the host system when the processing of a cutoff event is complete.

Depending on the integration with the host system, the host system may instead request for the manifest on demand. This is typically done when the host also initiates the cutoff event, polls to check its state periodically, and when complete, requests for its corresponding manifest.

Close Entire Shift

With depot cutoff events, earlier depot runs are closed while later runs remain open. This type of operation can thus accommodate same day changes as depot runs are still open later in the day.

If your business operations needs to close and lock the schedule of all runs for a crew's shift some time prior to the shift's planned start time, you don't need to use depot cutoff events at all. Instead you should leverage the advance dispatching functionality and set up your shifts to use the "close and dispatch all" advance dispatch mode. When using this mode, the system locks down the entire schedule for the shift, including all depot runs scheduled to it, at the specified time prior to shift start.

Refer to [Advance Dispatching](#) for more information.

About Depot Transfer Activities

Depot transfer activities may be used to transfer goods from one depot to another.

This functionality is optional and depends on the nature of your business and how it handles stock inventory. When used, the host system sends in transfer requests that specify the exact goods and quantity that should be transferred within a specific time frame as well as the source and target depots. Transfers may only be requested by a host system based on its knowledge of inventory levels, timings and needs.

Transfer activities are interfaced, validated and monitored until completion, similar to regular activities. However, they are scheduled and worked in a completely different way.

Refer to [About Activities](#) for more information.

This section describes concepts related to managing transfer activities.

A Depot Transfer Is a Job

A depot replenishment transfer request is actually a pair of depot transfer activities bound as a job. The first activity manages the task to load the goods at the source depot and the other activity handles the unloading of the goods at the destination depot.

Refer to [Job](#) for more information

Scheduled to Bays, Not to Resources

Unlike a regular activity, a transfer activity is not scheduled to a crew shift but rather to a depot bay. This is the most significant difference between transfer activities and regular activities as it impacts the entire work flow and behavior of such tasks.

The scheduling process is not responsible for finding a compatible resource to assign the transfer activity to but rather to book time for handling the goods, loading or unloading, at a compatible bay at the depot. A regular activity is incorporated into the schedule of a crew on a specific shift while a transfer activity is incorporated into the schedule of a specific bay at a specific time window at the depot.

Refer to [Bays](#) for more information.

The scheduler is completely agnostic of who would actually perform the transfer work. Managing the allocation of work to resources to perform transfer work is managed outside of the application.

Refer to [Depot Transfers Are Managed Externally](#) for more information.

Depot Transfers Are Managed Externally

The host system is responsible for figuring out who is to perform each depot replenishment transfer and how to handle all communications with the assigned driver. Typically this type of work is given to contractor companies that distribute the work amongst their fleet of drivers.

Drivers performing depot replenishment work are not defined up front in the application nor have any shifts planned for them. They are assumed to perform transfer jobs as one off assignments that are not part of any broader planned route.

Interaction with the driver, from their arrival to load the goods and until the transfer job is complete, is also managed by the host system. Status updates related to the load and unload activities are sent in from the host to properly reflect their progress.

When a driver reports to the loading depot, the host system sends in a status update for the loading activity indicating that the crew has arrived at the depot to perform it. It is at this point only that the system is notified of the identity of the assigned driver and associates it with the transfer activities. This information is required for tracking and reporting purposes only.

The execution of the transfer job once started as well as the travel between the depots is tracked as follows:

- A crew record for the reporting driver is created, if such record is not already defined for them. The host system may or may not send in along with the transfer job also the specific contractor company that is responsible for performing it. If this information is available the system associates the crew record with the company they work for.
- A transfer shift for the crew is automatically created to track the crew's loading and unloading assignments as well as their current location along the way.

NOTE: The driver may be an internal company worker. However, since depot transfer work is not scheduled against resource capacity the system will not schedule the work to any of the worker's shifts nor would it ensure they are indeed available to perform the work.

Refer to [Transfer Workflow](#) for a better understanding of how a depot replenishment job is worked.

Total Break Time

While driving from the loading depot to the unloading depot, the driver may be entitled for a certain amount of break time. Depot transfer work is not scheduled to crew shifts therefore the existing functionality handling planned shift breaks is not applicable.

Refer to [Scheduled To Bays, Not To Resources](#) for more information.

Instead, an estimated total break time may be sent in from the host system along with the transfer unload activity information. When specified, the scheduler adds the specified break time to the overall drive time between the depots when estimating the driver's planned arrival time to unload the goods.

Refer to [Estimated Arrival Time To The Unloading Depot](#) for more information.

Estimated Arrival Time to the Unloading Depot

The scheduler initially plans the crew's arrival to the unloading depot based on their travel time as well as any break time they are planned to take along the way.

However, once the crew has started the travel, the application has no knowledge of how much break time was actually taken. As the crew gets closer to the depot they are most likely to have taken all their planned break time and therefore their arrival time can be estimated solely based on their current location and the remaining travel time from there to the unloading depot.

When the crew starts a transfer job at the loading depot, the system creates a transfer shift record to track their progress and travel between the depots. The system relies on the transfer shift's offline state to indicate whether or not the crew has taken all their breaks. When the shift is offline, the crew is assumed to have not yet taken all their break time and their arrival time is calculated based on travel and break time only. If the shift is set to online, the crew is assumed to have taken all their break time and their current crew location is used to estimate their arrival time.

NOTE: By default the transfer shift is created in the offline state. To leverage crew location reports for a more accurate estimate of the crew's arrival time to the unloading depot, custom rules need to be implemented to mark the crew's transfer shift as being online at the appropriate time.

Refer to [Crew Location](#) for more information.

Crew Location

A transfer job starts when a driver arrives at the depot to load the goods. Once loaded the driver then travels to unload them at another depot. GPS reports of the crew's current location during this travel may be sent in to increase visibility as to their whereabouts. These reports need to be associated with the driver's crew record.

Refer to [Depot Transfers Are Managed Externally](#) for more information.

NOTE: Unlike regular activities, the scheduler does not estimate a crew's current location at any point in time for a crew performing a depot transfer job. Their current location is solely based on GPS reports. In absence of GPS reports, the last known location of the crew is the depot where they have loaded the goods.

Transfer Shift

Unlike a regular crew's shift, a transfer shift does not represent crew capacity and is not sent to the scheduler for planning. A transfer shift is created automatically by the system when a driver arrives at a loading depot to start a depot transfer job for tracking purposes only. Such shift can only be associated with depot transfer work.

Transfer Task Duration

Similar to a regular activity, the transfer activity may specify the estimated amount of time it should take a crew to perform the requested work. If not explicitly specified, this information is defaulted from the activity type.

The duration of a transfer task reflects the total time the crew spends at the depot from when they arrive at the main entrance until they leave. Duration calculation includes the following time components:

- A [site](#) delay for visiting the specific depot once, if applicable.
- [Time to arrive](#) to the scheduled bay, if applicable.
- The estimated time for loading or unloading goods as specified on the activity.
- For loading goods only, and depending on the type of load handled, the duration of the actual transfer may also be a function of the quantity requested and the bay's service rate for that load category. The calculated transfer time by service rate is added to the estimated duration on the activity. Refer to [Handling Various Categories of Load](#) for more information.
- [Time for clearance](#) after leaving the bay and before leaving the depot, if applicable.

For example:

- A depot transfer job requests to deliver 7200 gallons of petrol from one depot to another.
- There is a site delay of 10 minutes to arrive at both depots and pass through security.
- It takes 5 min to get to a depot bay at both depots.
- The loading activity is sent in with an estimated minimal duration of 5 min. The service rate at the loading bay is 5 gallons a second, which means it would take additional 24 minutes to load the fuel.
- The unload activity is sent in with an estimated duration of 20 minutes,
- The crew has to fill out clearance forms for 10 minutes in both depots.

The estimated duration of the loading activity is $10 + 5 + (5 + 24) + 10 = 54$ minutes.

The estimated duration of the unloading activity is $10 + 5 + 20 + 10 = 45$ minutes.

Transfer Job Workflow

Depot transfer activities are performed as one off jobs that are not incorporated as part of any crew's shift route. The host system is responsible for assigning transfer jobs to drivers, letting them know when to report to the loading depot as well as tracking their progress once they start the job. A driver assigned to perform a transfer job arrives at the specified time to the one depot, loads the goods and immediately continues from there to unload the goods at the other depot.

Refer to [Scheduled To Bays, Not To Resources](#) for more information.

NOTE: Depot replenishment work is not handled by the base product's mobile application. All status updates related to the progress of the load and unload activities are assumed to be reported by the host system.

The following describes the overall steps involved in handling a depot transfer job:

- **Request.** A transfer job is interfaced from the host system, containing a pair of load and an unload depot transfer activities. The activities describe the amount of goods, the depots involved and the time window for performing the work.
- **Schedule.** The scheduler incorporates the transfer activities into the appropriate depots' bay schedule and continues to improve on its scheduling decisions as part of an ongoing optimization process.
- **Lock Load Time.** Some amount of time prior to the loading activity's earliest time window, the activity is [queued for dispatch](#) and its schedule is locked in time. Notice that while the time window is locked the system still allows bay assignment optimization changes. Refer to the [Being Scheduled](#) state in the transfer activity workflow for more information.
- **Communicate Time To Drivers.** The host system requests the [schedule manifest](#) of all transfer load activities for a specific depot that are queued for dispatch in order to communicate their locked scheduled time to the contractor company or driver they have allotted the work to.
- **Load The Goods.** The following describe the overall execution of the loading activity:
 - The assigned driver arrives at the depot, states which activity they came to perform and waits to be told the bay to go to. Crew arrival is reported by the host system using a web service (Start Depot Transfer Job). From this moment on, the system keeps track of the crew's progress in performing the activity in the form of an assignment, setting it to the arrived state.
 - Once notified of the arrival of the driver, the scheduler locks the activity's schedule to a specific bay.
 - The host system polls for the confirmed bay using a web service (Task Bay Inquiry). Once confirmed the driver is informed of the bay to use.
 - The driver gets to the bay and starts to load the goods. The host system reports that the crew has started with the loading which then sets their assignment to the Started state.
 - When goods are loaded the driver gets back to the depot entrance to complete any clearance forms. When the host system reports the crew has completed this step the loading activity is completed.
- **Unload The Goods.** The following describe the overall execution of the unloading activity:
 - Once the load activity is complete the unload activity is dispatched, creating an assignment to track the crew's progress transferring the goods and unloading it at the target depot. The unloading assignment is automatically set to the Enroute state.
 - The driver arrives at the depot, states which activity they came to perform and waits to be told the bay to go to. Crew arrival is reported by the host system which sets their unloading assignment to the Arrived state.
 - Once notified of the arrival of the driver, the scheduler locks the activity's schedule to a specific bay. The host system polls for the bay for the activity and once locked the information is shared with the driver.
 - The driver gets to the bay and starts to unload the goods. The host system reports that the crew started the unloading which then sets their assignment to the Started state.
 - When goods are fully unloaded the driver gets back to the depot entrance to complete any clearance forms. When the host system reports the crew has completed this step the unloading activity is completed.

The Dispatch Process

Depot transfer activities are performed as one off jobs that are not incorporated as part of any crew's shift route. The job involves loading goods at one depot and immediately traveling to another depot to unload them.

Refer to [Transfer Job Workflow](#) for more information.

The dispatch process for the loading and unloading activities is different than regular activities.

The following describes the overall steps involved in dispatching a loading transfer activity.

- **Step 1 - Queue For Dispatch.** Drivers need to be notified in advance of transfer jobs assigned to them so they can plan accordingly. A loading transfer activity starts the transfer job and as such its scheduled time has to be locked so that the host system can communicate the information in advance to an available driver. This is achieved by placing the activity in the [queue ready for dispatch](#). This step is performed in batch, but other processes may do so as well.

NOTE: This step relies on the task being periodically monitored by the *Task Monitor* batch process. Refer to [M1-TSKTR](#) batch control for more information.

- **Step 2 - Wait For A Driver To Arrive.** A loading transfer activity remains queued until a driver arrives at the depot ready to start the job. The host system is responsible for initiating this event.
- **Step 3 - Dispatch.** Once notified by the host system that the driver is ready to start the job, the system creates an assignment for the loading activity and sets the activity to the [Assignment In Progress](#) state where it waits for the outcome of this assignment.
- **Step 4 - Acknowledgement.** The mobile application does not support crews working transfer activities. As such the dispatch of their assignment is automatically considered acknowledged.

The unloading activity is automatically queued for dispatch and immediately dispatched when the loading activity is completed. The system creates an assignment for the unloading activity and sets the activity to the "Assignment In Progress" state where it waits for the outcome of this assignment.

Refer to [About Depot Transfer Assignments](#) for more information.

Transfer Activity Workflow

A transfer job is interfaced from the host system as a pair of load and an unload depot transfer activities.

Depot transfer activity states are depicted in the following diagram:

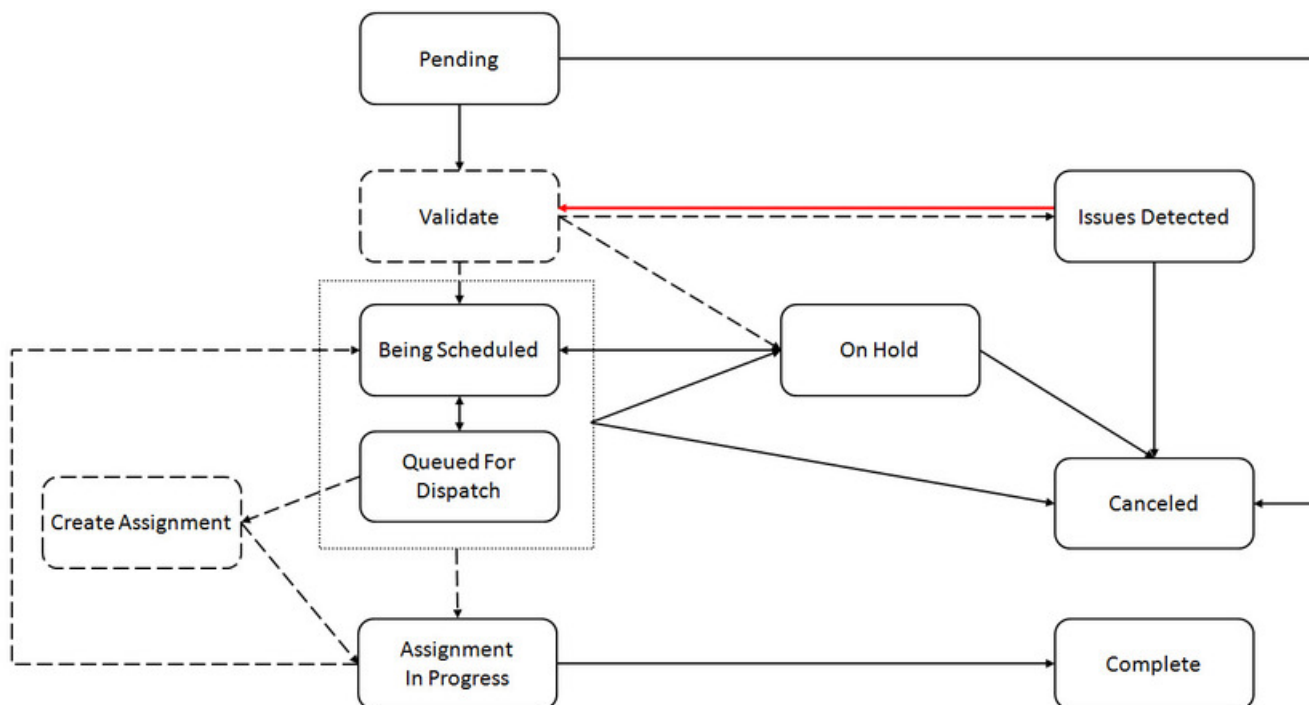


Figure 8: Depot Transfer Activity States

While the overall behavior of a transfer activity is very different from a regular activity, it does share similarities mainly around creation, validation and completion states.

