



DeepSee End User Guide

Version 2017.2
2020-06-25

DeepSee End User Guide
Caché Version 2017.2 2020-06-25
Copyright © 2020 InterSystems Corporation
All rights reserved.

InterSystems, InterSystems IRIS, InterSystems Caché, InterSystems Ensemble, and InterSystems HealthShare are registered trademarks of InterSystems Corporation.

All other brand or product names used herein are trademarks or registered trademarks of their respective companies or organizations.

This document contains trade secret and confidential information which is the property of InterSystems Corporation, One Memorial Drive, Cambridge, MA 02142, or its affiliates, and is furnished for the sole purpose of the operation and maintenance of the products of InterSystems Corporation. No part of this publication is to be used for any other purpose, and this publication is not to be reproduced, copied, disclosed, transmitted, stored in a retrieval system or translated into any human or computer language, in any form, by any means, in whole or in part, without the express prior written consent of InterSystems Corporation.

The copying, use and disposition of this document and the software programs described herein is prohibited except to the limited extent set forth in the standard software license agreement(s) of InterSystems Corporation covering such programs and related documentation. InterSystems Corporation makes no representations and warranties concerning such software programs other than those set forth in such standard software license agreement(s). In addition, the liability of InterSystems Corporation for any losses or damages relating to or arising out of the use of such software programs is limited in the manner set forth in such standard software license agreement(s).

THE FOREGOING IS A GENERAL SUMMARY OF THE RESTRICTIONS AND LIMITATIONS IMPOSED BY INTERSYSTEMS CORPORATION ON THE USE OF, AND LIABILITY ARISING FROM, ITS COMPUTER SOFTWARE. FOR COMPLETE INFORMATION REFERENCE SHOULD BE MADE TO THE STANDARD SOFTWARE LICENSE AGREEMENT(S) OF INTERSYSTEMS CORPORATION, COPIES OF WHICH WILL BE MADE AVAILABLE UPON REQUEST.

InterSystems Corporation disclaims responsibility for errors which may appear in this document, and it reserves the right, in its sole discretion and without notice, to make substitutions and modifications in the products and practices described in this document.

For Support questions about any InterSystems products, contact:

InterSystems Worldwide Response Center (WRC)
Tel: +1-617-621-0700
Tel: +44 (0) 844 854 2917
Email: support@InterSystems.com

Table of Contents

About This Book	1
1 Introduction to the User Portal	3
1.1 Purpose of the User Portal	3
1.2 Logging On to DeepSee	4
1.2.1 Selecting a Preferred Language	4
1.3 Introduction to the User Portal	4
1.4 About Pivot Tables and Dashboards	5
1.5 Summary of Options in the User Portal	7
1.6 Security	8
1.6.1 Security for Functionality	8
1.6.2 Security for Pivot Tables and Dashboards	8
1.6.3 Security for Model Elements	8
2 Using the Portal Features	9
2.1 Displaying Folder Items	9
2.1.1 Covers View	9
2.1.2 List View	11
2.1.3 Options Available in the Main Area of the User Portal	11
2.2 Visibility of Folders and Folder Items	12
2.3 Customizing Covers	13
2.3.1 Cover Parts	15
2.3.2 Cover Options	16
2.4 Customizing the Covers View	17
2.5 Using Worklists in the User Portal	18
2.6 Hiding or Showing the Worklists	20
2.7 Registering or Unregistering for Alerts	20
2.8 Sending an Alert	21
2.9 Viewing Alerts	21
2.10 Sending Email	22
2.10.1 How DeepSee Supports Email	23
2.11 Managing Workflow Tasks	23
2.12 Creating Links to Web Pages and Zen Reports	24
3 Orientation to Dashboards	27
3.1 Overview of Dashboards	27
3.2 Widget Variations	29
3.2.1 Pivot Table Widgets	29
3.2.2 Scorecard Widgets	30
3.2.3 Meter Widgets	30
3.2.4 Map Widgets	31
3.2.5 Calendar Widgets	31
3.2.6 Controls Widgets	32
3.2.7 Custom Widgets	32
3.3 Data Sources for Widgets	32
3.4 Widget Basics	33
3.5 Locations of Controls and Buttons	33
3.6 Empty Widgets	34

4 Using Dashboards	35
4.1 Toggling the Display of the Worklists	35
4.2 Managing the Widget Windows	36
4.2.1 Window Buttons	36
4.2.2 Resizing Widgets	36
4.2.3 Moving Widgets	36
4.3 Filtering Data	37
4.3.1 Using a Filter Control	37
4.3.2 Filter Control for a Time Level	38
4.3.3 Filter Control for a Day Level	40
4.3.4 Filter Control for a Measure	41
4.3.5 Filter Control for a Computed Dimension	41
4.3.6 Filtering via OnClick	42
4.3.7 Data Source Differences	43
4.4 Understanding the Effects of Filters	43
4.5 Sorting Data	45
4.6 Specifying the Row or Column Count	45
4.7 Changing the Display Type	46
4.8 Using the Chart Zoom Feature	47
4.9 Using Combo Charts	47
4.10 Displaying Listings	49
4.10.1 Map Listing Variation	49
4.10.2 Displaying iKnow Measure Values	50
4.11 Drilling Down	52
4.11.1 Drilling Down with Double-Click	52
4.11.2 Using the Drilldown Button	53
4.12 Using the Dimension List	54
4.13 Drilling to a Different Dashboard	55
4.14 Refreshing a Widget	56
4.15 Reloading a Dashboard	56
4.16 Exporting Data to Microsoft Excel	56
4.17 Printing Data	57
4.17.1 Requirements for Printing	57
4.18 Using a Map Widget	57
4.19 Other Tasks	58
4.20 Saving Your Changes	59
5 Using the Pivot Analysis Window	61
5.1 Using the Pivot Analysis Window	61
5.2 Cluster Analysis	62
5.3 Distribution Analysis	63
5.4 Regression Analysis	64
5.5 iKnow Content Analysis	66
5.6 iKnow Entity Analysis	67
5.6.1 Overview Tab	68
5.6.2 Cell Breakdown Tab	70
5.6.3 Entities Tab	72

About This Book

This book describes, to end users and to developers, how to use the DeepSee User Portal and dashboards. It includes the following chapters:

- [Introduction to the User Portal](#)
- [Using Portal Features](#)
- [Orientation to Dashboards](#)
- [Using Dashboards](#)
- [Using the Pivot Analysis Window](#)

For a detailed outline, see the [table of contents](#).

The other user books for DeepSee are as follows:

- [Creating DeepSee Dashboards](#) describes how to create and modify dashboards in DeepSee.
- [Using the DeepSee Analyzer](#) describes how to create and modify pivot tables, as well as use the Analyzer in general.

The following books are primarily for developers who work with DeepSee:

- [Getting Started with DeepSee](#) briefly introduces DeepSee and the tools that it provides.
- [DeepSee Developer Tutorial](#) guides developers through the process of creating a sample that consists of a cube, subject areas, pivot tables, and dashboards.
- [DeepSee Implementation Guide](#) describes how to implement DeepSee, apart from creating the model.
- [Defining DeepSee Models](#) describes how to define the basic elements used in DeepSee queries: cubes and subject areas. It also describes how to define listing groups.
- [Advanced DeepSee Modeling Guide](#) describes how to use the more advanced and less common DeepSee modeling features: computed dimensions, unstructured data in cubes, compound cubes, cube relationships, term lists, quality measures, KPIs, plugins, and other special options.
- [Using MDX with DeepSee](#) introduces MDX and describes how to write MDX queries manually for use with DeepSee cubes.
- [DeepSee MDX Reference](#) provides reference information on MDX as supported by DeepSee.
- [Tools for Creating DeepSee Web Clients](#) provides information on the DeepSee JavaScript and REST APIs, which you can use to create web clients for your DeepSee applications.

Also see the article [Using PMML Models in Caché](#).

For general information, see the [InterSystems Documentation Guide](#).

1

Introduction to the User Portal

This chapter introduces the User Portal. It discusses the following topics:

- [Purpose of the User Portal](#)
- [How to log on to the DeepSee User Portal](#)
- [A first look at the User Portal](#)
- [Introduction to dashboards and pivot tables](#)
- [Summary of the tasks you can perform in the User Portal](#)
- [Security in DeepSee](#)

Be sure to consult the online [InterSystems Supported Platforms](#) document for this release for information on system requirements for DeepSee.

1.1 Purpose of the User Portal

The User Portal is intended for direct use by end users (in contrast to such back end tools as Studio and the Management Portal). The User Portal is designed to enable you to do the following tasks:

- View and use dashboards and pivot tables, all containing key information. For Ensemble users, this information can include Ensemble business metrics. For DeepSee users, this information presents aggregated views of your data that you can investigate in various ways.

You can organize these items into folders for easier management. Also, you can attach keywords to these items, so that you can find items again more quickly.

- Share your dashboards and pivot tables with other users.

The User Portal is a shared environment. That is, the same items are visible to all users who have access to the User Portal for a given Caché namespace.

- Create dashboards, which generally display pivot tables and other elements such as meters and scorecards. Dashboards are typically intended to address specific business needs.
- Access the DeepSee Analyzer, in which you can create pivot tables or perform ad hoc analysis of data.
- Share and access links to web pages and Zen reports.
- (For Ensemble workflow users) Manage your workflow tasks.

1.2 Logging On to DeepSee

In a supported browser, go to the following URL:

`http://localhost:57772/csp/samples/_DeepSee.UserPortal.Home.zen`

Where `localhost:57772` is the server and port used by the Caché web server and `samples` is the Caché namespace that you are using. When prompted, enter your Caché username and password.

1.2.1 Selecting a Preferred Language

This user interface can be displayed in different languages. Initially, the default preferred language (that is, the session language) is the language specified for the browser, or English if the browser language is not supported.

To select your preferred language, select **About** and then select an option in the list **Preferred language for this session**.

Note that after you select a language here, the system continues to use that language even if you change the language used by the browser.

1.3 Introduction to the User Portal

In the SAMPLES namespace, the User Portal looks like the following:

From	Subject	Date
Olivia Randolph-Erickson	Sample Alert	Today at 14:58:08
NBAKER	PLEASE REVIEW THIS ASAP	Today at 14:58:08
L. Jones	See this recent activity	Today at 14:58:08
NBAKER	NEED YOUR INPUT HERE	Today at 14:58:08
Max Wilson	Sample low-priority alert	Today at 14:58:08
Peter Wang	Sample Alert	Today at 14:58:08
Max Wilson	Sample high-priority	Today at 14:58:08

Name	Type
Basic Dashboard Demo	DASHI
Drill Options	DASHI
Dynamic Dashboard	DASHI
HoleFoods Compound Cube	DASHI
Listing with Filters	DASHI
Listing with OnClick Filter	DASHI
Patients Compound Cube	DASHI
Real Time Updates	DASHI
Two Subject Areas Together	DASHI
User Defined Listing	DASHI
Calculated Members	Fold
Alternative Avg Allergy Count	PIVOT
Average Patient Count per Decade	PIVOT
Primary Color Example	PIVOT
Unique Doctor Count	PIVOT

Notice the following elements in this page:

- The left area contains worklists, which are rectangular areas that contain tools. For details, see [“Using Worklists in the User Portal.”](#)

In SAMPLES, these areas are initially blank.

