



# Contents

<b>Copyright.....</b>	<b>6</b>
<b>Acumatica Framework Guide.....</b>	<b>7</b>
<b>Acumatica Framework Overview.....</b>	<b>8</b>
Acumatica Cloud xRP Platform.....	8
Acumatica Framework Development Tools.....	11
Runtime Architecture of an Application Based on Acumatica Framework.....	14
<b>Getting Started with Acumatica Framework.....</b>	<b>17</b>
Data Querying.....	18
Business Logic Controller Declaration.....	19
Data View and Cache.....	20
Data Modification Scenarios.....	20
Business Logic Implementation.....	29
<b>Designing the Application.....</b>	<b>32</b>
Designing the Database Structure and DACs.....	32
System and Application Tables.....	32
Table and Column Naming Conventions.....	32
Common Columns and Data Types.....	34
Primary Key.....	35
Foreign Keys and Nullable Columns.....	36
Audit Fields.....	36
Concurrent Update Control.....	37
Attachment of Additional Objects to Data Records.....	37
Preservation of Deleted Records.....	37
Multitenancy Support.....	38
Designing the User Interface.....	39
Development Environment Options.....	39
Form and Report Numbering.....	40
Item Grouping on the Form Toolbar.....	41
Naming the Graphs and Event Handlers.....	41
Graph Naming.....	42
Naming Conventions for Event Handlers Defined in Graphs.....	42
<b>Configuring ASPX Webpages and Reports.....</b>	<b>43</b>
Overview of ASPX Webpages in Acumatica Framework.....	43
Processing of a Button Click.....	44
Technical Overview of the Modern UI.....	47
Configuring Containers.....	51
Use of the DataMember Property of Containers.....	52

Use of the SkinID Property of Containers.....	53
Use of the Caption Property of Containers.....	54
Use of Multiple Data Views for Boxes in Containers.....	54
Configuring Tables.....	55
Use of the SyncPosition Property of PXGrid.....	55
Use of the DisplayMode Property of PXGridColumn.....	55
Use of the Type Property of PXGridColumn.....	55
Configuring Tabs.....	56
Conditional Hiding of a Tab Item.....	56
Configuring Boxes.....	57
Use of the CommitChanges Property of Boxes.....	57
Use of the DataField Property of PXGroupBox.....	58
Use of the Caption Property of PXGroupBox.....	58
Use of the RenderStyle Property of PXGroupBox.....	59
Configuring Layout and Size.....	59
Predefined Size Values.....	59
Use of the StartRow and StartColumn Properties of PXLayoutRule.....	60
Use of the ColumnWidth, ControlSize, and LabelsWidth Properties of PXLayoutRule.....	62
Use of the ColumnSpan Property of PXLayoutRule.....	63
Use of the Merge Property of PXLayoutRule.....	63
Use of the GroupCaption, StartGroup, and EndGroup Properties of PXLayoutRule.....	64
Use of the SuppressLabel Property of PXLayoutRule.....	65
Maintaining Reports.....	67
Display of Reports.....	67
Display of Analytical Reports.....	71

## **Accessing Data..... 73**

Querying Data in Acumatica Framework.....	73
BQL and LINQ.....	73
Data Access Classes.....	75
PXView and PXCache of the Data View.....	76
Data Query Execution.....	77
Translation of a BQL Command to SQL.....	80
Comparison of Fluent BQL, Traditional BQL, and LINQ.....	85
Fluent BQL and Traditional BQL Equivalents.....	86
To Execute BQL Statements.....	93
To Process the Result of the Execution of the BQL Statement.....	96
Creating Fluent BQL Queries.....	97
Fluent Business Query Language.....	97
Data Access Classes in Fluent BQL.....	99
Search and Select Commands and Data Views in Fluent BQL.....	100
Constants in Fluent BQL.....	102
Parameters in Fluent BQL.....	103
To Select Records by Using Fluent BQL.....	104
To Use Parameters in Fluent BQL Queries.....	107
Creating Traditional BQL Queries.....	111