This section describes key stages in the overall lifecycle of a depot transfer activity.

Validation

A transfer activity is validated very much like a regular activity. Refer to Activity Workflow for more information.

Issues Detected

A transfer activity handles data issues very much like a regular activity. Refer to Activity Workflow for more information.

Being Scheduled

A transfer activity in this state is available for scheduling but its schedule is not finalized yet. Unlike a regular activity, a transfer activity is not scheduled to a crew shift. It is instead incorporated into the schedule of one of the depot's bays.

Refer to [Scheduled To Bays, Not To Resources](#) for more information about this process.

While the transfer activity is in the **Being Scheduled** state:

- The scheduler keeps improving on its scheduling decisions. During this continuous optimization phase, the activity's schedule may frequently change between being assigned to a bay, reassigned to another bay, its time changed within its currently assigned bay's time windows, as well as become unassigned again.
- For a loading activity, this goes on until it is time to [queue the activity for dispatch](#) so its scheduled time can be communicated to a driver by the host system.
- For the unloading activity part of the transfer job, this goes on until the loading activity is complete which automatically dispatches the unloading activity.
- The activity's time windows are checked and if the activity has expired, it may be automatically extended or [canceled](#) as per configuration on the activity type. Refer to [Expired Activities](#) for more information.
- The activity's time windows are recalculated to take into account new or canceled applicable override time window events, if any exist. Refer to About Override Time Windows for more information.
- Business rules may put the activity on hold.

NOTE: Many of these business rules rely on the task being periodically monitored by the *Task Monitor* batch process. Refer to [MI-TSKTR](#) batch control for more information.

Queued For Dispatch

Setting the activity to the **Queued For Dispatch** state is the first step in the dispatch process. The scheduler stops further optimization of the activity's scheduled time so it can be communicated to the assigned driver.

Note that while the time is locked when the activity is queued for dispatch, its scheduled bay may still change to allow the scheduler more time in finding an optimal queue the crew should go to right until they actually arrive at the depot.

Refer [The Dispatch Process](#) for more information.

The following business events queue transfer activities for dispatch:

- A loading activity is queued for dispatch at a predefined time prior to its earliest time window. This is controlled by a monitor algorithm configured on the activity's business object that is executed while the loading activity is being scheduled. The algorithm provided with the product considers a configurable cutoff time as well as an amount of time in advance of the activity time window to define when locking should be performed. Refer to the algorithm type *Advance Dispatch Transfer Activity* [MI-ADDSPTAC](#) for information.
- An unloading activity is automatically queued for dispatch when goods are loaded and the loading activity is completed.
- A dispatcher may manually mark an activity as ready for [dispatch](#). For example, the dispatcher may decide that the crew should perform the activity earlier than originally planned by the scheduler for some reason.

Queuing an activity for dispatch should not be confused with the actual process of dispatching the activity to track the crew's progress working the activity. The former locks the scheduled time whereas the latter tracks the actual execution when the crew is ready to perform the activity

Actual dispatching can only happen once the crew has started the transfer. A load activity cannot be dispatched until a driver arrives to the depot to start the transfer. When notified by the host system that the crew has arrived to the depot to start the transfer, the system immediately dispatches the activity.

An unload activity cannot be dispatched until the loading activity is completed. Once the loading is complete the system immediately dispatches the unload activity.

When ready to be dispatched, the system creates an assignment for the activity and its currently assigned crew, and sets the activity to the **Assignment In Progress** state. At this point the crew is in charge of [working the activity](#).

Along with the assignment for the loading activity the system also creates the following entities:

- A crew record is created for the driver, if one does not already exist.
- An transfer shift is created for the crew to track their [current location](#) on the way to unload the goods.

Both the loading and the unloading assignments are associated with the crew's transfer shift.

NOTE: Dispatching a transfer activity, load and unload, does not rely on any background process.

Assignment In Progress

When it is time for tracking the progress of the crew working a transfer activity, the system creates an assignment for the activity and sets the activity to the **Assignment In Progress** state. An assignment is a copy of the activity information created by the system to track the crew's progress in accomplishing the work.

Refer to [About Depot Transfer Assignments](#) for more information.

While the activity is in this state the ownership of the work is with the crew and the system waits for their resolution.

Possible resolutions may be as follows:

- If the crew performs the work and completes their assignment. the activity is automatically [completed](#). The activity is transitioned to the **Completed** state and the completion information is sent to the requesting host system.
- If the current crew was incorrectly associated with the loading activity, their assignment may be declined by the host system. This sets the activity back to being queued for dispatch and waits for the actual crew to arrive. Note that an unloading assignment cannot be declined since the goods were already loaded and the unload assignment has to be completed.

Completed

The transfer activity is set to the **Completed** state when:

- The crew performs the work and completes their activity assignment.
- The crew completes their activity assignment but marks it as incomplete because actual work could not be performed; for example, when details on the activity are invalid or there was some problem getting access to the work site. A completion message is sent to the host indicating that the system will no longer attempt to perform the work. This allows the host system to cancel the activity, or correct and resend it as needed.

When goods are on board and the loading activity is complete, the system automatically dispatches the unloading activity of the transfer job. An unload assignment is created stating the driver is en route to the destination depot.

If however there was a problem with loading the goods and the loading activity was completed with an incomplete reason, the unloading activity is automatically canceled.

Canceled

Canceling an activity is typically initiated by the host system for the purpose of preventing the work from being performed for some reason. The system can only confirm a cancellation when the activity is not yet dispatched. Once dispatched, the system attempts to cancel the activity but only the crew has the latest information as to its current state.

Refer to [Activity Cancellation By Host](#) for more information about how such a request is handled.

Expired activities may also be automatically canceled based on the activity type configuration.

On Hold

There are situations where the activity is not yet ready to be scheduled as it waits for some business event to occur or for proper timing, etc. An activity that is placed on hold is not scheduled; it is not sent to the scheduler and as such, would not be automatically assigned a resource to perform it.

The product does not implement specific business rules for placing a transfer activity on hold, but custom business rules may be introduced to do so in order to prevent the activity from being scheduled.

Transfer Activity Actions

A dispatcher or a business process may perform various actions on a transfer activity that may impact the activity's state and schedule.

The following section describes these actions and the business flows associated with them.

Lock To Bay

This action is only applicable if the crew has arrived at the depot to perform the transfer activity. It allows a user or a process to explicitly allocate the task to a specific bay.

Transfer Scheduling Manifest

The host system manages the assignment of depot transfer jobs to drivers, asking them to report to the loading depot at a specified scheduled time to start their job. Transfer jobs start with loading the goods at the distribution depot. At a predefined cutoff time, the scheduled time for these activities is therefore locked in advance to support this process.

A web service is provided for the host system to request a manifest of all upcoming loading activities for a depot that are locked and queued for dispatch. The host system should send such request on a daily basis some time after the cutoff time.

For example, assume the time window of a loading activity starts tomorrow at 9 AM. Also assume the system is configured to lock the scheduled time of transfer jobs if their time window starts within the next 24 hours as of a cutoff time of 10 AM. The next Task Monitor batch process that runs today, on or after 10 AM, would lock the scheduled time of the transfer activity since its time window starts within the next 24 hours. The host system may send in a transfer manifest request sometime after 10 AM, allowing enough time for the batch process to run and lock the schedule.

NOTE: Logic to lock the scheduled time of transfer loading activities resides in a monitor algorithm defined on the activity business object's [Being Scheduled](#) state.

Refer to [Transfer Job Workflow](#) for more information.

About Depot Transfer Assignments

An assignment is a copy of an activity created by the system to track the crew's progress in accomplishing the work. The system creates a loading assignment when a driver arrives at the depot to start a transfer job. The assignment is used to track

their progress loading the goods. Once loading is complete, the system creates an unloading assignment for the crew to track their travel and progress of unloading the goods at the destination depot.

Refer to [About Depot Transfer Activities](#) for more information.

The lifecycle of an assignment affects the lifecycle of its activity. This section describes concepts related to managing loading and unloading transfer assignments.

Loading Assignment Workflow

Loading assignment states are depicted in the following diagram:

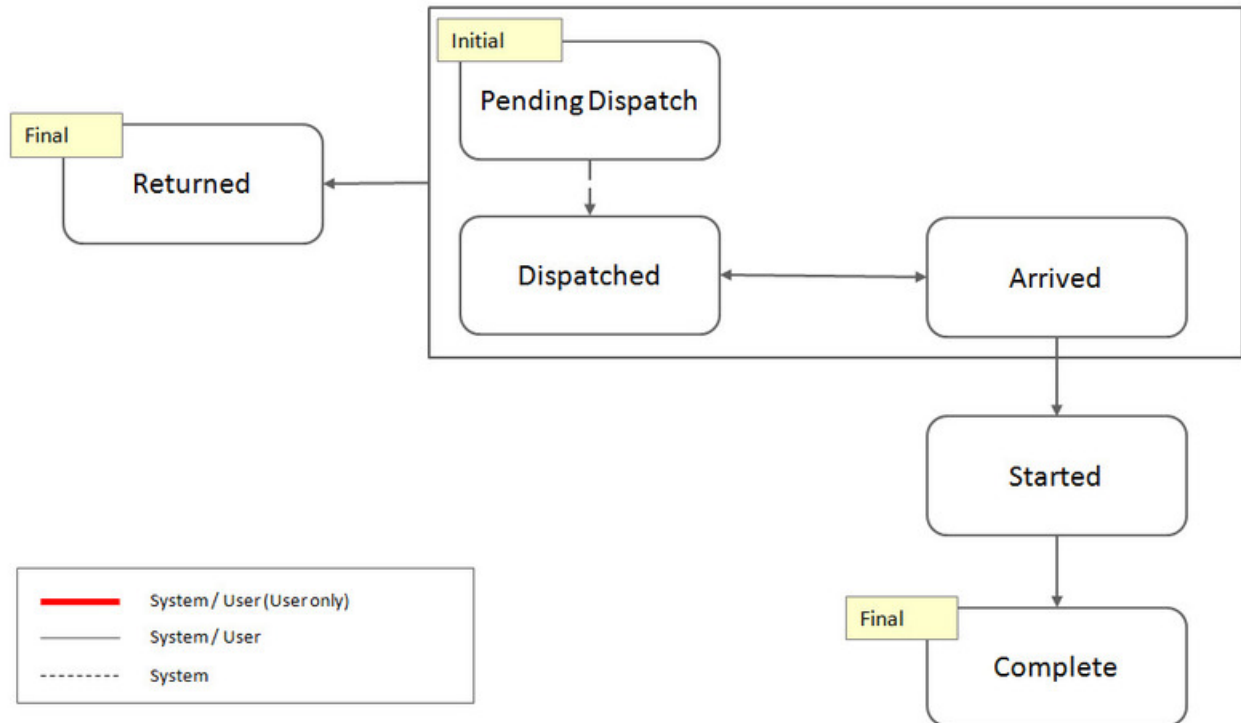


Figure 9: Loading Assignment States

The following sections describe the loading assignment overall lifecycle.

Pending Dispatch

An assignment is initially created in this state waiting to be acknowledged as [dispatched](#).

Dispatched

Typically, an assignment is considered dispatched if the assigned crew acknowledges its receipt. This is the last step in the process of dispatching an activity. However, since the product's mobile application does not support the workflows needed to dispatch and work depot transfer activities, transfer assignments are not sent to any mobile device and therefore are automatically marked as dispatched by default.

Arrived

This state indicates that the crew has arrived at the depot and waits to be told which bay to go to. The assignment is automatically set to this state upon creation.

The scheduler is notified of the crew's arrival and locks the scheduled bay for the activity. The host system polls for this information and once available, conveys it to the crew. If the bay is not available at this time they crew may be asked to wait shortly.

Started

This state indicates that the crew is on site and has started to load the goods.

The host system provides this [status update](#).

Completed

This state indicates that the crew has loaded the goods and left the depot. If for some reason goods could not be loaded, the loading assignment is completed with an incomplete reason code.

The host system provides this [status update](#).

Completing the assignment causes the activity to complete. Refer to the transfer activity [complete](#) state to better understand the impact on the entire transfer job.

Returned

An assignment in this state indicates that the crew is no longer responsible to perform the activity.

An activity assignment may be returned in the following situations:

- The transfer job has been associated with the incorrect driver by mistake.
- The system automatically recalled the activity due to a host request to cancel the activity. Refer to Activity Cancellation By Host for more information about how such request is handled.

Returning the assignment causes the activity to reevaluate its state. Refer to the transfer activity [Assignment In Progress](#) state for more information.