- The main area on the right shows the public DeepSee pivot tables, dashboards, and other items.

Here you can open items, delete items, and mark items as favorites. You can also expand or collapse the folders. The top area provides a search option.

For a dashboard, this area displays the title of the dashboard, if it has a title; otherwise, this area displays the name of the dashboard.

This area does not display items that are marked as private.

If you are an Ensemble workflow user, the main area also displays an item labeled **Workflow Inbox**.

- The **Menu** link in the upper left contains options to create and open dashboards and pivot tables and perform other tasks in DeepSee.

If you are assigned to the `%Developer` role, this menu includes an option to access the Management Portal.

- The **Home** link at the top displays this page.
- The **Logout** link enables you to log out.

The User Portal is a shared environment. Many of the items you see are also visible to other users.

In the ENSDEMO namespace (in Ensemble installations), the User Portal displays a different set of sample dashboards. These dashboards display Ensemble business metrics but look like the other dashboards discussed in this book.

1.4 About Pivot Tables and Dashboards

Pivot tables are central to InterSystems DeepSee; they select and aggregate data. The following figure shows an example pivot table that shows the number of patients and the average allergy count per patient, grouped by age and gender.

Age Group		Female		Male	
		Patient Count	Avg Allergy Count	Patient Count	Avg Allergy Count
0 to 29	0 to 9	680	0.60	750	0.63
	10 to 19	756	0.66	769	0.69
	20 to 29	661	0.64	648	0.61
30 to 59	30 to 39	815	0.63	735	0.65
	40 to 49	728	0.68	741	0.61
	50 to 59	586	0.58	552	0.62
60+	60 to 69	397	0.64	319	0.68
	70 to 79	304	0.58	242	0.56
	80+	217	0.57	100	0.66

Using this example as a reference, let us discuss the key terms in DeepSee:


- A *level* is used to group the source data. A level has *members*. Each member, in turn, corresponds to a specific set of records in the source data.

For example, the `Age Group` level has the members `0-29`, `30-39`, and `60+`. The `Gender` level has the members `Female` and `Male`.

- A *measure* is a value displayed in the body of the pivot table; it is based on values in the source data. For a given context, a measure aggregates the values for all applicable source records and represents them with a single value. For example, the measure `Patient Count` is the number of patients, and the measure `Avg Allergy Count` is the average number of allergies per patient.

The preceding example shows a preview of a pivot table within the Analyzer, which provides a large set of tools for modifying the pivot table, exporting to Microsoft Excel, toggling to a chart format, and so on.

You typically see a pivot table within a *dashboard*, as follows:




Notice that this dashboard includes multiple filters (shown in the **Filters** box), which you can use to filter the items on this dashboard.

Dashboards can include other kinds of elements, including *meters* and *scorecards*. A scorecard shows multiple rows of data in a tabular format that also includes features such as value-dependent lamps and arrows. For example:



A meter shows one or more values in graphical objects like the following:



The next chapter discusses dashboards in more detail; see “[Using Dashboards.](#)”

1.5 Summary of Options in the User Portal

In the User Portal, you can do the following tasks, depending upon your permissions:

- Send alerts to other users.

Note: These alerts are unrelated to Ensemble alerts.

- Review alerts that you have received.
- Send email to other users (if your system is configured for this).
- Manage Ensemble workflow tasks (only for Ensemble workflow users).
- Mark items as favorites, for easy access.
- Create, modify, view, and work with dashboards.

For information on creating and modifying dashboards, see [Creating DeepSee Dashboards](#).

- Create, modify, view, and work with pivot tables.

For information on creating and modifying pivot tables, see [Using the DeepSee Analyzer](#).

- Add links to web pages and to Zen reports.

- Delete public items displayed in the User Portal.
- Display the Analyzer and perform additional analysis there. See *Using the DeepSee Analyzer*.

Except where noted, these topics are discussed in this book.

1.6 Security

DeepSee has a formal mechanism for managing access to DeepSee functionality, DeepSee items, and your data. This mechanism is based on the underlying Caché security framework. This section summarizes how this mechanism affects you.

For details, see the *DeepSee Implementation Guide*.

1.6.1 Security for Functionality

The User Portal and the Analyzer are secured as separate entities. You may or may not have read and write access to both these components. For example, you might have both read and write access for the User Portal (which would let you view and create dashboards), but only read access for the Analyzer (which would let you change a pivot table but not save the changes).

1.6.2 Security for Pivot Tables and Dashboards

Pivot tables, dashboards, and other items displayed in the User Portal are known as folder items and can be secured via Caché resources. This works as follows:

- An administrator performs the following setup:
 - Creating resources in the Management Portal.
 - Creating roles in the Management Portal. These roles can have either READ permission on a resource or both READ and WRITE permissions on a resource.
 - Assigning users to roles, which enables them to read (or read and modify) the items or items associated with a resource.

The administrator should then notify the DeepSee users about the available resources and how they are to be used.

- When you save a folder item, you can type the name of the applicable resource into the **Access Resource** field.

1.6.3 Security for Model Elements

DeepSee model elements can also be secured via Caché resources. This works as follows:

- An administrator creates resources and roles, and assigns users to roles as described in the previous section.
- Developers specify the resource (if any) that secures each cube, subject area, and KPI (key performance indicator).

2

Using the Portal Features

This chapter describes how to use the basic features of the User Portal. It discusses the following topics:

- [Ways to display folder items](#)
- [Visibility of folders and folder items](#)
- [How to customize the covers](#)
- [How to customize Covers View](#)
- [How to use the worklists](#)
- [How to hide or show the worklists](#)
- [How to register or unregister for alerts](#)
- [How to send an alert](#)
- [How to view an alert](#)
- [How to send email](#)
- [How to manage Ensemble workflow tasks](#)
- [How to add links to web pages and Zen reports](#)

Note: Alerts in the User Portal are unrelated to Ensemble alerts.

2.1 Displaying Folder Items


The main area of the User Portal displays the dashboards and pivot tables in this Caché namespace, with different icons for dashboards and pivot tables. This item can also display links to web pages and links to Zen reports. All these items are known as *folder items*, because they can be grouped into folders (which are used primarily for security).

At the top of this area, use the Covers and List options to display [Covers View](#) and [List View](#), respectively.


You can also use **Menu > Find Dashboard...** to open a dashboard.

2.1.1 Covers View

The following shows the Covers View in the SAMPLES namespace:




When you hover the cursor over a cover, DeepSee displays an arrow on the cover, as follows:



If you select the cover arrow, DeepSee displays the dashboard or other item.

If you select anywhere except for the arrow, DeepSee displays an informational dialog box like the following:



To close this dialog box, select the X in the upper right or select the screen outside of this dialog box.

2.1.2 List View

The following shows the List View in the SAMPLES namespace:

Name	Type	Keywords	
Basic Dashboard Demo	DASHBOARD	Patients	★ X
Drill Options	DASHBOARD		★ X
Dynamic Dashboard	DASHBOARD	Patients	★ X
HoleFoods Compound Cube	DASHBOARD	HoleFoods	★ X
Listing with Filters	DASHBOARD	Patients	★ X
Listing with OnClick Filter	DASHBOARD	HoleFoods	★ X
Patients Compound Cube	DASHBOARD	Patients	★ X
Real Time Updates	DASHBOARD	Patients, KPIs	★ X
Two Subject Areas Together	DASHBOARD	Patients	★ X

This view groups folder items by the folders to which they belong.

2.1.3 Options Available in the Main Area of the User Portal

In the main area of the User Portal (the right area), you can do the following:

- (Only in [List View](#)) Expand or collapse a folder. To do so, select the triangle to the left of the folder. Covers View does not display folders.

- View an existing folder item. To do so, select the cover arrow (in [Covers View](#)) or select the name of the item (in [List View](#)).

Then, for dashboards, see the chapter “[Using Dashboards](#)” for information on your options. If you have the appropriate permission, you can then edit the dashboard. See [Creating DeepSee Dashboards](#).

Or, for pivot tables, if you have the appropriate permission, you can then edit the pivot table. See [Using the DeepSee Analyzer](#).

- Search for an item. To do so, type your search text into the box at the top of the right area and then select **Go**. The system finds all items whose names or keywords include the given text; the search is not case-sensitive.
- Display different groups of folder items.

Show: **All** Personal Shared Public

For details, see the [next section](#).

- Delete a folder item. In [Covers View](#), select the cover and then select the X button. In [List View](#), select the X button in the row for that item.

In either case, the system prompts you to confirm this action.

Note: When you delete a folder item, DeepSee moves it into a folder named \$TRASH. Or, if the item is already in a folder, then \$TRASH/ is prepended to the folder name. In either case, the item is no longer visible in the web-based interfaces. If you have access to Studio, you can restore it. For details, see “[Restoring Deleted Folder Items](#)” in the chapter “[Packaging DeepSee Elements into Classes](#)” in the *DeepSee Implementation Guide*.

- Mark a folder item as a favorite. In [Covers View](#), select the cover and then select the Favorites button . In [List View](#), select the Favorites button in the row for that item.

2.2 Visibility of Folders and Folder Items

The User Portal lists only a subset of the folders and folder items. You can use the **Show** options to display different groups of folder items.

Show: **All** Personal Shared Public

The options are as follows:

- All** — Shows all public folders and items. This is controlled by the **Public** setting of an item.

Important: Note that the **Public** setting controls only whether the item is displayed in the User Portal. It does not otherwise control access.
- Personal** — Shows all public items that you own. This is controlled by the **Owner** setting of an item. Note that, by default, the owner of an item is the Caché user who created it.
- Shared** — Shows all public items that you do not own and that have assigned resources to which you have permission. This is controlled by the **Resource** setting of an item.
- Public** — Shows all public items that you do not own and do not have assigned resources.

The following table lists all the possible variations and indicates the visibility of the folders and folder items, depending on what you choose for **Show**:

Marked as Public?	Owner	Resource	User Portal Displays Item to You?
Yes	You	Any	Yes (if you select All or Personal)
	Someone else	No resource	Yes (if you select All or Public)
		Resource that is available to you	Yes (if you select All or Shared)
		Other resource	No
No	Any	Any	No (but the item might be available via the Menu , depending on its owner and resource)


Similarly, to define an item so that only you can see it, do the following:

- Set yourself as the **Owner**.
- Specify a **Resource** that is available only to you.
- Optionally select **Public** to display the item in the User Portal.

For background information, see “[Security in DeepSee](#),” in the previous chapter.