Traditional Business Query Language.....	111
Data Access Classes in Traditional BQL.....	112
PXSelect Classes.....	112
The Classes That Compose BQL Statements.....	114
Parameters in Traditional BQL Statements.....	116
Traditional BQL and SQL Equivalents.....	117
To Select Records By Using Traditional BQL.....	119
To Filter Records.....	121
To Order Records.....	124
To Query Multiple Tables.....	125
To Group and Aggregate Records.....	127
To Use Parameters in Traditional BQL.....	127
To Use Arithmetic Operations.....	131
To Compose a BQL Statement from an SQL Statement.....	132
Creating LINQ Queries.....	137
LINQ in Acumatica Framework.....	137
Deferred LINQ Query Execution.....	139
Fallback to the LINQ to Objects Mode.....	139
To Select Records by Using LINQ.....	140
To Append LINQ Expressions to BQL Statements.....	143
Defining Relationships Between DACs.....	144
Establishment of a Relationship Between Two DACs.....	144
To Define a Primary Key.....	146
To Define a Foreign Key.....	147
Working with Data in Cache and Session.....	148
Session.....	148
Session Sharing Between Application Servers.....	149
Storing of Graph Data in the Session.....	151
Use of Slots to Cache Data Objects.....	153

## **Implementing Business Logic..... 158**

Working with Events.....	158
Event Handlers.....	158
Execution of Event Handlers.....	159
Data Manipulation Scenarios.....	161
Sequence of Events: Insertion of a Data Record.....	162
Sequence of Events: Update of a Data Record.....	164
Sequence of Events: Deletion of a Data Record.....	166
Sequence of Events: Display of a Data Record.....	168
Sequence of Events: Saving of Changes to the Database.....	168
List of Events.....	170
Working with Attributes.....	171
Code Reuse Through Attributes.....	171
Mandatory Attributes.....	173
Use of Attributes.....	174
Bound Field Data Types.....	174

Unbound Field Data Types.....	175
UI Field Configuration.....	176
Default Values.....	177
Complex Input Controls.....	178
Referential Integrity and Calculations.....	179
Ad Hoc SQL for Fields.....	179
Audit Fields.....	180
Data Projection.....	180
Access Control.....	180
Notes.....	180
Report Optimization.....	181
Attributes on DACs.....	181
Action Attributes.....	181
Attributes on Data Views.....	182
Working with Attachments.....	182
To Allow Attachments to a Particular Webpage.....	182
To Display an Attached Image on the Webpage.....	184
Configuring the UI from the Back End.....	185
To Configure an Input Mask and a Display Mask for a Field.....	185
To Display a Dialog Box.....	186
Executing Code Asynchronously.....	188
Asynchronous Execution.....	188
To Add a Button to the Processing Dialog Box.....	193
To Not Display the Processing Dialog Box.....	194
Localizing Applications.....	194
Localization.....	195
Strings That Can Be Localized.....	195
To Prepare DACs for Localization.....	196
To Localize Application Messages.....	196
To Work with Multi-Language Fields.....	197
To Optimize Memory Consumption of Localized Data.....	199
Reusing Business Logic.....	203
Reusable Business Logic Implementation.....	203
Mapped Cache Extensions and the Application Database.....	207
Reusable Business Logic and the Application Website.....	208
Use of Generic Graph Extensions by the System.....	210
Generic Graph Extensions Declared in Acumatica ERP.....	211
To Insert Reusable Business Logic That Has Already Been Declared.....	212
To Sort Multiple Generic Graph Extensions.....	214
To Implement Reusable Business Logic.....	214
<b>Troubleshooting Acumatica Framework-Based Applications.....</b>	<b>217</b>
To Debug Acumatica Framework-Based Applications.....	217
<b>Glossary.....</b>	<b>219</b>

# Copyright

---

© 2019 Acumatica, Inc.  
**ALL RIGHTS RESERVED.**

No part of this document may be reproduced, copied, or transmitted without the express prior consent of Acumatica, Inc.

11235 SE 6th Street, Suite 140  
Bellevue, WA 98004

## **Restricted Rights**

The product is provided with restricted rights. Use, duplication, or disclosure by the United States Government is subject to restrictions as set forth in the applicable License and Services Agreement and in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013 or subparagraphs (c)(1) and (c)(2) of the Commercial Computer Software-Restricted Rights at 48 CFR 52.227-19, as applicable.

## **Disclaimer**

Acumatica, Inc. makes no representations or warranties with respect to the contents or use of this document, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. Further, Acumatica, Inc. reserves the right to revise this document and make changes in its content at any time, without obligation to notify any person or entity of such revisions or changes.