Unloading Assignment Workflow

Unloading assignment states are depicted in the following diagram:

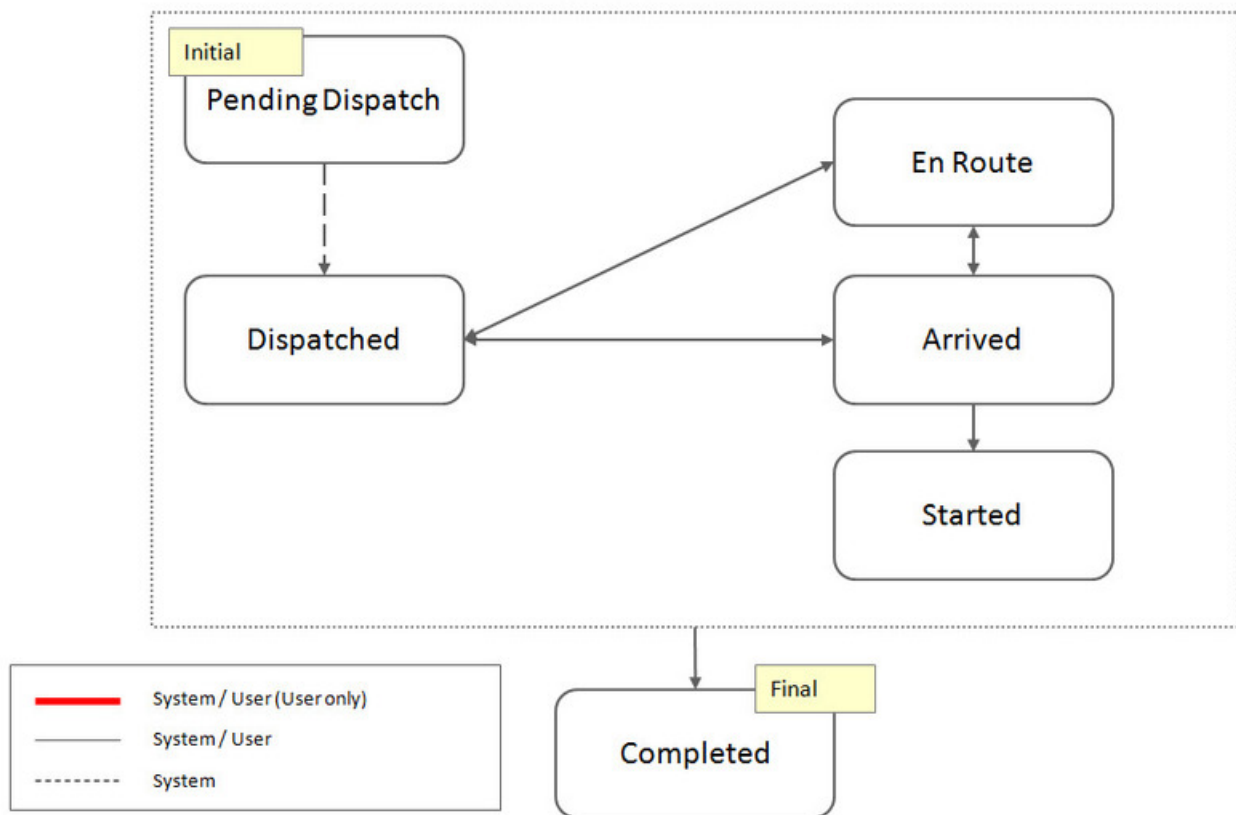


Figure 10: Unloading Assignment States

The following sections describe the unloading assignment overall lifecycle.

Pending Dispatch

An assignment is initially created in this state waiting to be acknowledged as [dispatched](#).

Dispatched

Typically, an assignment is considered dispatched if the assigned crew acknowledges its receipt. This is the last step in the process of dispatching an activity. However, since the product's mobile application does not support the workflows needed to dispatch and work depot transfer activities, transfer assignments are not sent to any mobile device and therefore are automatically marked as dispatched by default.

En route

This state indicates that the crew is currently on its way to the depot. The assignment is automatically set to this state upon creation.

Arrived

This state indicates that the crew has arrived at the destination depot and is ready to unload the goods. The host system provides this [status update](#).

The scheduler is notified of the crew's arrival and locks the scheduled bay for the activity. The host system polls for this information and once available, conveys it to the crew. If the bay is not available at this time they crew may be asked to wait shortly.

Started

This state indicates that the crew is on site and has started to unload the goods. The host system provides this [status update](#).

Completed

This state indicates that the crew has unloaded the goods and left the depot. The host system provides this [status update](#). Completing the assignment causes the activity to complete.

Maintaining Depot Tasks

This portal is used to view and maintain a depot task.

You can access a depot task from the "Scheduled Tasks" zone on the crew shift portal or a Gantt zone.

Refer to [About Depot Runs](#) for more information.

The following zones may appear as part of the portal's **Main** tab page

- **Depot Task.** This zone provides information about the depot task, including the various maintenance actions applicable to its current state.
- **Activities In Depot Run.** This zone list all the activities scheduled as part of the depot run.
- **Assignment List.** This zone lists load tracking assignments associated with activities scheduled to the depot run.

Maintaining Depot Time Windows

This portal is used to view and maintain depot time window overrides.

You can add a new override or access an existing override depot time window from the "Override Time Windows" zone on the Depot portal.

Refer to [About Depots](#) for more information.

The **Depot Time Window** zone on the portal's **Main** tab page provides information about the time window, including the various maintenance actions applicable to its current state.

Maintaining Depot Cutoffs

This portal is used to view and maintain a depot cutoff request.

You can add a new cutoff request or access an existing request from the "Cutoffs" zone on the Depot portal.

Refer to [About Depot Cutoff](#) for more information.

The following zones may appear as part of the portal's **Main** tab page

- **Cutoff.** This zone provides information about the depot cutoff request, including the various maintenance actions applicable to its current state.
- **Closed Depot Runs.** This zone lists the depot tasks closed by the cutoff event.
- **Activities Locked To Depot Run.** This zone lists delivery activities locked to the depot tasks due to the processing of the cutoff event.
- **Remaining Depot Tasks To Cutoff.** This zone lists depot tasks that still remain to be processed by this cutoff event.

Contractor Management

About Contractors

Contractor management functionality supports having outside contractors or a supplemental workforce completing work for your organization. If your organization uses contractors to work activities, you might use contractor management functionality to assign, schedule and complete this work.

Two models for working with contractors are supported: a shift based model and a capacity based model. You might use both models since each contractor company is set up and managed on an individual basis.

This section describes concepts related to contractor management functionality in general and those specific to each model.

Shift Based Model

In this model the contractor company allocates mobile resources to work for your company. Their mobile resources, i.e. mobile workers, vehicles, crews and their shifts, are set up, configured, and scheduled as if they are internal employees. Similarly, contractor crews work their scheduled shifts just like company crews; no additional management nor approval of work is needed.

Unlike internal company resources, contractor resources may only work as part of their own company's crews or your company's crews, but never as part of another contracting company's crew. In addition, a contractor mobile worker is restricted to access designated parts of the application and view only work committed to their contracting company.

The contractor company is defined as a [resource group](#) and all of their mobile workers, vehicles and crews must reference the company's resource group as part of their resource definition record.

NOTE: Aside from mobile workers, all other user roles such as dispatchers, resource planners, etc. are assumed to be company employees. Refer to [Contractor Users](#) for more information.

Capacity Model

Instead of providing resources to work as part of your mobile workforce, in this model contractor companies agree to provide a certain capacity of work to your company. These companies are then assigned work by your company based on such agreed upon capacity.

Contracted capacities are expressed in terms of number of hours or activities of a certain type, in certain areas, for a period of time. For example, a contractor company may offer to perform 1000 activities of a certain type, in one or more areas, on a specific day or week, or any period of time.

The contractor companies receive work from your company and they are responsible for managing these activities, scheduling them to their crews, and forwarding back to your company the corresponding completion information when work is complete.

The contractor company is defined as a [resource group](#) in your system.

NOTE: While the capacity contractor company is defined within the system, their individual resources are not. Their crews are fully managed and scheduled outside of your system.

This section describes concepts related to capacity based contracting model.

Scheduling Work To Capacities

Your company establishes capacity records representing work agreements with external contractor companies. Each capacity record specifies a period of time during which the contractor company agrees to take on work of a certain type in specific areas.

Refer to [About Capacities](#) for more information.

The scheduler considers capacities as well as internal crew shifts when scheduling a particular activity based on the whether or not the activity is allowed to be worked by a contractor and other preference options. During this phase, the schedule is constantly optimized and work is moved around between crew shifts and capacities, between your company and contractor companies.

It is important to understand the difference between scheduling work to a shifts vs. capacities. When the scheduler assigns work to a specific crew on a specific shift, it has all the information needed to choose an optimal assignment in terms of the capabilities the crew possess on that shift and it can incorporate travel time from the crew's current location to the various locations on their route. On the other hand, capacity defines a "bucket" of work the contractor company may receive from your company which may be completed by an unknown number of crews anytime over the agreed upon period. As such, scheduling to a capacity cannot account for travel time nor can it consider crew capabilities.

Refer to [Activity Contracting Restrictions](#) and [Contractor Eligibility Rules](#) for more information.

Some time prior to each capacity's start time, at the notification time defined for the contractor company, the system closes the capacity and dispatches ("locks") all the activities booked to it so the scheduler does not make any further changes to their schedule. Dispatching ("locking") the activities to the capacity commits them to the contractor company. Once committed, the activity is communicated to that company.

Activities may also be manually committed to the contractor company earlier than the agreed upon notification time by allocating or dispatching the activity to a capacity for that company.

Refer to [Advance Notification](#) for more information.

Advance Notification

Contractor companies may need to be notified in advance of work your company commits to them so they can plan accordingly. Each contractor company may specify the amount of lead time prior to the beginning of the contractor's capacity period that the contractor must be notified of activities committed to them for the specific capacity. This period of time is called the notification horizon.

At the beginning of the notification horizon relative to the capacity's start time, the system **closes** and dispatches ("locks") the capacity's schedule; no additional activities can be added to the capacity nor removed from it.

For example, if the capacity period begins on Thursday at 8:00 AM, and the notification horizon for this contractor is set to 72 hours, the capacity is closed on Monday at 8:00 AM and the contractor is notified of all committed work booked to the capacity.

Notification horizon is specified on the contractor's record but may be overridden for an individual capacity record as needed.

NOTE: This functionality relies on the shift monitor batch process to be run at an appropriate frequency.

Refer to [Outsourcing An Activity To The Contractor](#) for more information about communicating an activity to the contractor company.

Outsourcing An Activity To The Contractor

Activities committed to a capacity based contractor company need to be communicated to that company for approval. Refer to [Committing An Activity To A Contractor Company](#) to better understand the business meaning of this commitment.

The contractor company has to approve taking on each activity committed to them or decline the work, within a specified approval amount of time also specified on the contractor's record.

If the contractor has not responded in a timely manner, the activity is considered declined by the contractor. Once declined, the activity is no longer committed to the contractor company and they are excluded from performing the specific activity.

If the contractor company has approved taking on the activity, they are responsible for performing the work. When work is complete, the contractor company must provide the necessary completion information needed to complete the activity in the system. If for some reason the contractor company cannot perform the work they agreed to take on previously, they may decline it.

Refer to [Managing Committed Work](#) to learn more about the options available for a contractor company to manage the work committed to them by your company.

Activities can also be reverted manually as necessary by a dispatcher. The company can always pull back activities and allocate them elsewhere (with the risk that the contractor has already done the work).

The system creates an assignment copy of the activity when it commits it to an external contractor company to track their approval and progress working that activity. The assignment is established at commitment time and is finalized when the activity is no longer committed to that the contractor company. This is similar to the notion of a crew's assignment to track the crew's progress working an activity.

Refer to [Contractor Assignment](#) for more information.

Managing Committed Work

Interacting with a contractor company with respect to an activity committed to them involves these steps:

- The system sends the activity information to the contractor company for approval. The system may resend the activity information to the contractor company again if key information has changed.
- The contractor company approves or declines taking on the activity.
- The contractor company sends the necessary completion information when they complete the activity.

Large contractor companies may use their own system to manage their work. The product supports a standard web service based integration between their system and yours. When integrated, all communications with the contractor company is handled by message exchanges between the systems.

If the capacity contractor company's system is integrated with your company's system, their external system definition has to be referenced on their resource group record.

Refer to [Contractor Management Integration](#) for more information.

Smaller contracting companies that do not use a system of their own may manage their work from within your company's application as follows:

- Their managers may access the [Contractor Work Management](#) portal to approve, reject, and fully manage work your company has committed to their company. When using this portal, a contractor can only see their work; they cannot see other contractor's information nor your company's information.
- Their workers may use the [contractor's mobile application](#) to perform work assigned to them by their managers.

Refer to [User Access](#) for more information about allowing contractor users to access your application.

User Access

Any contractor manager who needs to access the Contractor Work Management portal must be set up as a user within the system. Refer to [Contractor Users](#) for more information.

If contractor workers utilize the contractor mobile application to perform work given to them by their manager, they must access the application using a valid user ID. Your company may set up each of these workers with their own user ID or allow all workers from the same contractor company to share a user ID.

To assist contractors managers in distributing work amongst their workers, the Contractor Work Management portal allows them to tag work to each of their workers. A worker's tag may be their own user ID or some free form text unique identifier they may assign to them. If your company has set up each contractor worker as a user, then their tag is their unique user ID.

Otherwise, if your company has chosen to provide a shared user ID to all workers from this contractor company, the tag must be a free text identifier set by the contractor manager. Free text identifiers must be unique across your company. As such, it is a good practice to have your contractor companies prefix their identifiers with the company name. For example, ABC company's tags might be "ABC John", "ABC Sally", and "ABC Paul" as opposed to just "John", "Sally", and "Paul".

NOTE: The shared user ID option is not applicable when using the contractor mobile application in "connected MCP" mode.

The option chosen by your company with respect to individual or shared use of user IDs is specified as the tag type on the contractor company's resource group record.

NOTE: Aside from contractor managers and workers, all other user roles such as dispatchers, resource planners, etc are assumed to be company employees.

Activity Contracting Restrictions

Not all work may be contracted. Each activity indicates whether it should be worked by a company crew only, any internal crew (company or shift based contractor), or any resource. If not specified, the activity is not eligible for contracting. This is defaulted from the activity type, but can be overwritten on the activity itself.

NOTE: Emergency activities are never eligible for capacity contractors.

The host may request that an activity can be assigned to specific contractors (either capacity or shift-based) by supplying a list of [Allowed Contractors](#) on the activity.

A dispatcher may intervene and set activities that are otherwise eligible for contracting as not eligible for contracting using the [Revert To Company](#) action. They can also [exclude the contractor](#) currently assigned to work an activity from performing it.

Contractor Eligibility Rules

Contractor eligibility rules define scheduling prioritization between your company and contractor companies. These rules determine how your organization wants to schedule work for a particular service area and service class.

Defining contractor eligibility rules is optional. If not specified, the scheduler automatically gives preference to internal crews, company and shift based contractors, over external capacity based contractors.

Eligibility rules specify an overall preference between company and shift-based contractors vs. external capacity contractors. The eligibility rule can also provide a ranked list of resource groups, company and contractors, to be assigned work in the specified order of preference. If multiple groups are assigned the same rank, they are equally preferred. If a resource group is not listed, it is assumed to be least preferred.

The ranked preference list is stamped onto activities that are eligible to be worked by contractors when the activity is created or rescheduled, and forwarded to the scheduler for consideration. If the system detects a change to the eligibility rules while the activity is still being scheduled, it re-stamps the preferences to reflect the change.