2.3 Customizing Covers

When you display the [Covers View](#) in the User Portal, DeepSee displays a rectangular cover for each item, and there is a default cover for each type of item. For example, the default cover for a dashboard looks like the following example:



To customize a cover:


- If you are currently viewing the User Portal, first select the **View Covers** option. Then select the cover (but do *not* select the arrow that is temporarily shown in its bottom area). DeepSee displays an information dialog box:



Then select the Cover Design button .

- (For a dashboard) If you are currently viewing the dashboard, select **Menu > Design Book Cover**.

In either case, DeepSee displays the Book Cover Editor, which looks like this:



Note the following:

- To customize part of a cover, either select a tab at the top or select an area in the cover preview. In either case, the left area then displays a list of options.
- To control the size of the cover preview, use the **1x**, **2x**, and **3x** buttons at the right.
- To undo a change, select **Undo** (not shown here).
- To clear all customizations and restore the default cover, select **Clear**.
- To close the editor without saving changes, select **Cancel**.
- To save changes and close the editor, select **Save**.
- To upload an image file, select **Upload Image**.

DeepSee copies the image files to the directory *install-dir/CSP/broker/covers*. The image files are available in all Caché namespaces.

- To hide the list of options on the left, select the **<** button.

2.3.1 Cover Parts

A cover has the following parts that you can customize separately:

- **Background** — The entire area within the rectangle, including the areas behind the header and footer (which can be hidden).
- **Header** — A rectangle at the top. By default, this is shown in a contrasting color and displays the type of the item (dashboard, pivot, and so on).

- **Title** — By default, this is displayed below the header and displays the name of the item (dashboard, pivot, and so on). For a dashboard, this displays the dashboard title, if any; otherwise it displays the dashboard name.
- **Image** — An image included on the cover. By default, this image is in the middle of the page. DeepSee uses a default image for each type of folder item.
- **Subtitle** — By default, this is blank. The default position is below the image.
- **Footer** — A rectangle at the bottom. By default, this is shown in a contrasting color and displays the owner of the item.

2.3.2 Cover Options


The following table summarizes the options available for each part:

Option	Notes	Available for Background?	Available for Header, Title, Subtitle, and Footer?	Available for Image?
Visible	Optionally clear this option to hide this element	yes	yes	yes
Text	Type a string or select a token to represent specific text from the selected item	no	yes	yes
Color	Select a color to use for the text	no	yes	yes
Image	Select an image or type the URL of an image	yes	no	yes
Background	Select a color for the fill pattern for this element	yes	yes	yes
Make Image Full Size	Select this option to make the image fill the background, if it does not already do so	yes	no	no
Font	Select a typeface	no	yes	yes
Font Size	Use the slider to specify the font size or select the check box to reset to the default	no	yes	yes
Text Style	Select a text style	no	yes	yes
Align	Select an option to make this element left-aligned, centered, or right-aligned	no	yes	yes
Top	Use the slider to specify the top position of this element or select the check box to reset to the default	yes	yes	yes

Option	Notes	Available for Background?	Available for Header, Title, Subtitle, and Footer?	Available for Image?
Left	Use the slider to specify the left position of this element or select the check box to reset to the default	yes	yes	yes
Height	Use the slider to specify the height of this element or select the check box to reset to the default	yes	yes	yes
Width	Use the slider to specify the width of this element or select the check box to reset to the default	yes	yes	yes
Margin	Use the slider to specify the margin around this element	yes	yes	yes
Padding	Use the slider to specify the padding around this element	yes	yes	yes
Opacity	Use the slider to specify the opacity of this element	yes	yes	yes
Radius	Use the slider to specify the radius of the corners of this element	yes	yes	yes

2.4 Customizing the Covers View

You can customize the font and colors used in the background of the [Covers View](#) in the User Portal. For example:



To do so:

1. Select the < button at the left.
It is not necessary to be in [Covers View](#).
2. Select **Background**.
3. Specify the following options:
 - **Color** — Specify the font color for labels in the background (such as the default category name, **All**, shown here).
 - **Background** — Specify the color for the background.
 - **Font** — Specify the typeface for labels.
 - **Font Size** — Specify the font size for labels.
 - **Text Style** — Specify the text style for labels.

Or select **Reset Styles** to remove all customizations.

These changes are visible to all users in this namespace.

2.5 Using Worklists in the User Portal

The home page of the User Portal always includes two *worklist areas* on the left, and these display different *worklists* depending on your site configuration. For any worklist area, the upper right corner might display icons to indicate which worklists it can display. For example:




If no icons are shown, this worklist area can display only one worklist.

The highlighted icon indicates which worklist is currently displayed in the worklist area. You can select a different icon to display the corresponding worklist in this area instead.

For the User Portal, the available worklists are as follows:


- The **Alerts** worklist displays recent alerts sent to you. For example:



From	Subject	Date	
Keith Madison	for your demo to management	Today at 10:03:36	X
Olivia Randolph-Erickson	Sample Alert	Today at 10:03:57	X
NBAKER	PLEASE REVIEW THIS ASAP	Today at 10:03:57	X
	See this recent	Today at	X

For information, see “[Viewing Alerts.](#)”


- The **Details** worklist displays details for a pivot table or dashboard. For example:



Details
Type
Dashboard
Name
Demo Real Time Updates
Created on
Today at 04:16:45
Last modified
Today at 04:16:45


To display details for an item, select the icon to the left of the item name in the main area of the User Portal. Then the system refreshes the **Details** worklist.

- The **Favorites** worklist displays any items that you have marked as favorites. For example:



Favorites	
Basic Dashboard Demo Today at 04:16:45	X
Sales Against Targets Today at 04:16:48	X

- The **Recent items** worklist displays items that you have recently accessed, in descending order by the date you accessed them. For example:




Recent items	
Patients Compound Cube Today at 10:28:09	X
Demo Trend Lines Today at 10:27:05	X
Basic Dashboard Demo Today at 10:25:29	X

In all cases, to open an item, select it. To remove an item from a worklist, select the X button.

2.6 Hiding or Showing the Worklists

The upper right corner of the main area displays a << button that you can use to toggle the display of worklists:



Select the << button to hide the worklists. Then select the >> button to display them again if needed.

2.7 Registering or Unregistering for Alerts

To change whether you are registered for alerts:

1. Access the User Portal home page.
2. Select **Menu > Register for Alerts**.

The system displays a dialog box like the following:

 A screenshot of the 'Alert Registration' dialog box. The title bar is blue with a circular logo and the text 'Alert Registration' and 'Register for dashboard alerts'. Below the title bar, there are three input fields: 'User Id' with the value 'DeepSeeDeveloper', 'Full name' with the value 'DeepSee Developer', and a note 'Name used for alerts you send'. At the bottom, there is a checked checkbox labeled 'Alerts enabled'.

3. Specify the following options:
 - **Full name** — Type the full name that you want other users to see when they send alerts.
 - **Alerts enabled** — Select this option to receive alerts.
4. Select **OK**.

2.8 Sending an Alert

You can send alerts to other DeepSee users, if those users have registered for alerts. To do so:

1. Open a dashboard that contains the data about which you want to send the alert.
2. Select **Menu > Send Alert**.

The system displays a dialog box like the following:


3. Specify the following options:
 - **Send to** — Select each user who should receive this alert.
 - **Subject** — Optionally type a short message to display in the Alerts worklists of the other users.
 - **Priority** — Optionally select a priority; this affects how the alert is displayed. The default is Normal priority.
 - **Comments** — Optionally type a longer message.
4. Select **OK**.

2.9 Viewing Alerts

To view your alerts:


1. In a worklist area, select the Alerts button .

You might see something like this:



Note the following items:

- A star on the left indicates that you have not yet opened this alert.
 - High-priority items are shown in red font, normal-priority items are shown in black font, and low-priority items are shown in gray font.
2. Select anywhere in the row for an alert to open it. The system then opens the associated dashboard and displays the alert details in the **Alert Details** worklist, as follows:



2.10 Sending Email

From within a dashboard, you can send email that contains a link to the dashboard, if your system is configured for this. To do so:

1. Open a dashboard that contains the data about which you want to send the alert.
2. Select **Menu > Send Email**.

If this option is not present, then your system is not configured to support email.

3. The system then does one of the following:

- Displays the default email client on your machine, with a message that contains a link to the dashboard.
Edit the message, specify one or recipients, and send it as usual.
- Displays a dialog box.

Type the email address of the recipient, type a brief message to include along with the generated link, and select **OK**.

If you have previously sent email, you can choose a recipient from the **Send to** drop-down list.

2.10.1 How DeepSee Supports Email

You can be configured to support email in either of two ways:

- It can use *client-side email*. In this configuration, when you send email, DeepSee accesses the default client-side email program such as Microsoft Outlook. You use this to send email in the usual way, and you have access to your address books.
- It can use *server-side email*. In this configuration, when you send email, you must type the email address of the recipient. DeepSee routes the message to the configured email server.

It can also be configured not to support email. Consult your implementers, who should read the [DeepSee Implementation Guide](#).

2.11 Managing Workflow Tasks

If you are an Ensemble workflow user, the main area of the User Portal displays an item labeled **Workflow Inbox**.

You use this to manage your workflow tasks. If you select **Workflow Inbox**, the User Portal displays something like the following:

#	Priority	Subject	Message	Role	Assigned to	Time Created	Age
1		Test this problem from Mercy H Wellbeing	Can someone please call b	Demo-Testing	HDonovan	Today at 19:39:39	00w 0d 00h 52m 09s
2		Problem reported by Nelson Q Jefferson	Need some help!!!!	Demo-Development	HDonovan	Today at 19:39:22	00w 0d 00h 52m 25s
3	NEW	Problem reported by A Watson	Need some help!!!!	Demo-Development		Today at 18:54:21	00w 0d 01h 37m 27s

This list shows all unclaimed workflow tasks (marked as **NEW**) and all open tasks assigned to you.

To manage a workflow task, select it to see its details, which are then displayed in the lower area of the page:

+ Details for selected item

Relinquish Save Corrected Ignored

Subject	Priority	Assigned To	Time Created	Role
Problem reported by Mercy H Wellbeing		HDonovan	Today at 20:38:14	Demo-Development

Message
Can someone please call back about my broken window? Thx.

Comments

This example shows the built-in buttons and one custom editable field (**Comments**); you might see additional buttons and fields, depending on the task definition.

For a selected task, you can do any of the following:

- Take ownership of the task. To do so, select **Accept** (not shown in the previous figure, in which the task has already been accepted). This indicates that you intend do perform the associated work. The task becomes assigned and is no longer visible in the Workflow Inboxes of other users.
- Give up ownership of the task. To do so, select **Relinquish**. This causes the task to become unassigned and thus visible to all users in the same workflow role.

You can do this action and the following actions only if you own the task.

- Indicate that you completed the task. To do so, select **Corrected**.

The task is then removed from the list.

The workflow process might then generate additional tasks, depending on how it is defined.

- Mark the task as ignored. To do so, select **Ignored**.

The task is then removed from the list.

You use this option to indicate that you have encountered an error condition (such as not being able to contact the customer) and therefore the task cannot be completed.