## **Trademarks**

Acumatica is a registered trademark of Acumatica, Inc. HubSpot is a registered trademark of HubSpot, Inc. Microsoft Exchange and Microsoft Exchange Server are registered trademarks of Microsoft Corporation. All other product names and services herein are trademarks or service marks of their respective companies.

Software Version: 2019 R1

Last updated: March 14, 2019

# Acumatica Framework Guide

---

In this guide, you can find information about how to develop applications based on Acumatica Framework.

## In This Guide

- [\*Acumatica Framework Overview\*](#)
- [\*Getting Started with Acumatica Framework\*](#)
- [\*Designing the Application\*](#)
- [\*Configuring ASPX Webpages and Reports\*](#)
- [\*Accessing Data\*](#)
- [\*Implementing Business Logic\*](#)
- [\*Troubleshooting Acumatica Framework-Based Applications\*](#)
- [\*Glossary\*](#)

# Acumatica Framework Overview

---

Acumatica Framework provides the application programming interface (API) and tools for developing cloud business applications. Acumatica Framework is a part of the Acumatica Cloud xRP Platform, which provides various opportunities for developing add-on applications that interact with Acumatica ERP through the web services API, applications embedded into Acumatica ERP through the built-in customization tools, and completely new applications based purely on Acumatica Framework.

In this part of the guide, you can find an overview of the Acumatica Cloud xRP Platform and the place of Acumatica Framework in this platform. This part also includes an overview of Acumatica Framework tools and a high-level overview of the runtime architecture of applications based on Acumatica Framework.

## In This Part

- [Acumatica Cloud xRP Platform](#)
- [Acumatica Framework Development Tools](#)
- [Runtime Architecture of an Application Based on Acumatica Framework](#)