Eligibility preferences are used by the scheduler when scheduling the activity. Scheduling cost is associated with not honoring this preference option. For example, if your preference is capacity over internal resources, then it would be less expensive to use a capacity over an internal resource.

Activities can further limit which contractors are eligible to work the activity by specifying one or more [Allowed Contractors](#) on the activity. The scheduler may only assign the activity to a company in that list.

Committing An Activity To A Contractor Company

Regardless of the contracting model used, committing an activity to a contractor company means it becomes their responsibility. Once committed, contractor users, irrespective of their contracting model, can view and act upon activities committed to their company and their company only.

NOTE: Until an activity is committed to a contractor company, only internal company users may access it!

[Allocating](#) or [dispatching](#) an activity to a contractor crew's shift or contractor company's capacity commits it to that company. Commitment stamps the activity with the resource group identifier of the contractor company and is key to ensuring contractors access their committed work only.

An activity remains committed to a shift based contractor company as long as it remains allocated or scheduled to one or the contractor company's crew shifts. If a contractor crew could not perform an activity and had to return it, or if the activity was recalled from them for any reason, the activity is no longer committed to that contractor company and can be rescheduled to any resource.

An activity remains committed to a capacity based contractor company until they decline it or a dispatcher recalls the activity from them for any reason. If the contractor company declines the activity, they are also explicitly excluded from being scheduled that activity again.

Refer to [Managing Committed Work](#) to learn more about the options available for a capacity based contractor company to manage the work committed to them by your company.

Resource Groups

Resource groups define information about the contractor companies doing work for your organization. This information includes name, contact information, and the base service area where they work.

Your organization must also be defined as a resource group of class "company" if you need to establish scheduling preference ranks between contractor companies and your own company using [Contractor Eligibility Rules](#).

The contracting model as well as whether this is your own company is defined on the [resource group type](#).

Regardless of the contracting model used, any individual who needs to access the system as a contractor needs have a user record set up to reference the contractor company they work. Once users are set up as contractors you may use this portal to associate them with the company they work with.

Refer to [Configuring Contractor Management](#) for more information.

CDI and Contractors

Contractor crews are highlighted on the Scheduling Gantt, and the activity search supports filtering by contracting related criteria.

The Contractor Capacity tab allows a dispatcher to manage contractor capacities they monitor.

About Capacities

A capacity defines a period of time during which a [contractor company](#) has agreed to perform a certain amount of work of a certain type, in certain areas, assigned to them by your company.

For example, your company may have the following work capacity agreements with contractor companies:

- Allow contractor company A to perform up to 500 hours worth of inspection work in service areas "Center" and "North" for a specific period of time.
- Allow contractor company B to perform up to 1000 activities of a specific type in service areas "Center" and "South" for a specific period of time.

Refer to [Capacity Model](#) to better understand how work capacities are used.

This section describes concepts related to managing contractor work capacities.

Capacity Type Controls Everything

Each capacity record references a [capacity type](#) which defines attributes common to all capacities of that type, for example whether the amount of work is expressed as a number of activities or number of hours.

Creating Capacities

Capacities are assumed to be set up by your company based on agreements made with the contractor company. Similar to the notion of crew shifts, capacity records may be generated based on weekly templates or created manually as needed to reflect agreements between your company and the contractor company. Capacities for a specific contractor company may be generated by a single weekly template at any given point in time.

Refer to [About Capacity Weekly Templates](#) for more information.

Capacity Workflow

The following describes key stages in the overall lifecycle of a capacity.

Pending

A capacity is initially created in this state. While in that state it is not sent to the scheduler and is considered a draft until marked as fully planned.

Planned

A capacity in this state is no longer a draft and is available for scheduling as per its configuration.

If generated by a template, the capacity is automatically set to the planned state. A user creating a capacity online has to explicitly mark it as planned when the set up is complete and the capacity should be sent to the scheduler.

While in this state, the scheduler may assign work to the capacity and change its decisions as needed to optimize and improve on the overall schedule. When it is time to [notify](#) the contractor company of the final list of activities committed to them, the system closes the capacity.

Refer to [Closing The Capacity](#) for more information.

Expired

When the end of the capacity's period is reached, it becomes expired and reaches a final state. This has no impact on the work scheduled to the capacity nor the responsibility of the contractor to that work.

Canceled

Until it expires, a capacity can be canceled and removed from the scheduler. A user or any process can cancel a capacity as long as no work is committed to it. Once started, it can only be completed.

Activities currently scheduled to the canceled capacity will be rescheduled.

Controlling The Schedule

This section describes the various aspects of a capacity that are used by the scheduler to control the activities assigned to that capacity.

Capacity Period

A capacity is offered for a specific period of time. Only activities that can be performed anytime during that time frame may be scheduled to the capacity. Refer to [activity time windows](#) for more information.

The time zone on a capacity indicates the time zone in which the period of time is specified.

Covered Areas

Contractor companies may be associated with a list of service areas where they are eligible to work. Capacities defined for that company may only cover these areas.

Similar to a crew's shift, a capacity must cover at least one service area and the scheduler would only assign work to capacity within the areas it covers.

Covered Work

A capacity defines the type of work it covers by either specifying an activity type or a service class, which is a group of activity types. If the amount of work is expressed as a number of activities, you may only specify an activity type.

The scheduler would only assign work that belongs to these activity types up to the number of activities or hours specified by that capacity.

Depot Restrictions

If depot related activities are covered by the capacity, you may limit the work to specific depots only. If no depots are specified, work from any depot may be considered.

NOTE: Assigning depot work to a contractor capacity assumes the contractor manages the travel to and from the depot to handle the goods as part of their scheduling process. In other words, the system does not schedule depot runs to for the contractor, nor leverages any transfer of goods functionality at all.

Capacity Actions

A dispatcher or a business process may perform various actions on a contractor's capacity that may impact the capacity's state and schedule.

The following section describes these actions and the business flows associated with them.

Close

Closes the capacity and [commits](#) all scheduled activities to the contractor company. The scheduler is prevented from removing any of the allocated activities or adding more activities to the capacity.

Capacities can be closed in the following ways:

- When it is time to [notify](#) the contractor company of the work committed to them, the system closes the capacity and commits all the work in to the contractor company.
- A dispatcher can close a capacity to finalize its schedule earlier than planned.

Open

Reopens a closed capacity's schedule to allow the scheduler to assign more activities to it or remove some activities from it, as needed. A capacity may only be reopened before the notification window is reached for the contractor company.

Commit All

This action [commits](#) to the contractor all of the activities scheduled to the capacity that have not yet been committed to them. This action does not close the capacity.

About Capacity Weekly Templates

While capacities can always be manually created to address a specific situation, they are typically planned and generated ahead of time based on recurring templates. Templates provide a tool for quickly generating large numbers of capacities based on pre-defined parameters.

A weekly template defines a cyclical weekly pattern made up of weekday capacity templates for a specific contractor company during a specific period of time.

For example, you might define a weekly template where the contractor company agrees to take on a specific type of work in a specific area on Mondays, Tuesdays and Wednesdays from 9 AM to 5 PM in one week, and cover another type of work in another area on Thursdays and Fridays from 9 AM to 5 PM the following week. Thus, the weekly template would rotate every 2 weeks.

Each contracting company may be associated with multiple effective dated capacity weekly templates representing their agreements with your company over time, but only one template is effective at any point in time.

This section describes concepts related to setting up and maintaining capacity weekly templates.

A Weekly Template Is Made Of Capacity Templates

A capacity weekly template is made of one or more capacity templates that describe the recurring capacity patterns. The details for each capacity within a cycle are captured on a capacity template record. For example, a weekly template might use a different capacity template on Fridays than the one used on other days.

Capacity weekly templates are dedicated to a contractor company. As such, each weekday capacity template is fully defined with everything that the capacity generator needs to generate a capacity for this contractor company.

The initial date of the cyclical weekly pattern is the effective date of the template.

Actual vs. Template Capacities

A capacity template defines the details needed to generate actual capacity records based on that template. Any detail you can define for an actual capacity you can also predefine on a capacity template.

To enforce the close pairing between a capacity template and the capacities it creates, both the actual and the template records are stored in the same table so they can share similar details. The same capacity type defines both the business object used to create actual capacities of this type, as well as the business object used to create capacity templates of that type.

Refer to [Defining Capacity Types](#) for more information.

A generated capacity maintains a reference to the capacity template that was used to generate it so that changes made to a template may be reapplied to those capacities.

Refer to [Capacity Generation](#) for more information.

Defining and Applying A Template

When you create a new capacity weekly template, its status is **Pending**. No capacities are generated while the template is in this state allowing you some time to properly define the weekday capacity templates and review them as needed.

After you have defined your capacity templates, you must apply the weekly template so that actual capacities will be generated from it, based on the template's effective period. Applying the template changes its status to **Applied** and the system right away generates future capacities starting from the template's effective date (or from now if that date is in the past) and no later than the template's expiration date (if any).

NOTE: Generating capacities online is limited to a preconfigured number of days in order to avoid the apply action from taking too long and timing out. The parameter is defined on the capacity generation algorithm.

Once applied, the capacity generation background process manages the generation of capacities going forward.

Refer to [Capacity Generation](#) for more information about the generation algorithm as well as the overall process.

Making Changes To An Applied Template

Extending or shortening the effective period of an applied template adjusts future capacities accordingly right away when you make the change.

If you attempt to add, delete, or modify a capacity template for an applied weekly template, the system automatically creates a draft copy of the entire weekly template it belongs to; your changes are made to this draft copy, leaving the original template unchanged. This allows you to make changes in a more controlled way.

You make changes to the draft copy and preview the capacities that would be generated should you accept the draft. If you decide to accept the draft, the original template is updated to match the draft and changes are immediately applied to already generated future capacities accordingly.

Refer to [Draft and Original Templates](#) for more information.

One Effective Weekly Template For A Contractor

A contractor company's capacity weekly templates are either effective for a specified number of weeks starting from their effective date, or they may be perpetual, which means they have an effective date but no end date.

The contractor company can have been associated with multiple weekly templates over time but their capacities may be generated by a single weekly template at any given point in time. In other words, a contractor company's weekly templates must never overlap in time.

Capacity Generation

The system generates future capacities for contractor companies based on the weekly templates your company has set up to reflect agreements made with them. Capacities are generated based on the cyclical pattern of the weekly template starting from the weekly template's effective date (or from now if the effective date is in the past) and up to a predefined number of days known as the planning horizon, or earlier, if the template expires before then.

The planning horizon is defined as a [global configuration](#) parameter and is used by the capacity generation process to control the maximum number of days worth of capacities it generates and has available for planning purposes, at any point in time.

If the template references a work calendar, the process skips generating capacities that even partially overlap with a public holiday.

The first few capacities may be initially generated online when you apply a weekly template, but the ongoing process of generating additional capacities to meet the planning horizon is handled by the weekly template monitor batch process.

NOTE: The capacity generation process relies on the capacity weekly template being periodically monitored by the **Shift Weekly Template Monitor** batch process as often as appropriate to meet your organization's needs. Refer to [M1-SHWKT](#) batch control for more information.

On a daily basis, the background process is designed to add new capacities as needed. However, if any capacities template was changed the process would automatically regenerate future capacities that were generated by that template to match its latest version. By default, even capacities that were manually overridden are regenerated to match the latest template version, but you may explicitly indicate to retain these changes when making a capacity template change.

Refer to [Overridden Capacity Details](#) for more information.

Unlike crew shifts that represent working hours for individual resources, capacities are allowed to overlap in time because they represent working hours for multiple contractor resources.

NOTE: Logic used to generate capacities from a capacity template resides in a "Generate Shifts" algorithm on the capacity template's business object. Click [here](#) to see the algorithm types available for this system event.

Overridden Capacity Details

In a situation where capacity details should differ from their template on a particular date, generate the capacity initially from the template and then you can manually adjust it to override the template settings. This includes canceling a capacity that is no longer needed. These modified capacities are called overridden capacities, as they no longer match the template from which they were generated.

By default, any changes made to a capacity template are automatically reapplied to all future planned, but not yet closed, capacities that were generated from that template, and any manual overrides will be lost. If you wish to retain overridden capacities, you may indicate so on the capacity template being changed.

Draft and Original Templates

If you attempt to add, delete, or modify a capacity template for a capacity weekly template that has been applied, the system automatically creates a draft copy of the entire weekly template it belongs to. This allows you to make your changes to the draft and preview the capacities that would be generated if you accepted the draft.

You can also see if any manual edits were made to previously generated capacities. By default, such overrides are lost when the capacity template that generated them is changed when a draft is accepted. You may explicitly request to retained overrides when making the capacity template change.

If you have reviewed this information and are sure you want to make the proposed changes, you may accept the draft. Otherwise you may delete the draft copy to discard the proposed changes.

When a draft template is accepted, the original template is updated to match the draft:

- If a capacity template was changed, future capacities generated by this template are regenerated to match the changed template.
- If a capacity template was deleted by the draft, the corresponding original template is also deleted if not yet used to generate capacities; otherwise it is marked as inactive. Future capacities generated by this older template are deleted.
- New capacities are generated for any capacity templates that were added by the draft.

By default, changes made to a capacity template impacts all future planned, but not yet closed, capacities generated by it, including overridden templates.

NOTE: Changes are reapplied to future capacities right away when you accept the draft.

Refer to [Overridden Capacity Details](#) for more information.

Once a draft has been accepted, any further changes made to the original template will result in a new draft being created.

Maintaining Resource Groups

This portal defines a contractor or your company resource group.

Refer to [Resource Groups](#) for more information.

You can access the portal from the **Main Menu > Resource Management > Resource Group**

The following zones may appear as part of the portal's **Main** tab page

- **Resource Group.** This zone displays information about the resource group.
- **Capacity Weekly Templates.** This zone lists capacity weekly templates for the contractor. It is only visible if the resource group record is for a capacity based contractor.
- **Users.** This zone lists all contractor users that work for this contractor. It is only visible if the resource group record is for a contractor.
- **Unassigned Contractor Users.** This zone lists all contractor users that are not yet associated with any contractor company. You may associate selected users with the selected contractor company. It is only visible if the resource group record is for a contractor.

The **Calendar** zone on the portal's **Calendar** tab page provides a calendar view of the contractor's work capacities for the current month. It is only visible if the resource group record is for a capacity based contractor.

Refer to the Calendar (F1-CALENDAR) zone type for more information.

Maintaining Capacities

This portal is used to display and maintain a contractor work capacity.

Refer to [About Capacities](#) for more information.

You can access the portal from the **Main Menu > Resource Management > Capacity**. You are brought to a query portal with options for searching for a specific capacity. Once a capacity has been selected, you are brought to the maintenance portal to view and maintain the selected record.

The following zones may appear as part of the portal's **Main** tab page:

- **Capacity.** This zone provides information about the capacity, including the various maintenance actions applicable to its current state.
- **Assigned Work.** This zone lists activities assigned to the capacity.