- Perform other actions by selecting other buttons in this area. The preceding figure shows only the built-in action items.
- Enter values into editable fields. The preceding figure shows one custom editable field (**Comments**).

The purpose of these fields is to convey information to the next person who receives the task, to add information to the permanent record, or both. Therefore, be sure to provide meaningful entries.

2.12 Creating Links to Web Pages and Zen Reports

In the User Portal, you can also create links to web pages and to Zen reports. To do so:

1. Select the **+** button at the top.
Or select the **<** button at the left and then select **Add**.
2. Select **Add Link** or **Add Report Shortcut**, depending on what you want to add.
3. Specify the following options:
 - **Folder** — Specify the folder to contain this item.
Also see “[Preparing for Folder Item Localization](#)” in the *DeepSee Implementation Guide*.
 - **Name** — Specify the name of this item.

Also see “[Preparing for Folder Item Localization](#)” in the *DeepSee Implementation Guide*.

Note that this name is used (along with **Folder**) as the logical name of the item; DeepSee displays this logical name when you use the **Save As** option, for example.

- **Title** — Specify the title of the dashboard.

If you specify a title, the main area of the User Portal displays the title. Otherwise, it displays the name.

Also see “[Preparing for Folder Item Localization](#)” in the *DeepSee Implementation Guide*.

- **Link URL** (only for links to web pages) — Specify the URL.
- **Report Class** (only for links to Zen reports) — Select the report class.
- **Owner** — Specify the Caché user who owns this item. If an item has an owner, then only the owner can specify the **Access Resource** value for the item; see the next item.
- **Access Resource** — Specify the Caché resource that is used to control access to this item. See “[Visibility of Folders and Folder Items](#),” earlier in this chapter.
- **Public** — Select this option to control whether this item is displayed in the User Portal main area.


Important: This option does not control access to the item. Access is controlled instead by the **Owner** and **Access Resource** options, discussed earlier in this list.

- **Description** — Briefly describe the item.
- **Keywords** — List any keywords or phrases to help you find this item later. Specify a comma-separated list of keywords or phrases.

These keywords are displayed in the User Portal in [List View](#).

- **Category** — Specify the category to which this item belongs.

If you specify categories for items, [Covers View](#) groups items by category. For example:



3

Orientation to Dashboards


This chapter reviews the contents of the dashboards that you may already have in your system. It discusses the following topics:

- [Overview of dashboards, their layout, and their possible contents](#)
- [Widget variations you might see in your dashboards](#)
- [Data sources for widgets](#)
- [Widget basics](#)
- [Locations of controls and buttons](#)
- [Empty widgets](#)

This chapter is intended as an orientation to the possible contents of your dashboards. For information on using the elements described here, see the chapter “[Using Dashboards](#).”

3.1 Overview of Dashboards

The following example shows a sample dashboard:




A dashboard consists of the following areas:

- The upper left displays the name of the dashboard and (if defined) its title.
- Depending on the system configuration and on the individual layout of a dashboard, a dashboard can include zero, one, or two worklist areas on the left. For any worklist area, the upper right corner displays icons to indicate which worklists it can display. For example:



The highlighted icon indicates which worklist is currently displayed. You can select a non-highlighted icon to display the corresponding worklist in this area instead.

The **Filters** worklist is specific to the dashboard. To access this, select the Filters icon . For example:



- The right area contains one or more widgets. Each *widget* is a rectangular panel that displays data in some form.

3.2 Widget Variations

This section shows the contents of different types of widgets that you might see in your dashboards.


3.2.1 Pivot Table Widgets

A pivot table widget displays data in one of three formats:

- As a table with aggregated values. For example:

Age Group		Female		Male	
		Patient Count	Avg Allergy Count	Patient Count	Avg Allergy Count
0 to 29	0 to 9	680	0.60	750	0.63
	10 to 19	756	0.66	769	0.69
	20 to 29	661	0.64	648	0.61
30 to 59	30 to 39	815	0.63	735	0.65
	40 to 49	728	0.68	741	0.61
	50 to 59	586	0.58	552	0.62
60+	60 to 69	397	0.64	319	0.68
	70 to 79	304	0.58	242	0.56
	80+	217	0.57	100	0.66

- As a chart. For example:



- As a *listing*, which is a table of the lowest-level values. For example:

<input type="checkbox"/>	#	PatientID	Age	Gender	Home City	Test Score
<input type="checkbox"/>	1	SUBJ_100631	0	F	Elm Heights	50
<input type="checkbox"/>	2	SUBJ_100781	0	F	Redwood	78
<input type="checkbox"/>	3	SUBJ_100820	0	F	Magnolia	89
<input type="checkbox"/>	4	SUBJ_100966	0	F	Cypress	91
<input type="checkbox"/>	5	SUBJ_101274	0	F	Pine	
<input type="checkbox"/>	6	SUBJ_101340	0	F	Redwood	
<input type="checkbox"/>	7	SUBJ_101466	0	F	Magnolia	81
<input type="checkbox"/>	8	SUBJ_101532	0	F	Pine	67
<input type="checkbox"/>	9	SUBJ_101587	0	F	Elm Heights	77
<input type="checkbox"/>	10	SUBJ_102327	0	F	Redwood	79

If there are more than 100 rows, the bottom area displays buttons that you can use to page through all the rows.


3.2.2 Scorecard Widgets

A scorecard widget displays one or more rows of data in a tabular format that also includes features such as value-dependent lamps and arrows. For example:




3.2.3 Meter Widgets

A meter widget displays one or more values, each in a graphical object as follows:



The preceding picture shows a speedometer. DeepSee also supports text meters. For example:



3.2.4 Map Widgets

A dashboard can include a map widget like the following:



The highlighted points typically correspond to locations that are relevant to your business.

3.2.5 Calendar Widgets


A dashboard can include a calendar widget like the following:

«	Day	Week	Month	»		
May 2011						
Sun	Mon	Tue	Wed	Thu	Fri	Sat
1	2	3	4	5	6	7
8	9	10	11	12	13	14
15	16	17	18	19	20	21
22	23	24	25	26	27	28
29	30	31	1	2	3	4

This widget is included purely as information; it is not connected to your data.

3.2.6 Controls Widgets

A controls widget consists only of toolbar controls and has no main body content. The following shows an example:



3.2.7 Custom Widgets

A dashboard can also include custom widgets called *portlets*. The following shows an example:

My Widget		
Sales	UP	12%
Costs	DOWN	-8%
Profits	UP	18%

3.3 Data Sources for Widgets


Most widgets use a data source, which is one of the following:

- A pivot table. Pivot tables are created in the Analyzer. A pivot table is a query that is created by drag and drop actions. In some cases, you can launch the Mini Analyzer and make local changes to the pivot table. Changes do not affect other widgets, other dashboards, or other users. For details, see “[Other Tasks](#)” in the [next chapter](#).
- A KPI (key performance indicator). A KPI is a more advanced query created by a programmer.

All of these data sources can be displayed in pivot table widgets, scorecard widgets, and meter widgets.

3.4 Widget Basics

A widget might or might not include a title bar with buttons, as follows:




The title bar may or may not include a title.

Most widgets include scroll bars when necessary. A pivot table instead includes paging buttons like the following example:



If a pivot table displays nested rows and if there are more than 100 rows, DeepSee initially accesses the first 100 rows and displays the **Show All** button instead of these paging buttons. If you select the **Show All** button, DeepSee accesses all the rows and then displays the paging buttons.


Depending on how the dashboard is configured, the lower right corner of each widget can include a resize handle:



3.5 Locations of Controls and Buttons

A dashboard can include controls and graphical buttons. Depending upon the dashboard design, you see these items in either or both of the following locations:

- In the **Filters** worklist of the dashboard. For example:



These items typically affect the entire dashboard.

- In the toolbars of the widgets included in the dashboard. For example:

Patients by Favorite Color	
Favorite Color	
None	221
Blue	129
Green	137

These items *typically* affect only the widget in which they are displayed, but they could possibly affect other widgets.


The [next chapter](#) provides information on using these options.

3.6 Empty Widgets

There are two scenarios in which you might see empty widgets.

In the first scenario, a widget (or the underlying pivot table, if applicable) has been deliberately configured not to automatically execute when you open the dashboard. (This configuration can be useful when a pivot table uses a long-running query.) In this scenario, the widget should include a refresh button , which you use to force the widget to display its contents. If there is no such button, contact the owner of the dashboard.

In the second scenario, if the system cannot find some of the data that the dashboard uses, the system displays as much of the dashboard as possible. For any affected widget, the system displays something like this:



This typically indicates that someone has either renamed or deleted the data source used by this widget; it could also indicate that you do not have permission for the data in this widget. Contact the creator of the dashboard or reconfigure the widget yourself; see [Creating DeepSee Dashboards](#).

4

Using Dashboards

This chapter describes how to work with the dashboards that you may already have. It discusses the following topics:

- [How to toggle the display of any worklists](#)
- [How to manage the size and position of the widgets](#)
- [How to filter the data](#)
- [Understanding what filters do](#)
- [How to sort the data](#)
- [How to specify the row or column count](#)
- [How to change the display type](#)
- [How to use the chart zoom feature](#)
- [How to use combo charts](#)
- [How to display a listing](#)
- [How to drill down to a lower level](#)
- [How to use the dimension list](#)
- [How to drill to a different dashboard](#)
- [How to refresh a widget](#)
- [How to reload a dashboard](#)
- [How to export data to Excel](#)
- [How to print data](#)
- [How to use map widgets](#)
- [How to use other options](#)
- [How to save changes](#)

4.1 Toggling the Display of the Worklists

Depending on the system configuration and the configuration of the dashboard, a dashboard might include zero, one, or two worklist areas on the right.

You can temporarily toggle the display of any worklists in the dashboard you are currently viewing. To do so, select **Menu > Hide/Show Worklists**. If you save the dashboard, this change is not saved.

4.2 Managing the Widget Windows

This section describes how to manage the size and placement of widgets. All these changes are discarded unless you do one of the following:

- Select **Menu > Save My Settings**. This option saves your changes locally so that only you can see them.
- Save the dashboard itself, as described in [Creating DeepSee Dashboards](#).


4.2.1 Window Buttons

In all widgets, to control the display of the widget window, use the buttons in the title bar:

Button	Action
	Temporarily hide the widget. (To restore it, redisplay the dashboard.)
	Maximize the widget.

4.2.2 Resizing Widgets

To resize a widget, drag its resize handle in the lower right corner, if present:



Your ability to resize widgets is also affected by the **Scalable Grid**; see the [next section](#).

4.2.3 Moving Widgets

Depending on how the dashboard is configured, it may be possible to move the widgets on it. There are three different possibilities:

- The **Scalable Grid** option is enabled. In this case, there is an invisible grid of possible positions for the widgets. You can move widgets and resize them according to this grid.
If you resize the dashboard, all the widgets are automatically resized. **Scalable Grid** is the only option that supports this behavior.
- The **Snap Into Place** option is enabled. In this case, if a dashboard contains only one widget, that widget is automatically positioned in the upper left and cannot be moved. If a dashboard contains at least two widgets, you can reposition them. To move a widget to the right, left, top, or bottom of another widget, select in its top bar and drag.
- Neither option is enabled. In this case, you can freely reposition and resize widgets.