Maintaining Capacity Weekly Templates

This portal is used to display and maintain a capacity weekly template for a contractor company.

Refer to [About Capacity Weekly Templates](#) for more information.

You can access the portal from the **Main Menu > Resource Management > Capacity Weekly Template**. You are brought to a query portal with options for searching for a specific capacity weekly template. Once a capacity weekly template has been selected, you are brought to the maintenance portal to view and maintain the selected record.

The following zones may appear as part of the portal's **Main** tab page:

- **Capacity Weekly Template.** This zone provides information about the capacity weekly template, including the various maintenance actions applicable to its current state.
- **Capacity Templates.** This zone defines the individual shift templates that comprise the rotation cycle for this weekly template.
- **Inactive Capacity Templates.** This zone lists capacity templates that are no longer active and is visible if such capacity templates exist for the weekly template.
- **Overridden Capacities.** This zone lists capacities that were generated by this template but were since modified manually and no longer matches the template from which they were generated. The zone is visible if such capacities exist for the weekly template.

The following zones may appear as part of the portal's **Review** tab page:

- **Preview Calendar.** This zone allows you to review a month's worth of potential capacities that would be generated if you applied the recent changes to the template. You may broadcast a specific capacity to displays its details in the capacity preview zone.
- **Capacity Preview.** This zone displays a capacity broadcasted from the preview calendar zone.

Refer to the Calendar (F1-CALENDAR) zone type for more information.

Maintaining Capacity Templates

This portal is used to view and maintain a capacity template as part of setting up a capacity weekly template.

Refer to [About Capacity Weekly Templates](#) for more information.

You can access the portal from the capacity templates zone on the [Maintaining Capacity Weekly Templates](#) portal.

The following zones may appear as part of the portal's **Main** tab page:

- **Capacity Template.** This zone provides information about the capacity, including the various maintenance actions applicable to its current state.
- **Overridden Capacities.** This zone shows capacities that were generated by this template but their setup has been manually changed from the setup defined by the template.

Maintaining Contractor Eligibility

This portal is used to view and maintain scheduling prioritization between your company and contractors.

You can also access this portal from the **Main Menu > Resource Management > Contractor Eligibility**.

Refer to [Contractor Eligibility Rules](#) for more information.

The **Contractor Eligibility** zone on the portal's **Main** tab page provides information about the contractor eligibility record.

Contractor Work Management

This portal helps contractor users to manage work committed to capacities offered by their company.

As a contractor manager, this portal should typically be set up as your home page, but you can also access the portal from the **Main Menu > Service Management > Contractor Work Management**.

NOTE: If you are a company user, you need to navigate to this portal for a specific contractor company in context. Navigate to this portal from the specific contractor's resource group record using the context menu associated with the contractor identifier in the record actions section.

The following zones appear as part of the portal's **Main** tab page:

- **Capacity Search.** This zone allows you to search for capacities set up for the contractor company.
- **Committed Work.** As a contractor manager, use this zone to manage activities assigned to the selected capacity and committed to your company. Refer to [Managing Your Work](#) for more information.

Managing Your Work

NOTE: Committed work must be reviewed and either approved or declined in a timely manner. Failing to do so in a timely manner would result in the system automatically removing that work from your company and handling it as if it has been declined.

To assist you with managing your company's approved work, you may use the **Maintain Workers** link on the main zone header to set up unique tag names for your workers. You only need to do this once or adjust as needed. Once set up, use these tags to assign approved activities to specific workers. This step is necessary if your company utilizes the contractor mobile application to complete work you have approved.

If you are not using the contractor mobile application to complete the work, you may navigate to each activity to complete your company's assignment providing the necessary completion information.

Mobile Application

About The Crew's Application

Crews may use the mobile application to perform their work. It allows them to fully manage their work schedule, get more work as needed, handle exceptions, as well as seamlessly report their whereabouts.

This section provides a high level overview of the mobile application as well as describes functionality supported by it that you may want to take advantage of.

NOTE: This is not a user guide to the mobile application itself.

Overview

Before a crew starts their shift, they may preview their first few tasks for the day on their mobile device. Refer to [Preview](#) for more information.

When the crew starts their shift, one or more of their scheduled tasks are dispatched to their device as per their shift configuration and appear on the mobile application's task list. Refer to [Auto-Dispatching](#) for more information.

Once tasks have been dispatched, the crew can work disconnected if wireless communication is not available. As they work each task, they enter status updates and completion details on the mobile device. The mobile application queues up these changes and once the device is connected, sends them to the server to update the corresponding master copy. These status updates are reflected on the dispatcher's Gantt, as well as used by the scheduler to adjust the crew's remaining schedule accordingly.

New tasks and updates to existing tasks may be received throughout their shift.

When the crew ends their shift, work assignments they did not manage to complete are returned to be rescheduled. Refer to [The End Of Shift Process](#) for more information.

Preview

The mobile application allows the crew to preview tasks on their schedule before starting their shift. Previewed tasks may be viewed, but no modifications are allowed.

The Preview functionality is available for shifts that are explicitly configured to allow one or more tasks to be advance dispatched. When enabled, activities that are manually allocated to the shift (but are not queued for dispatch) are also included in the preview list. If the shift is configured to close and advance dispatch the entire schedule then all tasks on that shift may be previewed. Refer to [Advance Dispatching](#) for more information.

The Mobile Device Is Known At Start Time

The device the crew uses to work their schedule is only known once they use it to start their shift.

Once known, the device identifier is stamped on the started shift and the system directs all communications for the crew to that device. If for some reason the crew has to switch to another device, the new device is stamped on the shift and the system would discard any messages that may still arrive from the crew's previous device.

Refer to [About Remote Messages](#) for more information about the message exchange mechanism.

The system provides security measures to ensure that only one device at a time can have the shift data. Only the device stamped on the shift may have access to the shift data. The system allows any member of the crew to work on the shift as long as it is made from this device.

If a mobile device breaks or is lost during a shift, the dispatcher can forcibly disconnect the device from the shift to allow the crew to work the shift using another device.

Refer to [Switching Devices Mid Shift](#) for more information.

Real-time Updates

The main benefit of using the mobile application is that it provides real time visibility to the crew's progress in working their schedule. Each and every significant state change is reported back to the server and gets reflected in the dispatcher's CDI portal, as well as forwarded to the scheduler for consideration.

The dispatcher and the scheduler may react to exceptions better and more efficiently when they are aware of them as soon as possible. For example, while on site working on an activity, the crew may use the **Send Now** action to provide an update as to the remaining time they think is needed to complete the work if the actual work is taking significantly more or less time than planned. This information allows the scheduler or the dispatcher to adjust the crew's remaining schedule accordingly. Note that the Send Now action forwards any interim completion information to the server. Based on your custom business rules, this information or a subset of it, may also be forwarded to the host system as needed.

Another benefit is the ability to notify the crew of changes made to the work they were asked to perform. This could be changes to activity details as well as schedule changes.

Note that if a message about a change in the details of a specific activity reaches the crew too late, i.e. they have already completed the task then the update is ignored on the device. If the crew is working on that activity at that moment, they get an alert that the information about the activity they are working on has changed and they can either reload the assignment and lose their changes or keep on with what they have.

Declining Work

The crew may use their mobile application to explicitly decline work dispatched to them for various reasons. For example, they may indicate that they don't have the proper skills or equipment needed to perform the work. In this case the system creates a To Do entry for a resource manager to follow up and ensure the capability information recorded for the crew's mobile workers and vehicles is accurate.

You may configure the system to add more reasons that best support your business practice and exception handling

NOTE: The decline action is hidden if the mobile worker's user does not have the proper application security permission to return an assignment.

Sending Current Location

The mobile application may be configured to periodically obtain Global Positioning System (GPS) data from the mobile device and forward this information to the server. This information is used to identify the most recent location of the crew utilizing this device.

Knowing the precise whereabouts of crews has the following main benefits:

- Provides dispatchers with a better picture on the CDI map of where the crews are relative to the crew's route, activities, other crews, and more.
- Allows the scheduler to more accurately adjust the crew's remaining schedule based on where the crew is as opposed to where they should have been. The scheduler estimates the crew's current location based on its last known location. The scheduler's estimation is more accurate when GPS reports are available.
- If the crew enters hazardous areas, their current location is used to raise alerts to dispatchers
- In case of an inquiry, allows the dispatcher to view the actual route a crew took on a specific shift compared with their planned route
- And much more.

Sending of GPS data from the mobile device can be enabled or disabled at the shift level.

Geographic Map

If the device is configured to support mapping, a geographic map icon appears on the task list user interface as well as on task related maps. When clicked, the icon displays a geographic map showing the crew's location as well as a route to their next tasks.

You can configure custom theme layers to be displayed on the maps. For example, you might configure layers to show the location of company assets. Different types of crews can be associated with different applicable layers.

NOTE: Logic to determine the custom themes to include on the mobile geographic map resides in a "May Layers" algorithm on the crew shift's business object. Click [here](#) to see the algorithm types available for this system event.

Enforce Procedures

Crews performing a certain primary function may be required to complete a set of procedures before they are allowed to start their shift, or when they change their primary function mid shift. These procedures may be defined on the service class used for primary function purposes.

Crews utilizing certain types of vehicles may be required to complete a set of procedures before they are allowed to use the vehicle. These procedures may be defined on such vehicle types.

Crews performing certain types of work may be required to complete a set of procedures on site before they are allowed to start the work. These procedures may be defined on such activity types.

Procedure enforcement is embedded as part of the following process flows:

- Request to start a shift
- Request to change a started shift's primary function
- Request to change a started shift's vehicles
- Every time the crew starts an activity of a specific type. If they suspended and start it later again, they would be requested to complete the procedures again.

A failed procedure may or may not be allowed to be manually overridden as passed based on your procedure type configuration. Once all procedures are passed or overridden the system proceeds to perform the requested change.

When applicable a procedure form may require the crew to take attendance to record who has participated in completing the procedure.

All procedures completed by the crew are sent to the server for audit purposes.

Refer to [About Procedures](#) for more information.

Handling Attachments

In the field, mobile users can upload and download attachments as needed to support their work.

Assignments with attachments, either common or specific to this activity, show a paperclip icon next to their entry in the mobile device **Task List** to indicate that an attachment is present. The crew can download their content on demand. All attachments can be downloaded on all devices; however, if the file type is not supported on the device, it cannot be opened or viewed.

The crew may capture pictures in the field while working an activity. Each capture is uploaded to the server and linked to the activity as an attachment. This attachment is also listed on the device as part of the assignment attachments list.

Refer to [About Activity Attachments](#) for more information.

Creating Activities in the Field

Depending on settings on the activity type, crews have the ability to create new activities in the field which may or may not be related to any existing activity on their schedule. This is useful when the crew is out in the field and come across an issue. Refer to [Related Activities](#) for more information.

The crew can choose to work the activity or send it in to the server to have it assigned to any crew. They may also indicate if the work requires an appointment to be coordinated with the customer. Refer to [Appointments](#) for more information.

Switch To Offline Mode

Crews that operate in areas with poor or no connectivity coverage need to work in offline mode. In this mode, their entire work schedule for the day is dispatched to them at shift start time as they won't be able to get additional work later. In addition, their shift schedule is closed so that no further work is assigned to them.

The crew requests to switch to offline mode and waits for all their work to be dispatched to them. The system notifies the crew when the process is complete so they can start traveling. The crew can also request to be online again, as needed.

Refer to [Offline Mode](#) for more information.

Request For Help

Crews may use the mobile application to request for immediate help.

They may click on the "Panic" icon to request immediate assistance for any reason. This action raises an alert for a dispatcher to follow up with the crew and send help as needed. These alerts are immediately assigned to an available dispatcher and should be given attention right away. If the mobile device has the capability, the alert is sent with the GPS location of the crew at the time that the panic request was issued to better locate the crew on the map.

NOTE: This feature is optional. It is only available on the mobile application if an active "Panic" alert type is referenced as the **Panic Alert Type**[global configuration](#) parameter.

The crew may also set a timer for a specified amount of time such that, if they don't manage to manually turn off the timer before it expires, a "Timed Event" alert is raised for a dispatcher to follow up with the crew and send help as needed. This is used when the crew feels they may be in a dangerous or hazardous situation and would like higher visibility from the dispatcher. The alert only becomes in effect and shows up on a dispatcher's alert queue once expired. If the crew manages to close the event before it has expired the alert is closed and does not appear on any dispatcher's alert queue.

NOTE: This feature is optional. It is only available on the mobile application if an active "Timed Event" alert type is referenced as the **Timed Event Alert Type**[global configuration](#) parameter.

Sending / Receiving Mail Messages

A crew can received mail messages on their device as well as reply and compose new messages using the mobile application.

The End Of Shift Process

When the crew completes their shift the following steps are involved:

- Activity assignments the crew did not manage to complete are returned to be rescheduled. Other types of tasks that were not completed are canceled. A remote message is sent for each such task with all the information captured on the device for the task for audit and reporting purposes.
- A remote message to complete the shift on the server is sent.
- If timesheet functionality is enabled, a message to prepare timesheets for the mobile workers is sent to the server. The crew may request to review their timesheet on the device or later at the office. Refer to [Reviewing Timesheets](#) for more information.

Reviewing Timesheets

If your organization uses timesheets functionality, timesheet records are created automatically at end of shift for each period of time the resource was part of the shift.

As part of the shift completion process on the mobile device, the crew is asked whether they would like to review their timesheet records on the device.

If they chose to do so, the application waits for the timesheet records to be prepared on the server and when ready, presents them to the crew for review. Whether the crew is allowed to adjust the reported time depends on your business practices and settings.

The crew can choose at any time to not review their timesheet records on the device and wait to do that later when they are back at the office.

Refer to [About Timesheets](#) for more information.

Switching Devices Mid Shift

If the crew needs to switch mobile devices for any reason, they need to contact the dispatcher to disconnect their old device from their shift.

Once the shift is forced logoff from the old device by the dispatcher, the crew can then logon from another device and request shift details from the server. All the tasks that were previously dispatched to the previous device are dispatched to the new device. The status of the tasks reflects the latest information sent to the server before the device was switched. Data on the previous device that did not reach the server is not recoverable.

Refer to [Force Logoff](#) for more information.

About The Contractor's Application

Work assigned to a contractor's work capacity may be managed by their workers using the contractor mobile application. They use the mobile application to request a list of tasks to work on and enter the completion information needed to complete the work.

This section provides a high level overview of the contractor mobile application as well as describes functionality supported by it that you may want to take advantage of.

NOTE: This is not a user guide to the mobile application itself.

Overview

A manager from the contractor company uses the [Contractor Work Management](#) portal to review, approve, and reject activities your company has committed to their company. They then can optionally manually tag each activity they have approved to be worked by one of their workers.