To access the **Scalable Grid** and **Snap Into Place** options, select **Menu > Save With Options....**

4.3 Filtering Data

Many dashboards include controls that you can use to filter the data shown in one or more widgets. Typically, if a control is in the **Filters** worklist, it affects the entire dashboard. In contrast, if it is in a widget toolbar, it typically affects only that widget.

Most filter controls are shown as searchable controls. The details depend on the type of data used in the control, as well as the source of the data; the following sections explain the variations.


Note: In some cases, a filter control might be read-only. In this case, its purpose is simply to tell you how the data is filtered.

4.3.1 Using a Filter Control


A filter control might initially be like this:



To use this control, select the Search button . In most cases, the system then displays something like this:



If this menu includes the word **Required**, then it is necessary to choose at least one item. In this case, the drop-down menu does not list the **All** choice. The following shows an example:




In this case, if you try to close the drop-down menu without selecting an option, the system displays a pop-up message saying that a value is required.

The **Range** option is not available if the [data source](#) for the widget is a KPI. The following subsections show other variations. In these filter controls, depending on their exact nature, you can do the following:

- Type into the search box and press **Tab**.

If you do, the control is refreshed to display only members that include the string that you typed. For example:



- Select an item or hold down the **Ctrl** key and select multiple items.

In some cases, it is not possible to select multiple items.

- Select **Exclude** to exclude the selected item or items.
- Select **Range** and then specify an inclusive range of members. When you select **Range**, the control changes to display two drop-down lists. In each drop-down list, select a member.


When you are done, select the check mark button to accept your selections. Or select the X button to discard them.

The widget or dashboard may or may not be refreshed immediately, depending on the dashboard configuration. Look for a standard Refresh button  or a button labeled `Refresh`, `Rerun`, or so on. If the dashboard has such a button, typically you must select that button to see the change due to the filters.


4.3.2 Filter Control for a Time Level

Depending on the [data source](#), if the filter refers to a time level, the list includes a **NOW** member, which always refers to the date on which the pivot table was run. (More specifically, **NOW** refers to the applicable member of the time level for the date on which the pivot table was run. For a year filter, **NOW** refers to the year, for a year and month filter, **NOW** refers to the year and month, and so on.)

For example, consider the following:



In this case, the **NOW** member refers to the current year, when this pivot table is run. For another example:



In this case, the **NOW** member refers to the current year and current month, when this pivot table is run.

Note: The **NOW** option is not available if the [data source](#) for the widget is a KPI.


[Day levels](#) are handled differently; see the next subsection.

For some kinds of time levels, if you click the **NOW** member and click **Range**, DeepSee displays the following:




Here, you can specify a range of dates relative to the current date. To do so, type into the boxes shown on the right. For each box, to specify a date in the past, enter a minus sign (-), followed by an integer *n*. This specifies a date *n* time units in the past, where the time unit is determined by the kind of time level (in this case, years). Similarly, to specify a date in the future, specify a plus sign (+) followed by an integer *n*.

For example, the following filter selects a range of dates from five years in the past until now:



4.3.3 Filter Control for a Day Level

Depending on the [data source](#), if the filter refers to a day level, the system shows the following control instead of a list:



S	M	T	W	T	F	S
					1	2
3	4	5	6	7	8	9
10	11	12	13	14	15	16
17	18	19	20	21	22	23
24	25	26	27	28	29	30
					«	»

If you select **Range**, the system displays two calendars, side by side. Below each calendar, the filter displays a **NOW** option. **NOW** represents the current date (that is, the date when the dashboard queries are executed). The following shows the **NOW** options:

All
 Exclude
 Range

From

November				2013		
S	M	T	W	T	F	S
					1	2
3	4	5	6	7	8	9
10	11	12	13	14	15	16
17	18	19	20	21	22	23
24	25	26	27	28	29	30
					«	»

NOW -365

To

November				2013		
S	M	T	W	T	F	S
					1	2
3	4	5	6	7	8	9
10	11	12	13	14	15	16
17	18	19	20	21	22	23
24	25	26	27	28	29	30
					«	»

NOW ±days

If you select **NOW**, you can specify an optional offset, which specifies an integer number of days before or after **NOW**. For example, **NOW - 365** represents the date that is 365 days before the dashboard query was executed.

4.3.4 Filter Control for a Measure

If the filter refers to a measure, the system shows the following control instead of a list:

Measures.Test Score


All

Test Score

In this case, select an operator from the drop-down list and type a value into the box to its right.

4.3.5 Filter Control for a Computed Dimension

If the filter refers to a computed dimension, the system shows the following control instead of the usual drop-down list:



In this case, you can either select an item in the list or select **Reset** to clear the selection.

A computed dimension is a dimension whose members are defined at runtime, typically via SQL SELECT statements.

4.3.6 Filtering via OnClick

In some cases, if you select a row in a widget, that widget filters another widget. This is a type of *onclick control*.

For example, consider the following dashboard:

	Patient Count			Patient Count
32006	2,203	^	Cedar Falls	1,155
32007	1,065		Centerville	1,055
34577	3,375		Cypress	1,096
36711	1,055		Elm Heights	1,147
38928	2,302		Juniper	1,069
			Magnolia	1,091
			Pine	1,188
			Redwood	1,065
			Spruce	1,134

For this dashboard, when you select a ZIP code row in the left widget, the system automatically refreshes the right widget. For example:

	Patient Count			Patient Count
32006	2,203	^	Cedar Falls	1,155
32007	1,065		Elm Heights	1,147
34577	3,375			
36711	1,055			
38928	2,302			

To clear the selection, select the left column header (for example, select above the **32006** cell).

4.3.7 Data Source Differences

For any widget that includes filter controls, if the data source is a KPI, a filter control has some or all of the following limitations:

- Does not include a **Range** option.
- Cannot be displayed in the calendar format.
- Does not include the system **NOW** option. (It is possible, however, for a KPI to include a filter option labeled **NOW** with similar behavior.)
- Does not necessarily enable you to select multiple items.

4.4 Understanding the Effects of Filters

In most cases, a filter includes or excludes members of a level. Each member is a set of base records, so a level filter includes or excludes records that correspond to the members you choose; the behavior is the same regardless of the contents and organization of the pivot table.

For example, consider the following dashboard:

ZIP		ZIP		ZIP	
<input type="text"/>		<input type="text"/>		<input type="text"/>	
ZIP	Patient Count	city	Patient Count	Age Group	Patient Count
32006	214	Cedar Falls	112	0 to 29	447
32007	112	Centerville	120	30 to 59	387
34577	337	Cypress	105	60+	166
36711	120	Elm Heights	105	Total	1,000
38928	217	Juniper	115		
Total	1,000	Magnolia	127		
		Pine	105		
		Redwood	112		
		Spruce	99		
		Total	1,000		

This dashboard contains three pivot table widgets, each of which displays a different pivot table. Each pivot table includes a total line. Each widget includes a level filter that uses the **ZIP** level (this is for demonstration purposes; more commonly, the **Filters** box would include a single filter that applied to all widgets).

Consider what occurs when we choose the same ZIP code in each of these widgets:

ZIP		ZIP		ZIP	
<input type="text" value="34577"/>		<input type="text" value="34577"/>		<input type="text" value="34577"/>	
ZIP	Patient Count	city	Patient Count	Age Group	Patient Count
34577	337	Cypress	105	0 to 29	142
Total	337	Magnolia	127	30 to 59	138
		Pine	105	60+	57
		Total	337	Total	337

The first widget shows patients grouped by ZIP code. This widget now shows one row, corresponding to the selected ZIP code. The other ZIP codes are no longer shown, because this pivot table is not configured to display null rows.

The second widget shows patients grouped by home city. This widget now shows the three cities that correspond to the selected ZIP code. The other cities are no longer shown, because this pivot table is not configured to display null rows.

The third widget shows patients grouped by age groups. This widget displays all age groups, because the selected ZIP code includes patients of all ages.

Notice that the total count is the same in each pivot table.

In each case, when we select a ZIP code, DeepSee selects only patients whose home address is in that ZIP code.

A measure filter refers to the value of a measure, instead of referring to members of a level. For example:



A measure filter examines each base record and includes or excludes it, depending on its value for the given measure. For example, consider the following dashboard:

Test Score		Test Score	
<input type="text"/>		<input type="text"/>	
Patient Group	Patient Count	Favorite Color	Patient Count
Group A	402	None	221
Group B	398	Blue	129
None	200	Green	137
Total	1,000	Orange	133
		Purple	127
		Red	107
		Yellow	146
		Total	1,000

Consider what occurs when we filter by Test Score in the same way in both widgets:

Test Score		Test Score	
<input type="text" value=">=65"/>		<input type="text" value=">=65"/>	
Patient Group	Patient Count	Favorite Color	Patient Count
Group A	299	None	118
Group B	268	Blue	74
Total	567	Green	76
		Orange	78
		Purple	76
		Red	65
		Yellow	80
		Total	567

The first widget shows patients grouped by patient group. This widget no longer shows the patients in group None, which contains the patients who have no test score; this row is gone because this pivot table is not configured to display null rows.

The second widget shows patients grouped by favorite color. This widget shows all the same rows it used to show, because there are patients with the selected score or higher in all these groups.

Notice that the total count is the same in each pivot table.

4.5 Sorting Data

In a pivot table widget, you can sort the rows by the values shown in any data column. To do so, double-click the column header. For example:

Age Group-	Age Bucket	Patient Count ▼	Avg Allergy Count
60+	80+	41	1.12
60+	70 to 79	62	0.97
60+	60 to 69	86	1
30 to 59	50 to 59	106	0.94
0 to 29	20 to 29	116	1.06
0 to 29	10 to 19	135	1.10
30 to 59	40 to 49	139	1.18
0 to 29	0 to 9	149	1.08
30 to 59	30 to 39	166	1.04

The first time you double-click the column header, the table is sorted in ascending order, by the values in that column, as shown in the preceding example.

If you double-click the column header again, the table is sorted in descending order instead. If you double-click a third time, the sorting is removed and the default order is restored.


In some cases, you might see a drop-down list that you use to control the sorting of the rows or column. Select **Decreasing**, **Increasing**, or **Sort** (removes the sorting). In this case, the control considers the values in the first data column or the first data row, as appropriate.

4.6 Specifying the Row or Column Count

In some widgets, you can specify how many rows or columns are shown. For example:



For another example:



Type or select a value as needed.

4.7 Changing the Display Type

In some widgets, you may be able to toggle the display between table format and chart format. To do so, use the following buttons:

Button	Action
	Displays the widget in table format.
	Displays the widget in chart format.

Note that for a chart, only the first hundred chart items are displayed.