Refer to [Managing Committed Work](#) for more information.

Note that these workers are not defined in your system as resources and as such the system does not schedule nor push work to them. Instead they are assigned a unique tag by their manager and use the mobile application to pull the list of activities their manager has tagged to them. These activities are pulled to the device for them to be worked on.

The worker reviews the list of activities and work each activity in no particular order.

Once tasks are on the device, the worker can work disconnected if wireless communication is not available. As they work each task, they enter the necessary completion details on the mobile device. The mobile application queues up these messages and when the device is connected, sends them to the server to update the corresponding master copy.

At the end of their day, the worker uses a menu option to clear up the data on the device.

No Dispatch Process

There is no shift to start nor end, as well as no dispatch process. The worker needs to periodically refresh their task list to reflect changes made to it by their manager.

Completion Only Updates

The system does not track the contractor individual workers nor their attempts to perform an activity. As such, interim status updates are not applicable. The only update the worker may provide with respect to an activity is completing it.

If the worker needs to decline an activity, they need to contact their manager to remove the activity from their list and then refresh their list on the device.

About Remote Messages

Remote messages are used to send data to and from a mobile device. This section describes the mechanism that allows the main application to send and receive messages to a mobile device.

Messages To A Mobile Device

Depending on your mobile application requirements, certain business events triggered by the main application may need to be communicated to a mobile user's specific mobile device.

When using the Hybrid mobile application, the payload and rules associated with such message are defined using a [Message To Device](#) mobile component. You may use the **Invoke Remote Script** (M1-InvokeRSIScript) business service to post a message to be processed by a mobile device. Refer to the description of this business service for more information.

The mobile device periodically polls for queued messages and process them in the order they were posted.

You may configure an expiry time as a parameter to the **Expire Old Remote Messages** (M1-EXOLDRMMS) algorithm type. If specified, messages that have not been picked up by the specified time are automatically discarded by the remote message monitor batch program. Callback and Error scripts may exist for these kinds of messages and they too are handled by the remote message monitor batch process.

NOTE: It is important to schedule the **Remote Message Monitor** batch process to expire messages and finalize the state of processed messages. Refer to [M1-RMMSG](#) batch control for more information.

Messages From A Mobile Device

Asynchronous messages from the field may be received continually throughout the day. The processing of such messages ensures against data loss by saving the message before processing it. This approach guarantees message delivery despite of potential processing errors. Refer to [Guaranteed Delivery](#) for more information.

Messages from the crew are received continually throughout the day. If message processing fails at any time, a To Do entry is created for a user to review and resolve the problem. You can view the content of any message, try to reprocess the message as is, or edit the message content and then reprocess.

Some errors may arise due to a temporary situation and may resolve on their own when the system tries to process the erroneous message again by the remote message monitor batch process.

NOTE: It is important to schedule the **Remote Message Monitor** batch process to retry failed messages periodically. For example, it should run every 5 minutes or so. Refer to [M1-RMMSG](#) batch control for more information.

Maintaining Remote Messages

This portal is used to view a messages sent to the crew or sent from the crew.

Refer to [About Remote Messages](#) for more information.

You can access this portal from the **Main Menu > Service Management > Remote Message**. You are brought to a query portal with options for searching for a specific remote message. Once a remote message has been selected, you are brought to the maintenance portal to view and maintain the selected record.

The **Remote Message** zone on the portal's **Main** tab page provides information about the remote message, including the various maintenance actions applicable to its current state.

To Do Processing

There are numerous events that trigger the creation of email-like messages describing work that requires human attention. For example, if a bill segment has an error, a message is created to highlight such. We refer to each message as a **To Do Entry**. In this section, we describe how to view, manage and work To Do entries.

FASTPATH: We recommend familiarizing yourself with the topics described in [The Big Picture of To Do Lists](#) to fully appreciate how To Do entries are managed by the system.

The Big Picture of To Do Processing

The topics in this section provide background information on To Do processing functionality.

To Do Role

Each To Do entry is assigned a specific To Do Role. The role defines the users who may work on the entry. The permissible roles that may be assigned to a To Do entry are defined on the entry's To Do type. Refer to [To Do Entries Reference A Role](#) to learn more about initial role assignment.

After an entry is created, its role may be changed to any role defined as valid for the entry's To Do type. Manual changes of a To Do entry's role are recorded on the [To Do entry log](#).

To Do Priority

A To Do entry is associated with a priority level representing its processing order compared to other entries. Refer to [The Priority Of A To Do Entry](#) for more information on how an entry's priority is initially assigned.

A user may manually override a To Do entry's priority at any time. Once overridden, automatic recalculation of the priority is no longer performed so as not to override the value explicitly set by the user. Manual overrides of a To Do entry's priority are recorded on the [To Do entry log](#).

Work Distribution

To Do entries may be manually assigned to users either by their supervisors or picked up by the users themselves. The system can also be configured to dynamically assign the highest priority work to users on demand whenever they are available for fixing errors.

NOTE: Assignment on demand is in a plug-in. You can customize the [next assignment](#) plug-in on the installation options to follow your organization's work distribution rules. For example, when a user asks for the next assignment, business rules may be used to only consider To Do entries the user has the proper [skills](#) to resolve, prioritizing them based age, impacted dollar amount, etc.

One of the main benefits of the assignment on demand approach is that work is allocated real-time based on most current information such as skills, priority, user availability, etc. This also means that work remains unassigned until a user is actually going to work on it making it available to any qualified user.

Manual assignments of To Do entries are possible as follows:

- Supervisors may use the [supervisor function](#) pages to review and assign To Dos to users.
- A user may use the [To Do list](#) pages to review and assign To Dos to themselves or forward any of their To Dos to someone else.
- All of the above assignment functions may also be performed using the [To Do Search](#) page.


Assignment "on demand" is possible from the [Current To Do](#) zone on the Dashboard portal. Refer to this zone to learn about all the functions it provides.

A User's Current To Do

Users have the option of drilling down on an entry from various list queries in the system. Refer to [Working On A To Do Entry](#) for more information on drilling into an entry.

Drilling down on an entry from the following To Do list queries (using the entry's message description only) also sets the entry as the user's **Current** To Do entry:

- [To Do List](#)
- [To Do Search](#)

Additionally, users may have the option to click on the Current To Do icon,  , on the Current To Do dashboard zone to request their next task. The returned To Do entry automatically becomes the user's **Current** To Do.

Regardless of how a To Do entry became current its information is displayed on the Current To Do zone.

FASTPATH: Refer to [Current To Do](#) zone to learn about all the functions it provides to efficiently resolve the current To Do and other entries that may potentially be [related](#) to it.

If a To Do entry becomes current as a result of drilling down from a list query, the system supports the ability to drill down to the **Next** and **Previous** entries respective to the **Current** To Do entry in that list.

NOTE: The **Previous To Do** button and **Next To Do** button in the Current To Do zone allow you to scroll through additional To Do items in your current work list without returning to the list query transaction. These buttons set the respective To Do entry to now being current and take you to its related transaction or to a script associated with the To Do Entry's message (not to the To Do entry itself). This is yet another way of drilling down on an entry.

If a To Do entry becomes current using the Current To Do zone's next assignment functionality, that entry is not associated with any work list to scroll through. The Next and Previous To Do buttons are therefore disabled in this case.

NOTE: Current Work List. Current work list refers to the list from which the current To Do entry was selected. Each new election of a current To Do overrides the previous work list with the new list the entry was selected from. If not selected from a list then the current work list is considered empty.

Related To Do Entries

A user working on a To Do may want to know of other To Dos that may be related to the same problem. If can be worked by the same user they may get resolved faster. It seems almost impossible to programmatically determine whether two To

Dos really stem from the same problem. The user should therefore be presented with a list of potentially related To Dos from which he/she can select the truly relevant ones.

A To Do entry is considered potentially related to a given To Do if they reference a common entity via a characteristic. For example, if a given To Do references a certain account then any To Do referencing that same account is considered related to it.

NOTE: The common entity is any characteristic value linked to the To Do entry where the characteristic type is a **Foreign Key** and the main table of the foreign key is classified as a **Master** or **Transaction** table.

The To Do Search page supports a **Related To Do filter** that lists all related To Dos for a given entry. Various pages showing To Do information often provide an indication as to how many non-complete entries may be related to a given To Do entry. Clicking on the indication's hyperlink transfers the user to the To Do Search query pre-populating it to list the related entries for the specific To Do entry.

FASTPATH: Refer to the [To Do Search](#) page for more information on all the functions it supports.

Complete A To Do

After you finish working on an entry, you can mark it as **Complete**. At this point, you may want to review potentially [related](#) To Do entries and complete those that have been addressed by resolving the problem associated with the current entry.

If you cannot resolve the problem, you can forward the To Do to another user, thus assigning the To Do to another user.

FASTPATH: Refer to [How To Complete To Dos](#) for more information.

NOTE: It is possible for the system to automatically mark an entry as **Complete**. Refer to [The Lifecycle Of A To Do Entry](#) for more information.

Managing To Do Entries

This section describes the various pages available to view and work on To Do entries.

To Do Supervisor Functions

This section describes To Do transactions that aid supervisors in the monitoring and assignment of To Do entries.

Supervisor To Do Summary

This Supervisor To Do Summary page summarizes To Do entries by To Do Type. Only entries that are **Open** (i.e., not assigned to a specific user) or **Being Worked On** are summarized.

NOTE: Only To Do entries that are assigned to a role to which you belong are summarized on this page. Refer to [To Do Entries Reference A Role](#) for more information.

Use **Menu > To Do > Supervisor To Do Summary** to open this page.

Description of Page

This page summarizes all To Do entries that are **Open** or **Being Worked On**.

- **To Do Type** contains the description of the To Do type.
- **Priority** provides a total count of To Do entries of a given type by priority. For each distinct priority the description includes the priority value followed by the corresponding number of entries. If all entries are of the same priority no additional count is displayed.
- **Total** contains the total of **Open** and **Being Worked On** entries. The hyperlink navigates to the [Supervisor To Do Assignment](#) transaction, setting it to display all **Open** and **Being Worked On** entries for the given To Do Type.
- **Open** contains the number of To Do Entries whose status is **Open**. The hyperlink navigates to the [Supervisor To Do Assignment](#) transaction, setting it to display all **Open** entries for the given To Do Type.
- **Being Worked On** contains the number of To Do entries whose status is **Being Worked On**. The hyperlink navigates to the [Supervisor To Do Assignment](#) transaction, setting it to display all **Being Worked On** entries for the given To Do Type.

The "age bars" that appear to the right of the totals highlight the age of the To Do entries. The colors red, yellow and green are used to highlight the relative number of old, middle-aged and new To Do entries. The age associated with each of these colors is [user-configurable](#).

Supervisor User Summary

The Supervisor User Summary page summarizes To Do entries by the user who is assigned to work on the entries. Only entries that are **Being Worked On** are summarized.

NOTE: Only To Do entries that are assigned to a role to which you belong are summarized on this page. Refer to [To Do Entries Reference A Role](#) for more information.

Use **Menu > To Do > Supervisor User Summary** to open this page.

Description of Page

This page contains a grid that displays all assigned To Do entries by User ID and To Do type. It indicates how many To Do entries of the displayed Type are assigned to the user (Assigned this Type) and how many To Do entries in total are assigned to the user (Assigned Total).

Clicking the hyperlinks transfer you to the [Supervisor To Do Assignment](#) transaction. On this transaction, you can assign / reassign To Do entries to a specific user. The difference is that clicking the To Do Type hyperlink causes only **Open** entries to be displayed on the Assignment transaction. Clicking the count in the Assigned this Type column causes **Being Worked On** entries assigned to the displayed user to be shown.

Supervisor To Do Assignment

This page lists To Do entries of a given type and allows a supervisor to assign / reassign them to a designated user. Use **Menu > To Do > Supervisor To Do Assignment** to open this page.

Description of Page

This page contains **Open** or **Being Worked On** To Do entries of a given type. These entries can be assigned to a specified user one-by-one or several at a time. The following points describe this functionality:

- You can select a different **To Do Type** by typing in a value or pressing the search button in the upper right corner. An **Info** description of the To Do type appears on the left.
- You can change the list of entries displayed according to their status. Select the **Filter by** status to be All (blank), **Open**, or **Being Worked On**.
- For **Being Worked On** entries, you may also **Filter by User** to limit the search to a specific user assigned to the To Do entries.

- If you want to re-sort the entries by a different value, click the search button adjacent to **Extra Information**. Each row's respective sort value is displayed in the **Extra Information** column.

NOTE: Extra Information. Each To Do list's default **Extra Information** is defined on the list's [To Do Type](#). Also note that each To Do list's available **Extra Information** differs.

- Click the **Select All** button to select all entries displayed in the grid. Click the **Clear All** to remove the selections in the grid.
- Use the checkbox if you want to assign specific entries.
- Entries are further sorted by their **Priority**. Highest in priority are listed first.
- **Create Date/Time** contains the date and time the To Do entry was created by the system. Click the Create Date/Time hyperlink to:
 - View more detailed information about the entry.
 - Add comments to the To Do entry.
 - Add a log entry of type **User Details** to record information about the To Do entry.
 - Forward the To Do entry to another user.
 - Send Back a To Do entry that was forwarded to you by another user.
- Click the number of **Related** entries hyperlink to navigate to the [To Do Search](#) query to list potentially related entries to the current entry.
- **Worked By** shows the user currently assigned to the To Do entry.
- **To Do Entry Status** is the current status of the To Do entry.
- **Message** contains a description of the To Do entry. Use the message description hyperlink to drill into the page on which the entry's respective object is displayed. If the To Do type has been configured to associate a script with this To Do entry's message number, then instead of being taken to the respective page, the associated script is launched. Refer to [Working On A To Do Entry](#) for more information.
- You can assign entries to a user:
 - Enter a user ID or search for a user using the **Assign To User** field. You can only select a user that is associated with a role that is associated with this **To Do Type**. Note that the current number of assignments is displayed adjacent.
 - Select one or more entries to assign to this user and click the **Assign** button. The entries' status becomes **Being Worked On**.

To Do List Functions

This section describes To Do pages that aid users in the assignment and working of entries in To Do lists.

To Do Summary

This page summarizes To Do entries that meet the following qualifications:

- The entry is **Open** and you belong to the entry's To Do role. Refer to [To Do Entries Reference A Role](#) for more information about roles.
- The entry is **Being Worked On** and you are the user to which it is assigned.

Use **Menu > To Do > To Do Summary** to open this page.

Keyboard Shortcut**Alternate Shortcut**

Alt+X

Alt+Shift+X

NOTE: Refer to [Shortcut Key Summary](#) for information about the alternate shortcut.

FASTPATH: Refer to [The Big Picture of To Do Lists](#) for an overview of To Do functionality.

Description of Page

This page is dedicated to a grid that summarizes the To Do entries on which you can work. The following information is displayed:

- **To Do Type** contains the description of the To Do type.
- **Priority** provides a total count of To Do entries of a given type by priority. For each distinct priority the description includes the priority value followed by the corresponding number of entries. If all entries are of the same priority no additional count is displayed.
- **Total** contains the total of **Open** and **Assigned to You** entries. The hyperlink navigates to the [To Do List](#) transaction, setting it to display all **Open** and **Being Worked On** entries assigned to you for the given To Do Type.
- **Open** contains the number of To Do Entries whose status is **Open**. The hyperlink navigates to the [To Do List](#) transaction, setting it to display all **Open** entries for the given To Do Type.
- **Assigned to You** contains the number of To Do entries that are assigned to you whose status is **Being Worked On**. The hyperlink navigates to the [Supervisor To Do Assignment](#) transaction, setting it to display all **Being Worked On** entries for the given To Do Type.

The "age bars" that appear to the right of the totals highlight the age of the To Do entries. The colors red, yellow and green are used to highlight the relative number of old, middle-aged and new To Do entries. The age associated with each of these colors is [user-configurable](#).

To Do List

This page lists To Do entries of a given type that are either **Open** (e.g., not assigned to a specific user) or are assigned to you.

NOTE: Restricted by role. Only those entries assigned to a role to which you belong may be seen on this page. Refer to [To Do Entries Reference A Role](#) for more information.

FASTPATH: We recommend familiarizing yourself with the topics described in [The Big Picture of To Do Lists](#) to fully appreciate how To Do entries are managed by the system.

Use **Menu > To Do > To Do List** to open this page.

Description of Page

This page contains a grid that summarizes the To Do entries of a given type on which you can work.

The **Info** section contains a concatenation of the **To Do Type's** description and the user's name.

The area beneath **Info** provides you with options that control which To Do entries appear and the order in which they are displayed. The following points describe the various options:

- Use **Sort By** to define how you want the To Do entries sorted. The following options are available:
 - **Assigned Date - On or After.** This option sorts the To Do entries in chronological order (based on the date the entry was assigned to you) starting from the input **Date**. If you choose this option, the system assumes you want to start



from the earliest To Do entry. If you want to start from a date other than the earliest To Do entry, enter the desired **Date** and click the adjacent search button.

- **Assigned Date - On or Before.** This option sorts the To Do entries in reverse chronological order (based on the date the entry was assigned to you) starting from the input **Date**. If you choose this option, the system assumes you want to start from the today's date. If you want to start from a date other than today's date, enter the desired **Date** and click the adjacent search button.
- **Create Date - On or After.** This option sorts the To Do entries in chronological order (based on the date the entry was created) starting from the input **Date**. If you choose this option, the system assumes you want to start from the earliest To Do entry. If you want to start from a date other than the earliest To Do entry, enter the desired **Date** and click the adjacent search button.
- **Create Date - On or Before.** This option sorts the To Do entries in reverse chronological order (based on the date the entry was created) starting from the input **Date**. If you choose this option, the system assumes you want to start from the today's date. If you want to start from a date other than today's date, enter the desired **Date** and click the adjacent search button.
- **Extra Information.** This option sorts the To Do entries by the value of the **Extra Info**.
- Use **Extra Info** to choose data related to the To Do Entry that should be used to sort the To Do entries. This information is displayed in each entry's **Extra Information** column.

NOTE: Extra Info. Each To Do list's available **Extra Info** is defined on the list's [To Do Type Sort Keys](#). Also note that each To Do list's available **Extra Info** differs based on the configuration of the To Do Type.

- Use **Filter by** to restrict the To Do entries based on their status. The following options are available:
 - **All.** This option shows all entries that are **Open** or **Being Worked On**.
 - **Being Worked On.** This option shows all entries that are **Being Worked On** (and assigned to you).
 - **Open.** This option shows all entries that are **Open**.

The grid that follows contains the To Do entries that match your criteria. The following information appears:

- Click the **Select All** button to select all entries displayed in the grid. Click the **Clear All** button to remove the selections in the grid.
- Use the checkbox if you want to mark specific entries as **Completed** or **Being Worked On** using the **Complete** or **Assign** buttons at the bottom.
- A Comments icon  appears if there are comments on the To Do item and a Log Entry icon  appears if there is a log entry of type **Forwarded**, **Sent Back** or **User Details** on the To Do item. This column appears if at least one To Do has either comments or a special log entry.
- Entries are further sorted by their **Priority**. Highest in priority are listed first.
- **Create Date/Time** contains the date and time the To Do entry was created by the system.
- Click the Create Date/Time hyperlink or either of the special icons to navigate to the [To Do entry](#) page.
- Click the number of **Related** entries hyperlink to navigate to the [To Do Search](#) query to list potentially related entries to the current entry.
- **Assigned Date/Time** contains the date and time the To Do entry was assigned to a user. This field is only populated on entries that are **Being Worked On**. If a user assigned the To Do entry, the user ID and name of the user who assigned the To Do entry are displayed beneath the Date/Time.
- **Message** contains a description of the To Do entry. Use the message description hyperlink to drill into the page on which the entry's respective object is displayed. If the To Do type has been configured to associate a script with this To Do entry's message number, then instead of being taken to the respective page, the associated script is launched. Refer to [Working On A To Do Entry](#) for more information.

NOTE: Current To Do. Drilling into an entry from this page sets it to be the user's current To Do. Refer to [A User's Current To Do](#) for more information on how to efficiently scroll through additional entries in your current work list.

CAUTION: Automatic Assignment! When you drill into an entry using the message description hyperlink from this page, the system changes the status of the entry to **Being Worked On** and assigns it to you.

- **To Do Entry Status** is the current status of the To Do entry.
- **Extra Information** contains the data for the entry that corresponds to the Extra Info chosen in the sort options above.
- **To Do ID** contains the unique identifier of the To Do entry.
- If you want to mark one or more entries as **Being Worked On**, check the desired entries and then click the **Assign** button.
- If you want to mark one or more entries as **Completed**, check the desired entries and then click the **Complete** button. Completed entries do not appear on the To Do list (but they are retained on the database for audit purposes).

To Do Entry

The topics in this section describe the pages on which a To Do entry and its associated To Do log are maintained.

To Do Entry - Add Dialog

The To Do Entry transaction features a dialog that can add new To Dos (if the system is configured with [manual To Do types](#)). This dialog appears if you open the **Menu > To Do > To Do Entry > Add** (it also appears if you click the clear button when on the To Do Entry page).

Description of Page

The **To Do Type** is the type of To Do entry that you want to create. Only To Do types that are set up for manual usage appear in the list.

The default **To Do Type Priority** is displayed adjacent. Use **Override Priority** to override the default [priority](#) assigned by the system.

Type a **Subject** for your To Do entry. Depending on how the To Do type is set up, the subject may appear as the message in the To Do entry and To Do lists.

Enter a **Comment** to provide a longer description of and more details about the To Do entry.

Send To indicates whether the new To Do will be assigned to a **Role** (the default) or **User**.

The **Role** indicates the role to which this To Do entry will be assigned. Only users of the specified role will be able to view and work on the To Do entry. The role is defaulted from the To Do type. If you want to assign the To Do to a different role, change the specified role. You can only assign the To Do entry to one of the roles specified on the To Do type.

User appears if **Send To** is set to **User**. Use this option if you want to assign the To Do to a specific user. The user must be a member of the selected role.

If the [To Do type](#) defines valid characteristic types, the **Characteristics** collection is visible. If any of the characteristic types are flagged as Default on the To Do type, the characteristic type with its default value (if applicable) are displayed. If applicable use the characteristics to define other information related to this new To Do entry. In Oracle Utilities Customer Care and Billing, for example, you can link the account and premise that you are attempting to start service for.

Click **OK** to add the To Do Entry.

To Do Entry - Main

You use this page to perform various actions or modify details about a To Do entry. Use **Menu > To Do > To Do Entry > Search** to open this page.

NOTE: Restricted by role. Any user may change Comments for any To Do Entry. In addition, any user may add User Log details. However, for all other changes, you may only change To Do entries that are assigned to you or ones that are **Open** and assigned to a role to which you belong. Refer to [To Do Entries Reference A Role](#) for more information about roles.

Description of Page

To Do contains a concatenation of important information about the To Do entry. **To Do ID** is the system-assigned unique identifier of the entry. These values only appear after the case is added to the database.

NOTE: Formatting may be performed by a plug-in. Refer to [To Do Info May Be Formatted By An Algorithm](#) for more information.

To Do Type displays the entry's To Do Type and its description. Refer to [To Do Entries Reference A To Do Type](#) for more information about To Do types.

Role displays the role of the users who may be assigned to the entry along with a description of the role. Refer to [To Do Entries Reference A Role](#) for more information about roles.

Message is a display-only field that describes why the entry exists. If the To Do entry is in the status of **Open** or **Being Worked On**, the message description allows you to drill into the page on which the entry's respective object is displayed. If the To Do type has been configured to associate a script with this To Do entry's message number, then instead of being taken to the respective page, the associated script is launched. Refer to [Working On A To Do Entry](#) for more information.

Status indicates if the entry is **Open**, **Being Worked On**, or **Complete**. Refer to [The Lifecycle Of A To Do Entry](#) for more information status.

Use **Priority** to manually override the To Do entry's current [priority](#). Any manual change to the priority is recorded on the [To Do entry log](#).

Assigned To displays the user who is working the entry.

Sent By displays the user who assigned this To Do to the current user.

Created By only appears if a background process that uses To Do entries to report object-specific errors created the To Do entry. The field displays the background process description and batch run number in which the entry was created. Refer to [To Dos Created for Object-Specific Error Conditions](#) for more information.

Extract Info only appears if the To Do is being sent out of the system. Extract Info describes the status of the extract (for example whether it has been processed or not). Refer to [To Do Entries May Be Routed Out Of The System](#) for more information.

Use **Comments** to describe anything interesting or unusual about the To Do entry. The **Comments** field is disabled if the To Do entry is complete.

If the To Do entry has non-complete [related entries](#) a summary description appears below the comments line. The description states the total number of related entries, how many are still **Open**, how many are assigned to you and how many assigned to others. Use this hyperlink to navigate to the [To Do Search](#) query to list these entries.

The **Actions** buttons allow you to **Complete** the To Do, **Forward** the To Do to another user or Role, or **Send Back** a To Do that has been forwarded to you by another user. All buttons are disabled once a To Do is complete. Refer to [To Do Entry Actions](#) and the [How To](#) section for more information on using these actions.

The **Characteristics** collection is visible for **manual** To Do entries or for **automatic** To Do entries that have characteristics. **Characteristic Type** indicates the type of characteristic. Specify a valid **Characteristic Value** for the characteristic type.

NOTE: Only characteristic types defined as permissible on a To Do entry may be linked to a To Do. Refer to [Setting Up Characteristic Types & Their Values](#) for more information about characteristic types. In addition, for **manual** To Do types, the characteristic type must also be defined as valid for the To Do entry's [To Do type](#).

The **Log** grid displays log entries (in reverse chronological order) that audit the progress of a To Do entry. You cannot modify or delete log entries that have been created by the system. You can manually add a log entry of type **User Details** and you may remove a log entry before it is saved. Once a **User Details** log entry is saved, you cannot modify or delete it.

The Log grid displays the following information:

Log Date/Time displays the date and time when the log entry was created.

Log Type indicates the To Do entry event for which the log entry was created. For example, when the To Do entry is created, the Log Type indicates **Created**. For more information refer to the [Log Entry Events](#) table below.

The **Details** column contains user-specified or system-generated information about the log entry. This field is blank for some types of log entries.

Log User displays the user who caused the log entry to be created.

Assigned To displays the user to whom the To Do entry was assigned (if any).

To Do Entry - Keys/Values

The Keys / Values page contains detailed information about the sort keys and drill keys associated with a To Do entry. This information may not be modified. It is shown to assist support personnel.

Use **Menu > To Do > To Do Entry > Search** to search for the desired To Do entry, then navigate to the **Keys/Values** page.

Description of Page

The **Sort Keys** grid contains an entry's sort keys and values. This information is display-only.

FASTPATH: Refer to [To Do Type - Sort Keys](#) for more information.

The **Drill Keys** grid contains an entry's drill keys and values. This information is display-only.

FASTPATH: Refer to [To Do Type - Drill Keys](#) for more information.

To Do Entry Actions

You can perform the following actions on a To Do Entry.

Complete

Clicking **Complete** causes the To Do entry's status to be **Completed**. When you complete a To Do, a [log entry](#) is created that indicates when the To Do was completed and by whom. This button is enabled for To Do entries with a status of **Open** or **Being Worked On**.

FASTPATH: Refer to [The Lifecycle Of A To Do Entry](#) for more information.

Forward

Clicking **Forward** causes the Forward To Do Entry dialog to appear, allowing you to forward the To Do to another user or role. This button is enabled for **Open** To Do entries assigned to a role to which you belong or **Being Worked On** To Do entries assigned to you.

Send To indicates whether the To Do will be forwarded to a **Role** or **User** (the default).

User appears if **Send To** is set to **User**. The **User** indicates the user to which this To Do entry will be forwarded. The user must be a member of one of the roles specified for the To Do type.

Role is enabled if **Send To** is set to **Role** or if set to **User** and a User is specified. It indicates the role to which this To Do entry will be forwarded. Only users of the specified role will be able to view and work on the To Do entry. The role is defaulted from the To Do type. You can only forward the To Do entry to one of the roles specified on the To Do type. If forwarding to a user the role must also be valid for the specified user.

The **Details** field provides an explanation of why a user is forwarding the To Do. This information appears in the To Do log.

Click **OK** to forward the To Do entry.

Send Back

Clicking **Send Back** causes the Send Back To Do Entry dialog to appear, allowing you to send back the To Do to the user who forwarded it to you. This button is enabled if the To Do entry's status is **Being Worked On** and it was forwarded to you by another user.

User is a display only field that indicates the user who assigned you the To Do.

Role may be used to also change the entry's To Do Role. The role must be a permissible role specified on the To Do type that is also valid for the user this entry is sent back to. By default, if the sent to user is not a member of the entry's current role, the system automatically defaults the role to the first permissible role that is also valid for the user.

The **Details** field provides an explanation of a why the user is sending back the To Do or any other desired information. This information appears in the To Do log.

Click **OK** to send the To Do back to the assigning user.

Log Entry Events

The following table contains the events that cause a log entry to be created by the system.