Or the widget might include a control that lets you choose the chart type. For example:



In this case, select an option from the drop-down list. The widget is then redisplayed in the selected format.

Or the widget might include buttons that let you switch to specific types. For example:




In this case, when you click a button, the widget is redisplayed in the selected format. From left to right, the buttons shown here display the data as a pivot table, a bar chart, or a column chart. There are similar buttons for other chart types.

4.8 Using the Chart Zoom Feature

For some kinds of charts, you can zoom in and out, to see more or less data at once. To do so, select the zoom in button (+) in the lower left corner:



The chart then shows less data and includes arrow buttons that you can use to scroll left and right. For example:




4.9 Using Combo Charts

A combo chart provides two features that are not available in other charts:


- It can have multiple y-axes.
- It can use different display types for different measures.

The following chart demonstrates both features:



Notice that the y-axis is labeled Patient Count. The bars that display the Patient Count measure are scaled to match the values on this axis. In contrast, the bars that display Avg Enc Count do not correlate with the y-axis — these bars are correctly proportioned with respect to each other, but are scaled overall to fit within the chart.

To use a combo chart, select the box next to the name of a series, in the key. For example, if we select the box next to Avg Enc Count, the labels on the y-axis change as follows:



4.10 Displaying Listings

If a widget includes the Listing button , you can display a listing, which shows the lowest-level data for the context you select. In some cases, the listing button might appear as a button with text (such as `Toggle Details`).

To display a listing:

1. Select one or more items.

- For a table, select one or more cells.

To select multiple cells, hold down the **Shift** key while selecting the cells.

To select an entire row, select the row label on the left. To select an entire column, select the column header.

The listing option is not available for cells in a total row or a total column.

- For a chart, select a single chart item. For example, select a bar in a bar chart.

The selected item or items are highlighted.

2.

Select the Listing button .

The widget then toggles to list the lowest level records for the selected cells. For example:

<input type="checkbox"/>	#	PatientID	Age	Gender	Home City	Test Score
<input type="checkbox"/>	1	SUBJ_100631	0	F	Elm Heights	50
<input type="checkbox"/>	2	SUBJ_100781	0	F	Redwood	78
<input type="checkbox"/>	3	SUBJ_100820	0	F	Magnolia	89
<input type="checkbox"/>	4	SUBJ_100966	0	F	Cypress	91
<input type="checkbox"/>	5	SUBJ_101274	0	F	Pine	
<input type="checkbox"/>	6	SUBJ_101340	0	F	Redwood	
<input type="checkbox"/>	7	SUBJ_101466	0	F	Magnolia	81
<input type="checkbox"/>	8	SUBJ_101532	0	F	Pine	67
<input type="checkbox"/>	9	SUBJ_101587	0	F	Elm Heights	77
<input type="checkbox"/>	10	SUBJ_102327	0	F	Redwood	79

If there are more than 100 rows, the bottom area displays buttons that you can use to page through all the rows.

The check boxes enable you to select rows so that you can perform actions (particularly custom actions) that use those rows as input. Actions are usually shown as buttons at the top of the widget.


If your pivot table uses a subject area that includes an iKnow measure, the iKnow Measure Value  button might be available; for details, see “[Displaying iKnow Measure Values for Selected Facts](#),” later in this chapter.

3.

To toggle back to the original view, select the Table button .

4.10.1 Map Listing Variation

In some cases, the widget displays a Map Listing button: . If you select this, the system displays a map that shows locations. For example, from the HoleFoods sample:



In this case, the displayed locations are locations of the sales for the selected records. Map listings are also known as *geo listings*.

4.10.2 Displaying iKnow Measure Values

When you display a listing for a pivot table that uses a subject area that includes an iKnow measure, the listing might include a column that displays the iKnow Measure Value  button. For example:

StateCode	City	AirportName	CoordsLatitude	CoordsLongitude	Report	R
GA	HOGANSVILLE	NEWNAN-COWETA COUNTY				Ir
AZ	MESA	FALCON FIELD	33.2739	-111.4342		C
AZ	Phoenix	Phoenix Deer Valley	33.413	-112.0456		T
FL	Naples	Naples Municipal	26.0915	-81.4653		C
FL	Melbourne	Melbourne International	28.061	-80.3845		C
AL	Sylacauga	Merkel Sylacagua	33.1031	-86.1833		C
OH	Swanton	Toledo Express	41.35	-83.4748		T
CA	TRUCKEE	TRUCKEE TAHOE	39.4340	-120.0007		C

If you select this button, DeepSee then displays a page like the following:

Record #1

On January 2, 1999, at 1618 eastern standard time, a Piper PA-34-200, N5016T, collided with the ground during an emergency landing near Hogansville, Georgia. The airplane was operated by the registered owner/pilot under the provisions of Title 14 CFR Part 91, and instrument flight rules (IFR). Instrument meteorological conditions prevailed at the accident site and an IFR flight plan was filed for the personal flight. The commercial pilot received minor injuries and the three passengers received no injuries. The airplane sustained substantial damage. The flight departed Manassas, Virginia, at 1035 the same day, destined for Columbus Metropolitan Airport, Columbus, Georgia. According to the pilot, he received a full weather briefing from the Leesburg, Virginia, Flight Service Station at 0430 on the morning of the flight. Additional weather information was obtained enroute via weather broadcasts over navigational radios, and from automated surface observations. While approximately 30 miles from Columbus Metropolitan Airport (CSG), the pilot was informed of a lowered ceiling and thunderstorm over the destination airport. He immediately notified Air Traffic Control (ATC) of his low fuel condition and was vectored to LaGrange Airport (LGC). Once stabilized on the ILS approach for LGC, the pilot realized he was on the back course and a missed approach was executed. The pilot was then given vectors towards Newnan-Coweta County Airport (CCO) which was reported "in the clear". According to the pilot, an emergency was declared to Atlanta Approach Control several times but no response was received. After starting a descent, the pilot reported that "the left engine ran out of fuel", and quit. The pilot feathered the left propeller, and he increased power on the right engine. Visual contact with the ground was made about 400 feet AGL. During the emergency landing, the airplane collided with the top of a tall pine tree. The airplane landed with the gear and flaps retracted on a hilly power line right-of-way. According to the Piper Aircraft PA-34 Information Manual, N5016T has a usable fuel capacity of 93 U.S. gallons and seven gallons unusable. The pilot reported that he departed with 100 gallons of fuel on board. While operating the airplane at a gross weight of 4000 pounds and 65% power, fuel consumption is 18.3 gallons per hour for both engines, or maximum flight duration of 5.08 hours. N5016T had flown 5.72 hours when the accident occurred. The pilot did not flight plan or make a refueling stop for this flight. No mechanical problems, with the airplane, were reported by the pilot. ...

Summary length:

number of sentences to display

Highlight nothing dictionary terms dominant terms

Dictionaries

 Maneuvres Weather conditions Pilot skills InjuriesDominant concepts relations CRCs paths

This page displays (by default), the complete text of an iKnow measure, for that record.

On this page, you can do the following:

- Select a tab to display the text for a different record.
- Select the most relevant sentences of this text. To do so, type an integer into **Summary length** and then select **summarize**.

The system then displays that number of sentences, starting with the most relevant sentence, as determined by the iKnow engine.

- Highlight specific kinds of elements, as determined by the iKnow engine. By default, nothing is highlighted.

Note the following points:

- A *dictionary term* is a word or phrase found in a specific dictionary loaded into the system. In general, a dictionary defines a mapping that typically associates multiple synonyms with each other. You can select the dictionary or dictionaries to consider, among the available dictionaries in your system; these are specific to your implementation
- A *dominant term* is a unit of text that is dominant within this text, as determined by the iKnow engine. To understand the choices, it is necessary to review some iKnow concepts.

For iKnow, a key concept is the *entity*. The iKnow engine processes the given text and identifies the entities in it. An entity is a minimal logical unit of text. It is a word or a group of words. Example entities are `clear skies` and `clear sky`. The iKnow language model identifies two kinds of entities: relations and concepts. A *relation* is a word or group of words that join two concepts by specifying a relationship between them. A relation is commonly but not always a verb. A *concept* is a word or group of words that is associated by a relation. A concept is commonly but not always a noun or noun phrase.

The page can also highlight dominant CRCs and paths. A *CRC* is a concept-relation-concept sequence, and a *path* is a longer sequence.

4.11 Drilling Down

In some widgets, you may be able to drill down to lower levels. In the DeepSee model, levels are contained within *hierarchies*, which are contained within *dimensions*. In some cases, there are multiple levels within a hierarchy. In these cases, a member of the higher level contains one or more members of the lower level, and you can drill down — navigate from the higher-level member to the lower-level members. You can do this via [double-click](#) or (if available) via the [Drilldown button](#).

It is not possible to drill down from a total row. Also note that the drilldown action has no effect if the pivot table uses a manually entered MDX query (which is a rare scenario).

4.11.1 Drilling Down with Double-Click

If a pivot table row represents a higher-level member, you can double-click the member name to view the lower level members. For example, consider the following pivot table widget:

Decade	Patient Count	Avg Enc Count
1910s	81	76.28
1920s	205	62.04
1930s	517	56.34
1940s	709	51.23
1950s	1,066	44.68
1960s	1,469	39.53
1970s	1,544	33.74
1980s	1,315	26.62
1990s	1,514	20.56
2000s	1,436	14.46
2010s	144	10.48
All Patients	10,000	33.05

If you double-click the cell 1930s, the widget then displays the following:

Decade	Patient Count	Avg Enc Count
« 1930	31	64.42
« 1931	51	65.69
« 1932	40	67.40
« 1933	62	57.82
« 1934	47	45.45
« 1935	43	50.44
« 1936	60	55.38
« 1937	72	47.31
« 1938	52	61
« 1939	59	55.80

If you select the << characters in any row, the system restores the previous state of the widget.


This kind of drilldown action is not available in charts.

4.11.2 Using the Drilldown Button

A widget might include the Drilldown button , which provides you with an alternative way to drill down in a hierarchy. In this case, to drill down:


1. Select a row in the pivot table. Or select a chart item if the widget is in chart format.

For example:



2. Select the Drilldown button .

The widget then displays data for the lower-level members associated with the selected item. For example:



4.12 Using the Dimension List

A widget might include the Dimension List button , which provides you with additional drill options.

If you select this button, the widget then displays a list of dimensions, to the left of the pivot table. For example:



This list may or may not be the same as the dimensions defined in the cube, depending on how it was created.

To use this list:

1. Expand the folders as needed.
2. Drag and drop from the list onto the gray part of any row.

For example:

	Patient Count	Avg Age	Avg Test Score	Avg Enc Count
« None	116	32.93	75.97	0
« I	129	35.95	72.88	0
« II	73	32.22	75.30	0
« III	80	40.83	72.93	0

For another example:

	Patient Count	Avg Age	Avg Test Score	Avg Enc Count
« Burroughs, Angelo	9	38.78	81.22	0

To reverse the effect, select the Reset button  above the dimension list or select the << button in the pivot table.