Event	Log Type	To Do Entry Status	Content of Details	Content of Assigned To (on To Do entry and To Do Log)	Log User ID	Contents of Sent By (on To Do entry)
A To Do is created (by the system or by a user)	Created	Open			ID of user who created the To Do	
A To Do is created (by the system or by a user) in Being Worked On status	Created	Being Worked On		User ID of user who was assigned this To Do	ID of user who created the To Do	Creating user's ID
A To Do is completed (by the system or by a user)	Completed	Complete			ID of user who completed the To Do	

Event	Log Type	To Do Entry Status	Content of Details	Content of Assigned To (on To Do entry and To Do Log)	Log User ID	Contents of Sent By (on To Do entry)
A user takes an open To Do	Assigned	Being Worked On	Selected by user	Assigned To User ID	Assigned To User ID	ID of user who assigned the To Do
A supervisor assigns a To Do to a user (also happens if they reassign)	Assigned	Being Worked On	Assigned to user	Assigned To User ID	ID of user who assigned the To Do	ID of user who assigned the To Do
A user forwards an entry to a user	Forwarded	Being Worked On	<i>Log details entered by user</i>	Forwarded To User ID	ID of user who forwarded the To Do	ID of user who forwarded the To Do
A user forwards an entry to a role	Forwarded	Open	Forwarded to role <i>Role Description (Role ID) + Log details entered by user</i>		ID of user who forwarded the To Do	ID of user who forwarded the To Do
A user sends back a To Do	Sent back	Being Worked On	<i>Log details entered by user</i>	Sent Back To User ID	ID of user who sent back the To Do	ID of user who sent back the To Do
A user adds a log entry	User Details	Status remains the same	<i>Log details entered by user</i>		ID of user who added the entry	
A user manually overrides priority	Override Priority	Status remains the same	Priority changed from <i>old value</i> to <i>new value</i>	Assigned To User ID	ID of user who made the change	
Priority is updated by a calculate priority algorithm	Calculated Priority	Status remains the same	Priority changed from <i>old value</i> to <i>new value</i>		ID of user that added or updated the To Do or the user associated with the Recalculate Priority batch job.	

To Do Search

This page allows you to search for To Do entries using various search criteria, and perform various actions on selected ones.

NOTE: No restriction by role. This page returns all To Do entries that meet the search criteria, regardless of the To Do entry's role. Using the search, a user may drill into the details of a To Do entry even if the To Do is not assigned to the user. Refer to [To Do Entries Reference A Role](#) for more information about roles.

FASTPATH: We recommend familiarizing yourself with the topics described in [The Big Picture of To Do Lists](#) to fully appreciate how To Do entries are managed by the system.

To Do Search - Main

Use **Menu > To Do > To Do Search** to open this page.

The top half of the page is where you enter the criteria used to search for To Do entries.

NOTE: Multiple search criteria may be specified. You can search for To Do entries using a combination of search criteria. For example, if you enter both a **To Do Type** of **TD-NOMR** and a **Status** of **Being Worked On**, only those To Do entries for To Do type **TD-NOMR** that are currently **Being Worked On** are displayed.

CAUTION: Try to be as specific as possible when entering search criteria. Why? Because entering open-ended search criteria may have a severe impact on response times.

The following table describes each of the different search filters.

Search Filter	Description
Related To Do	<p>If you're searching for potentially related entries for a given To Do entry, choose Specific To Do for the Related To Do Filter. Enter the To Do ID.</p> <p>A filter value of Not Applicable defaults.</p>
To Do Type	<p>If you're searching for a To Do entry and you know the value of a To Do Type linked to the To Do entry, choose Specific To Do Type for the To Do Type Filter. Enter the To Do Type.</p> <p>A filter value of All defaults.</p>
Char Type	<p>If you're searching for a To Do entry and you know the value of a characteristic linked to the To Do entry, choose Specific Char Type for the Char Type Filter. Enter the Characteristic Type and Characteristic Value.</p> <p>A filter value of Not Applicable defaults.</p> <p>Only characteristic types that include a characteristic entity of To Do Entry and that have been marked to allow searching by characteristic value are available.</p>
User	<p>Use this search method if you know the Created By User, Assigned To User or Completed By User for the To Do entry. Enter the appropriate User Id.</p> <p>A filter value of Not Applicable defaults.</p>
Date Range	<p>Use this field if you're searching for a To Do entry and know the Date Created or Date Completed.</p> <p>A filter value of Date Created defaults. The Start Date/Time defaults to the current date, 12:00AM and the End Date/Time defaults to the current date, 11:59PM.</p> <p>If you enter a Date Range Filter of All, the system ignores this search method.</p>
Status	<p>Use this search method to restrict the search to To Do entries with a specific status value. You may search for To Do entries in the status of Open, Being Worked On or Completed. You may also perform a combination search for To Do entries that have a status value of either Open or Being Worked On.</p> <p>A filter value of All defaults.</p>



NOTE: The system allows any combination of values. The system does not try to validate your search combinations, even if they do not make sense. For example, if you enter a **User Filter** value of **Completed By User** and a **Status Filter** value of **Open**, you will not get an error. In this case, the system will find no records that match your criteria (because To Do entries that are Open do not have a Completed By User).

The system shows the total number of To Do entries that satisfy your search results immediately below the grid.

The first group of To Do entries is displayed in the grid at the bottom of the page.

- Click the **Select All** button to select all entries displayed in the grid. Click the **Clear All** to remove the selections in the grid.
- Use the checkbox if you want to assign, reopen or complete specific entries. Only **Open** or **Being Worked On** entries may be selected.

FASTPATH: Refer to [To Do Search Actions](#) and the [How To](#) section for more information on using these actions.

- A Comment icon  appears if there are comments on the To Do item and a Log Entry icon  appears if there is a log entry of type **Forwarded**, **Sent Back** or **User Details** on the To Do item. This column appears if at least one To Do item has either comments or a special log entry. Use either icon to navigate to the To Do entry page.
- Entries are sorted by their **Priority**. Highest in priority are listed first.
- **Create Date/Time** contains the date and time the To Do entry was created by the system. Use the hyperlink to navigate to the To Do entry page.
- Click the number of **Related** entries hyperlink to redisplay this page listing potentially related entries to the current entry.
- **Last Assigned Date/Time** contains the date and time the To Do entry was most recently assigned to a user, if applicable, along with the Assigned To user.
- **Message** contains a description of the To Do entry. If the To Do entry is in the status of **Open** or **Being Worked On**, the message description allows you to drill into the page on which the entry's respective object is displayed. If the To Do type has been configured to associate a script with this To Do entry's message number, then instead of being taken to the respective page, the associated script is launched. Refer to [Working On A To Do Entry](#) for more information.

NOTE: Current To Do. Drilling into the page by clicking the Message hyperlink does not automatically assign **Open** To Do entries to the user (as the To Do List page does when clicking on the Message). But it does treat this as the user's current to do, enabling the functionality in the Current To Do dashboard zone. Refer to [A User's Current To Do](#) for more information.

- **Status** contains the current status of the To Do entry.
- **Completion Date/Time** contains the date and time the To Do entry was completed and the user who completed the entry, if applicable.
- **To Do Type** contains a description of the To Do type for the To Do entry.

To Do Search Actions

You can perform the following actions on selected To Do Entries.

Assign Selected Entries

Use the **Assign** button to assign or reassign the selected entries to a user. Enter a user ID or search for a user using the **Assign To User** field. By default this field is populated with the current user's ID. Note that the user's current number of assignments is displayed adjacent.

Upon successful processing, the selected entries are assigned to the specified user and the status of **Open** entries becomes **Being Worked On**. Additionally, a corresponding [log entry](#) is added to each To Do.

NOTE: Only users with **Supervisor Assignment** access rights may reassign entries that are not currently assigned to them.

Reopen Selected Entries

Use the **Reopen** button to unassign the selected entries from their users forwarding them back to their role.

Upon successful processing, the selected entries become **Open** and a corresponding [log entry](#) is added to each To Do.

Complete Selected Entries

Use the **Complete** button to complete the selected entries.

Upon successful processing, the selected entries become **Completed** and a corresponding [log entry](#) is added to each To Do.

FASTPATH: Refer to [The Lifecycle Of A To Do Entry](#) for more information.

How To

The following sections describe how to perform common tasks for To Dos.

How To Assign a To Do to a User

There are several ways to assign a To Do entry to a user:

- A supervisor can use the [Supervisor Functions](#) pages to review, assign and reassign To Dos to users.
 - The [Supervisor To Do Summary](#) page summarizes by type the total, **Open** (unassigned) and **Being Worked On** (assigned) entries for roles to which you belong.
 - The **Total**, **Open** or **Being Worked On** column hyperlinks transfers you to the [Supervisor To Do Assignment](#) page where you can assign / reassign one or more of the corresponding entries to a specific user.
- Use the [To Do List Functions](#) pages to review and assign To Dos to yourself.
 - The [To Do Summary](#) page summarizes by type the total, **Open** (unassigned) and **Being Worked On** (assigned to you) entries for roles to which you belong.
 - The **Open** column hyperlink transfers you to the [To Do List](#) page where you can assign one or more of the **Open** entries to yourself.
- You may also use the [To Do Search](#) page to find certain To Do entries based on various criteria and assign selected entries to a user.
- Forwarding an entry to a user using the [To Do Entry](#) page is yet another way of assigning an entry to a user.
- The [Current To Do](#) zone next assignment functionality can be configured to dynamically assign an **Open** To Do to the requesting user based on your organization's work assignment rules. Refer to this zone for more information.

How To Work on To Do Entries


Depending on your organization's [work distribution](#) practice, there are several ways to work on To Do entries.

When work is manually assigned, a typical approach may be to review the entries in a list and work them one entry at a time.

- You can use the [To Do List](#) page to review and work on To Do entries by list type.
- You may also use the [To Do Search](#) page to find certain To Do entries based on various criteria and work on selected entries. For example, you may want to review and work on entries related to a given To Do.

Regardless of the above methods to define your work list, when ready to work on a specific entry, drill into its respective page using the message description hyperlink. Drilling into an entry from the above work lists makes it your current To Do entry. This entry's key information is displayed on the [Current To Do](#) dashboard zone and you may use special action bar buttons to scroll through your work list to the next and previous entries respective to your current To Do.

FASTPATH: Refer to [Working On A To Do Entry](#) for more information.

When work is assigned "on demand", click the Current To Do icon,  , on the Current To Do dashboard zone to request your next task. The returned To Do entry automatically becomes your current To Do. The entry's key information

is displayed on the zone but since the entry was not selected from a work list the next and previous action bar buttons are disabled.

FASTPATH: Refer to [Current To Do](#) zone to learn about all the functions it provides to efficiently resolve the current To Do and potentially other To Do entries that may be [related](#) to it.

Resolve the current To Do by fixing the problem indicated in the message and mark the entry as **Complete**.

FASTPATH: Refer to [How To Complete To Dos](#) for more information.

How To Add a New To Do

If your system is configured with [manual To Do types](#), you can add a To Do entry online. Online To Do entries may be used like electronic help tickets in the system. In Oracle Utilities Customer Care and Billing, for example, if you are having a problem starting service, you can create a To Do that describes the problem. The To Do can be assigned to a help resolution group that could either resolve the problem or send the To Do back to you with information describing how to resolve the problem.

FASTPATH: Refer to [To Do Entry - Add Dialog](#) for more information.

How To Complete To Dos

When you complete a To Do, a log entry is added that indicates when the To Do was completed and by whom.

There are multiple ways in which you can complete a To Do.

- The [To Do List](#) page allows you to complete one or more selected entries of a specific To Do Type.
- The [To Do Search](#) page allows you to search for To Do entries using various search criteria and complete selected ones.
- Use the [To Do Entry](#) page to complete a single To Do Entry.
- You may also use the [Current To Do](#) zone to complete the current entry you are working on and/or all [related](#) entries assigned to you.

NOTE: There are also situations where To Do entry is completed automatically by the system. Refer to [The Lifecycle of a To Do Entry](#) for more information.

How To Forward a To Do

If you have been assigned a To Do that you are unable to resolve, you can forward it to another user or role for help or additional instructions. When a To Do is forwarded, a log entry indicates who forwarded the To Do, to whom it was forwarded, and the specified reason.

- Use the **Forward** button on the [To Do Entry](#) page to forward a To Do entry to a role or a user.
- You may also use the **Reopen** button on the [To Do Search](#) page to forward selected entries back to their current role.

How To Send Back a To Do

If you have been forwarded a To Do by another user, you may send back the To Do with instructions to the forwarding user. When a To Do is sent back, a log entry is created that indicates when the To Do was sent back and by whom.

Use the **Send Back** button on the [To Do Entry](#) page to send back a To Do entry.

How To Add a Log Entry to a To Do

You can use the **Comments** section of a [To Do entry](#) to add a comment to the To Do as long as the To Do is **Open** or **Being Worked On**. However, the Comments field is disabled if the To Do entry is complete.

Alternatively, you can add a **User Details** log entry to a To Do at any time (even after it's completed). Refer to [To Do Entry](#) page for more information.

Reports

The following section describes the functionality provided for submitting reports on-line and for viewing reports submitted through a third party reporting tool.

FASTPATH: For information about this functionality, see [The Big Picture of Reports](#).

On-line Report Submission

The on-line report submission page enables you to request a specific report to be run. This is only available if your system is able to integrate with your third party reporting system on-line. The system comes configured to integrate with BI Publisher.

Report Submission - Main

This page allows you to submit a report on-line. Navigate to this page using **Menu > Tools > Report Submission..**

Description of Page

To submit a report, choose the **Report Code** for the process you wish to submit. Refer to [Defining Report Definitions](#) for information about defining reports in the system.

A description of the report is displayed.

The parameters defined for the report on the [report definition](#) page are displayed in the grid. To see a description of a parameter, use the search button adjacent to the parameter label.

NOTE: Default note. The default values displayed for the parameters are defined on the [report definition](#) page.

When you have entered the appropriate parameter values, click the **Submit** button to submit the report. Refer to [How To Request Reports](#) for more information.

NOTE: Reports submitted via the on-line report submission page are not logged in the report history repository and are therefore not visible from the [report history](#) page. Only reports submitted through your third party reporting tool's batch scheduler are logged in the report history, if supported by the tool.

Viewing Historic Reports

If your third party reporting tool stores historic copies of reports in a repository, you may configure the system to allow a user to view these historic reports on-line.

NOTE: This page currently does not display historical reports for BI Publisher.

Report History - Main

This page allows you to view a report stored in your third party reporting tool's report repository, if supported. Navigate to this page using **Menu > Tools > Report History**.

Description of Page

The top part of the page shows the **Report Code**, **Description** and **Long Description** for the selected report.

To select a report, search by either **Report Code** or **Description** and click the adjacent search button. The reports matching your search criteria are displayed. Choose the desired report from the selection list.

NOTE: Displaying Report History. When you choose a report, the system needs to interface with your reporting software to display a list of historical occurrences of this report.