4.13 Drilling to a Different Dashboard

In some pivot table widgets, you might be able to double-click a cell in order to display a different dashboard. The newly opened dashboard is filtered to the context in which you double-clicked. For example, suppose that the first dashboard includes the following widget:

Gender	Patient Count	Avg Age
Female	5,108	36.79
Male	4,892	34.57
All Patients	10,000	35.71

Depending on how this pivot table widget is configured, if you double-click a data cell, DeepSee might display a different dashboard, as follows:

- If you double-click in the `Female` row, the newly opened dashboard is filtered to female patients.
- If you double-click in the `Male` row, the newly opened dashboard is filtered to male patients.
- If you double-click in the `All Patients` row, the newly opened dashboard is not filtered by gender.

4.14 Refreshing a Widget

To refresh the display in a widget, select the Refresh button .

In some cases, this button might appear as a button with text. The text on the button might be `Refresh`, `Redisplay`, or any other suitable string.

In some cases, the widget might automatically refresh at a specified interval. In this case, there might not be a refresh button.

4.15 Reloading a Dashboard

To reload the dashboard itself, select the Reload Dashboard button .

In some cases, this button might appear as a button with text. The text on the button might be `Reload` or some other suitable string.

4.16 Exporting Data to Microsoft Excel

You can export data to Microsoft Excel from a pivot table widget, but not from other types of widgets. To do so, select the

Export to Excel button  in the widget.

The system generates an Excel file that contains all the data currently displayed in the widget. The file has the name `%DeepSee.UI.MDXExcel.zen.xls`, `%DeepSee.UI.MDXExcel.zen-1.xls`, or similar. Your browser then does one of the following, depending on your operating system, the configured file types on your machine, your browser, and your browser settings:

- Opens the file with Excel (or other selected program).
You specify the program to use by setting a browser option. For example, on Firefox, you use **Tools > Options > Applications**. This option also enables you to specify whether the browser should open the file with this program or prompt you to save it to the hard drive.
- Prompts you to open the file with Excel (or other selected program, as described in the previous bullet).

- Opens this file within the current browser window.

This is the default behavior for Windows operating systems earlier than Windows Vista. You can change this behavior on the **Edit File Type** dialog box (which you access via **My Computer > Tools > Folder Options** or **My Computer > View > Options**, depending on the operating system version).

This option is not available on Windows Vista and later Windows operating systems.

- Prompts you to save the file to the hard drive.
- Automatically saves the file to the hard drive.

When you export to Excel, note that it is not possible to export more than 250000 rows. Also, if the subject area uses an external table, there is a limit of 1000 rows in the listing.

Note: If no numeric formatting is specified, any numbers are formatted as integers by default. (The actual values are available in Excel, and you can change the formatting as needed to display them appropriately). If numeric formatting is specified in the model or in the pivot table definition, that formatting is used in Excel.

The pivot table name is shown in the exported data, unless the pivot table definition specifies a different print name; see “[Customizing Pivot Tables](#)” in *Using the DeepSee Analyzer*.

4.17 Printing Data

The dashboard might include one or more Print buttons . Depending on how it is configured, a Print button might print only a single widget or multiple widgets. In any case, when you use this button, DeepSee generates a .pdf file and opens it in your default browser. (If the output represents multiple widgets, the file contains one page for each widget.) You can then use options in the browser to print the file or to download it, for printing later.

Note that if you download the file, the browser saves the file to the default download directory for that browser; there is no mechanism for DeepSee to specify a different location.

If the pivot table is extremely wide, it is not possible to generate a .pdf file for it.

4.17.1 Requirements for Printing

When you invoke the Print option, Caché uses Java to call out to a third-party PDF rendering tool. This means that Java (or specifically JRE, Java Runtime Executable) is required on your local machine. For information on the requirements, which also apply to Zen Reports, see “[Configuring Zen Reports for PDF Output](#)” in *Using Zen Reports*.


4.18 Using a Map Widget

A map widget includes a different set of options that are not present in other widgets:



In a map widget, you can do the following:

- Toggle the display between the map view and the satellite view. To do so, select the **Map/Satellite** button.
- Display a different area. To do so, drag the cursor while holding down the left-select button.
- Zoom in or out. To do so, drag the slider on the zoom bar.
- Display additional details about a data point. To do so, select the data point. The widget then displays a popup window, if there is any additional data. For example:




To close this window, select the X in the upper right.


4.19 Other Tasks

In a dashboard, you may also be able to do the following:

- Reconfigure any widget. To do so, select the < button (if available) on the left of the dashboard to open the Dashboard Editor. Then make changes as described in [Creating DeepSee Dashboards](#).
- Select items that redefine the widget in different ways. There are many options. For example, the widget could include a drop-down list where you can choose a different pivot table to display:



For another example, it could include a drop-down list that controls what is shown for rows or columns:



- Display the Mini Analyzer and make local changes to the pivot table shown in a widget. Select the Mini Analyzer button , if included, and make changes as described in [Using the DeepSee Analyzer](#).
The changes do not affect other widgets, other dashboards, or other users.
Then to save your changes and exit the Mini Analyzer, select **OK** in the Mini Analyzer. Or, to discard your changes and restore the original pivot table, select **Reset**.
- Perform *custom actions*. The **Filters** box and the widgets might include additional buttons to perform custom actions.

In some cases, you might first need to select a row or rows in a widget. To do so:

- In a standard pivot table widget, to select a row, select the row. To select multiple rows, hold down the **SHIFT** key while selecting those rows.
- In a pivot table widget that displays a listing, select the check boxes at the left of the rows.

A custom action can do any of the following:

- Open a different dashboard in the same window.
- Open a URL in the same window or in a new window.
- Execute server-side code.
- Execute client-side commands such as refreshing the dashboard.

4.20 Saving Your Changes

To save the current state of the dashboard that you are viewing, select **Menu > Save My Settings**. This option performs a local save of the current dashboard. It saves the size and placement of each widget and the current state of all filters and other controls in the dashboard. These changes are saved only for you, and are not visible to other users.

If you select **Menu > Clear My Settings**, the system removes your settings for this dashboard. After using this option, use the browser's refresh button to display the dashboard in its base state.

There are also options for saving the base dashboard design itself so that all users see the changes (see [Creating DeepSee Dashboards](#)). You can use these only if you have the appropriate permissions.

For reference, the following table describes how the different dashboard changes can be saved.

Dashboard Change	Can Be Saved As My Setting?	Can Be Saved As Dashboard Design?
Maximizing or minimizing a widget	Yes	Yes
Moving or resizing a widget	Yes	Yes
Changes made via the Dashboard Designer	No	Yes
Using the Mini Analyzer to make local changes to the pivot table used in a widget	Yes	No
Specifying values for filters or for other widget controls	Yes	No
Other changes (sorting, drilling down, and others)	No	No

When you save the dashboard design, your settings for that dashboard are cleared, except for any changes made via the Mini Analyzer.

5

Using the Pivot Analysis Window

This chapter describes how to use the Pivot Analysis window, which you can access via the Analysis button  on pivot table widgets. This chapter discusses the following:

- [How to use the Pivot Analysis window generally](#)
- [Details for a cluster analysis](#)
- [Details for a distribution analysis](#)
- [Details for a regression analysis](#)
- [Details for an iKnow content analysis](#)
- [Details for an iKnow entity analysis](#)

5.1 Using the Pivot Analysis Window

DeepSee provides an analysis window that you can use for several specialized kinds of analysis. In each case, you first select one or more cells, and the analysis considers the lowest-level data associated with those cells. To access this window:

1. Select the data cells in the row or rows.

To select multiple cells, hold the **Shift** key down while selecting the cells.

To select an entire row, select the row label on the left. To select an entire column, select the column header.

The analysis option is not available for cells in a total row or a total column.

2. Select the Analysis button .

Depending on how the widget is configured, the system either provides a choice of analysis options or displays one of them without any choice.

3. (If applicable) For **Analysis Option**, select one of the following:
 - **Cluster** — Performs a [cluster analysis](#).
 - **Distribution** — Performs a [distribution analysis](#).
 - **Regression** — Performs a [regression analysis](#).

Or select **iKnow Plugins** and then select one of the following:

- **Content Analysis** — Performs an [iKnow content analysis](#).
- **Entity Analysis** — Performs an [iKnow entity analysis](#).

The iKnow options are applicable only if your cube includes iKnow measures.

The following sections provide the details.

Note: This window is also available in the Analyzer.

5.2 Cluster Analysis

For background on the Pivot Analysis window, see “[Using the Pivot Analysis Window](#).”

Cluster analysis or *clustering* is the assignment of a set of observations into subsets (called clusters) so that observations in the same cluster are similar in some sense. Clustering is a method of unsupervised learning, and is a common technique for statistical data analysis used in many fields, including machine learning, data mining, pattern recognition, image analysis, information retrieval, and bioinformatics.

For a cluster analysis, the top area of the page looks something like the following:

Measures To Include in Clustering:

- Age
- Allergy Count
- Avg Age
- Avg Allergy Count
- Avg Enc Count
- Avg Test Score
- Encounter Count
- Test Score

Input (X) Measure: Age

Response (Y) Measure: Encounter Count

Algorithm: CLARA Sample Size: 75 Number of Clusters: 3 Run

Items: 1,080 X: [:] Y: [:]


To use this page:

1. Specify details of the analysis in the top area of the page.

Details are beyond the scope of this book; it is assumed that the reader is familiar with cluster analysis.

2. Select **Run**.

The bottom area of the page displays the results. For example:



For general information on cluster analysis, see the Wikipedia page (http://en.wikipedia.org/wiki/Cluster_analysis). There are also many books available on this topic.


Also see the DataMining package in the SAMPLES database.

5.3 Distribution Analysis

For background on the Pivot Analysis window, see “[Using the Pivot Analysis Window.](#)”

A *distribution analysis* shows the number of occurrences of different values of specific measurement, across a set of records.

For a distribution analysis, the system displays something like the following:



To use this page, select a measure (such as `Age` in this example).

The horizontal axis then shows all the values of this measure for the selected cells. For each measure value, the vertical axis shows the number of source records (for the selected cells) that have that value.

The top of the page shows the following values:


- **Items** — Number of source records for the selected cells of the pivot table.
- **Missing** — Number of source records (for the selected cells) that do not have a value for this measure.
- **Min value** — Lowest value for this measure among these source records.
- **Max value** — Highest value for this measure among these source records.
- **Avg** — Average of these values.

5.4 Regression Analysis

For background on the Pivot Analysis window, see “[Using the Pivot Analysis Window.](#)”

A *regression analysis* attempts to determine the relationship between independent variables and a dependent variable. (The DeepSee regression analysis considers only one independent variable.)

For a regression analysis, the system displays something like the following:



To use this page, specify the following details:

- **Input (X) Measure** — Select the measure that you want to treat as independent (such as `Age` in this example). These values are shown on the horizontal axis.
- **Response (Y) Measure** — Select the measure that you believe is dependent on the first measure (such as `Encounter Count` in this example). These values are shown on the vertical axis.

When you do so, the page displays a data point for each lowest-level record associated with the selected cells.

- **Fit** — Select the type of fitting to perform:
 - **line** — Performs a linear regression.
 - **exp** — Performs an exponential regression.
 - **log** — Performs a logarithmic regression.
 - **power** — Performs a power regression.
- **Confidence Level** — Optionally select the confidence level to use when calculating error bars for the values being determined.

In each case, the system determines the details of the equation that predicts Y as a function of X, and it plots the predicted X-Y curve on the page. The area above the chart displays details of the fit. For example, for a linear regression:

$$\text{Equation: } Y = (0.70 \pm 0.059) * X + (8.98 \pm 2.51) \quad \text{Correlation Coefficient} = 0.57$$

For another example, the page displays the following details for an exponential regression:

$$\text{Equation: } Y = 3.99 * \exp(0.024 * X) \quad \text{Correlation Coefficient} = 0.19$$

For general information on regression analysis, see the Wikipedia page (http://en.wikipedia.org/wiki/Regression_analysis). There are also many books available on this topic.

5.5 iKnow Content Analysis

For background on the Pivot Analysis window, see “[Using the Pivot Analysis Window.](#)”

The *iKnow content analysis* option displays information about the most typical and least typical unstructured text values. This analysis option is applicable only if your cube includes iKnow measures.

For the iKnow content analysis, the system displays something like the following:

Measure		Display listing				
Report						
Most typical facts						
	EventId	EventDate	CountryCode	StateCode	AircraftModel	Score
Q	20021204X05552	2002-11-28 00:00:00	USA	NC	150L	255192
Q	20050519X00638	2005-04-26 00:00:00	USA	AK	180	248106
Q	20021115X05472	2002-11-07 00:00:00	USA	OK	J35	237502
Q	20030103X00010	2002-12-23 00:00:00	USA	AZ	R22 Beta	234611
Q	20010626X01255	2001-06-24 00:00:00	USA	AK	S-1B2	233878
Most breaking facts						
	EventId	EventDate	CountryCode	StateCode	AircraftModel	Score
Q	20001206X02438	1994-10-10 00:00:00	USA	TX	269C	11582
Q	20091103X85152	2009-11-02 00:00:00	USA	MI	PA-28-180	1463
Q	20001212X21503	2000-07-09 00:00:00	USA	CA	E-3	708
Q	20010110X00083	2000-12-20 00:00:00	USA	AK	C46A	514
Q	20050913X01440	2005-08-28 00:00:00	USA	FL	8GCBC	453

The **Most typical facts** section lists the records that have the most typical content, as determined by the iKnow engine. The **Most breaking facts** section lists the records that have the breaking content. Breaking content is content in which the dominant entities in that source are least similar to the dominant elements of the group of sources. For more information, see the chapter “Dominance and Proximity” in *Using iKnow*.

Each table lists the records and displays the fields listed in the selected detail listing.

On this page, you can do the following:

- See the entire unstructured text for a given record. To do so, select the magnifying glass icon shown by that record.
- Select a different iKnow measure to analyze. To do so, select a choice for **Measure**.

An iKnow measure is different from other kinds of measures and cannot be used in the same way as other measures. It is based on a source value that contains unstructured text such as a free-form text report. When you select an iKnow measure to analyze here, you are examining results of the iKnow engine for those source values.

- Select a different detail listing to use. To do so, select a choice for **Display Listing**.

The system initially displays the listing named `ShortListing` (if that is defined) or the default listing for the subject area.

This example uses the Aviation Events demo in the SAMPLES namespace. For reasons of space, this demo is not initialized when you install Caché. To set up the demo, open the Terminal and enter the following commands:


```
zn "SAMPLES"
do ##class(Aviation.Utils).Setup()
```

5.6 iKnow Entity Analysis

For background on the Pivot Analysis window, see “[Using the Pivot Analysis Window.](#)”

The *iKnow entity analysis* option displays information about the entities in your unstructured text values. This analysis option is applicable only if your cube includes iKnow measures.

For the iKnow entity analysis, the system displays something like the following:



This page provides three tabs, which display information about the records associated with the pivot table cells that you selected before you selected the Analysis button . These tabs work together as follows:

- The **Overview** tab displays information for the 20 most common entities among those records. You use the **Select by** option to control how DeepSee selects these entities; see the [first subsection](#).
Select a rectangle in this chart to see further detail; see the next bullet item.
- For a selected rectangle in the **Overview** tab, the **Cell breakdown** tab shows how that entity is distributed among the original pivot table cells. See the [second subsection](#).
Select a bar in this chart to see further detail; see the next bullet item.

- For a selected bar in the **Cell breakdown** tab, the **Entities** tab displays related entities. See the [third subsection](#).

This example uses the Aviation Events cube demo in the SAMPLES namespace. For reasons of space, this demo is not initialized when you install Caché. For information on setting it up, see the [previous topic](#).

5.6.1 Overview Tab

The **Overview** tab displays information about the top entities among the records associated with the pivot table cells that you selected before you selected the Analysis button . The chart on this tab displays one rectangle for each of the top 20 entities, according to your choice of metric.

5.6.1.1 Selecting the Top Entities

To determine how the system selects the top entities, select a metric from the **Select by** drop-down list. The options are as follows:

- **spread**, the default, lists the top entities by their spread (total number of facts in which the entity appears).
- **BM25** lists the top entities using a frequency calculation algorithm in descending order by calculated BM25 score. It calculates this score using an algorithm based on the Okapi BM25 standard, which combines an entity's Term Frequency with its Inverse Document Frequency (IDF), taking into account document length.
- **TFIDF** lists the top entities using a frequency calculation algorithm in descending order by calculated TFIDF score. It calculates this score by combining an entity's Term Frequency (TF) with its Inverse Document Frequency (IDF). The Term Frequency counts how often the entity appears in a single source. The Inverse Document Frequency counts the number of times the entity appears in the collection of sources, and uses this overall frequency to diminish the Term Frequency. Thus an entity that appears multiple times in a small percentage of the sources is given a high TFIDF score; an entity that appears multiple times in a large percentage of the sources is given a low TFIDF score.

Note: The **BM25** and **TFIDF** options are computation-intensive and can take some time to complete.

5.6.1.2 Color Coding

The colors in this chart indicate how well each entity serves as an indicator for a given pivot table cell, based on the Naive Bayes probability for that entity, as follows:

- If only one cell was originally selected in the pivot table, the colors indicate how well each entity serves as an indicator for that pivot table cell, compared to *all other records*.
- If multiple cells were originally selected in the pivot table, the chart displays the **Color-code by** drop-down list, which lists the pivot table cells by name and includes the **all** option.

In this case, the colors indicate how well each entity serves as an indicator for the pivot table cell that is currently shown in the **Color-code by** drop-down list, compared to *the records for all the other originally selected cells*.

To use a different pivot table cell as the basis of comparison, select an option from the **Color-code by** drop-down list. For example:



The **total** option considers all the originally selected cells as a group.


Green generally denotes good indicators, and red denotes bad indicators. Solid green means that the term is a very good indicator, pale green means that the term is a good indicator, pale red means that it is a bad indicator, and solid red means that it is a very bad indicator. Note that in some cases, an entity is a bad indicator because it is common to all sources and thus does not enable you to discriminate categories of sources.

For example, suppose that we use a pivot table that displays aircraft type as rows. As the starting point for analysis, we select the pivot table cells `Airplane` and `Helicopter`. If we use `Airplane` to color-code, this chart looks as follows:



Notice that the entities `airplane` and `runway` are good indicators for the pivot table cell `Airplane`. The other entities (especially `ground`) are bad indicators.

In contrast, if we use `Helicopter` to color-code, the chart would be colored as follows:



This shows that the entities `airplane` and `runway` are bad indicators for the pivot table cell `Helicopter`. The other entities (especially `ground`) are good indicators.

5.6.1.3 Other Options

On the **Overview** tab, you can also do the following:


- Select a different iKnow measure to analyze. To do so, select a choice for **Measure**.
An iKnow measure is different from other kinds of measures and cannot be used in the same way as other measures. It is based on a source value that contains unstructured text such as a free-form text report. When you select an iKnow measure to analyze here, you are examining results of the iKnow engine for those source values.
- Specify an additional string to analyze. To do so, type a string into **Analyze string** and then press **analyze**. The system then includes additional information for that string.
- Select a rectangle in this chart to display a breakdown for the corresponding entity. See the [next section](#).

5.6.2 Cell Breakdown Tab

The **Cell breakdown** tab is useful only if you started by selecting multiple cells of a pivot table.

This tab shows how an entity is distributed among the pivot table cells from which you started. When you select a rectangle in the **Overview** tab, the system displays the **Cell breakdown** tab with information for that entity.

The **Cell breakdown** tab looks like this:



For the given entity, this tab shows how that entity is distributed among the pivot table cells that you selected before you selected the Analysis button. For each cell, the chart shows five color-coded bars, which display the following series:

- **entities** — indicates the count of distinct, similar entities in records associated with the given pivot table cell or cells. This count includes all entities that are similar to the given entity.
- **entity freq** — indicates the *entity frequency*; that is, the number of times the given entity occurred in records associated with the given pivot table cell or cells.


If you select **Include similar entities**, this calculation also includes similar entities.

- **entity spread** — indicates the *entity spread*; that is, the total number of records in which the given entity appears, for the given pivot table cell or cells.

If you select **Include similar entities**, this calculation also includes similar entities.

- **cell probability** — indicates how well the given entity serves as an indicator for the selected pivot table cell or cells, based on the Naive Bayes probability for that entity.
- **dominance** — indicates how dominant the given entity is, within these records. For information, see “Semantic Dominance” in *Using iKnow*.

Important: These series are scaled separately so the chart can show them all. The scale shown on the chart applies only to the **entities** series. To see the actual numeric value for any bar, hover the cursor over the bar. For example:




On this tab, you can do the following:

- Select a different iKnow measure to analyze. To do so, select a choice for **Measure**.
- Include similar entities in the calculations for **entity freq** and **entity spread**. To do so, select **Include similar entities**.
- Display **entity spread %** (the entity spread as a percentage of the total number of records) instead of **entity spread**. To do so, select **Spread as %**.
- Specify an different string to analyze. To do so, type a string into **Analyze string** and then press **analyze**.
- Select a bar in this chart to display related entities. See the [next section](#).

5.6.3 Entities Tab

The **Entities** tab displays related entities in the given pivot table cell or cells. When you select a bar in the **Cell breakdown** tab, the system displays the **Entities** tab with the associated details.

The **Entities** tab looks like this:



On this tab, you can do the following:

- Select a different iKnow measure to analyze. To do so, select a choice for **Measure**.

- Specify a different string to analyze. To do so, type a string into **Analyze string** and then press **analyze**. The system then shows information for that string.
- Page through the results. To do so, use the << and >> buttons.