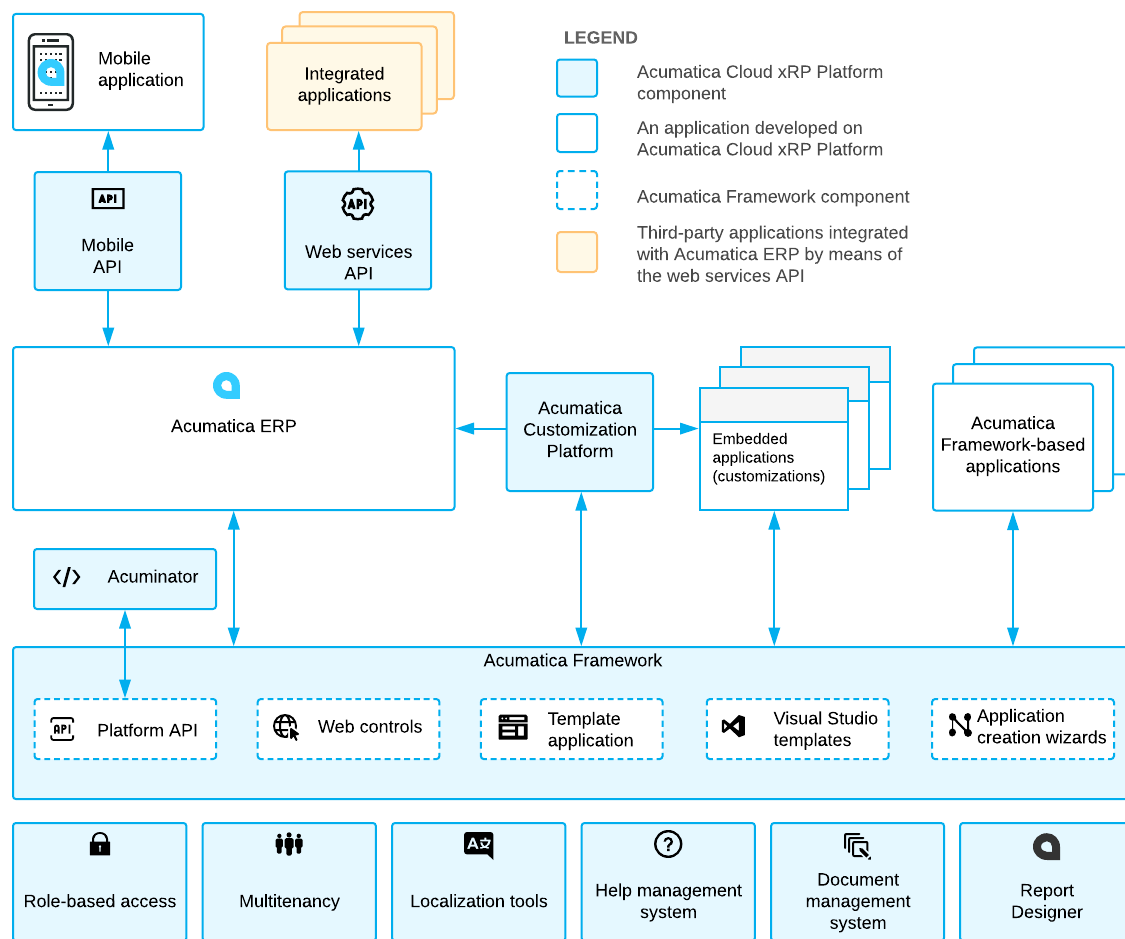
## Acumatica Cloud xRP Platform

---

The Acumatica Cloud xRP Platform is the platform provided by Acumatica that is used to build the Acumatica ERP application itself, any customizations of Acumatica ERP, the mobile application for Acumatica ERP, and applications integrated with Acumatica ERP through the web services API.

The Acumatica Cloud xRP Platform consists of a number of components, which are highlighted with light blue in the following diagram. These components serve different purposes, which are described in detail in this topic, and can be used either separately or combined to achieve your business purposes.





**Figure: Acumatica Cloud xRP Platform**

### Basic Components and Tools

The base of the Acumatica Cloud xRP Platform is formed by the components and tools that provide the basic application functionality, such as multitenancy support, role-based access, and localization tools. These components and tools are available out-of-the-box in Acumatica ERP, any embedded in Acumatica ERP applications, or applications based purely on Acumatica Framework applications. This means that you do not need to worry about implementing mechanisms similar to these components during the design or programming of your application based on the Acumatica Cloud xRP Platform.

Acumatica Cloud xRP Platform contains the basic components and tools listed in the following table.

Component or Tool	Description
Role-based access	A set of components responsible for user authorization, access rights verification, and audit on the data access and business logic levels. For more information, see <a href="#">Role-Based Access</a> in the System Administration Guide.
Multitenancy	A component responsible for hosting multiple tenants on a single application server. For details about multitenancy, see <a href="#">Support of Multiple Tenants</a> in the User Guide.
Localization tools	The tools that help you to perform the localization of the application to multiple languages. For more information about localization, see <a href="#">Translation Process</a> .

Component or Tool	Description
Help management system	The integrated wiki-based help content editing, management, and search system. For details about the help management system, see <a href="#">Wiki Overview</a> .
Document management system	The integrated document storage and management system. For details, see <a href="#">Managing Attached Files</a> .
Report Designer	A separate utility (which can be installed along with Acumatica ERP or Acumatica Framework) that you can use to design custom reports. For details on this tool, see <a href="#">Acumatica Report Designer Guide</a> .

## Acumatica Framework

Acumatica Framework provides the platform API, web controls, and other tools for the development of the UI and business logic of an ERP application. The platform API is used for the development of Acumatica ERP and any embedded applications (that is, customizations of Acumatica ERP). Acumatica Framework can also be used to develop an ERP application from scratch. You can find detailed information about the development of applications with Acumatica Framework in this guide.

Acumatica Framework consists of the development tools listed in the following table.

Component	Description
Platform API	The API for implementing application business logic.
Web controls	A set of web controls implementing access to business logic through the web UI.
Template application	The application and database structures providing frameset, layout, and navigation services. You can deploy a template application along with the template website by using the Acumatica Framework Configuration Wizard.
Visual Studio templates	The project template for the creation of a new application and a set of page templates that automate the creation of typical page layouts.
Application creation wizards	A set of components to automate the creation of the application data access classes from the database tables and the web forms during application development.

For details about the development tools, see [Acumatica Framework Development Tools](#).

## Acuminator

Acuminator is a static code analysis and colorizer tool for Visual Studio that simplifies development with Acumatica Framework. Acuminator provides diagnostics and code fixes for common developer challenges related to the platform API. Also, Acuminator can colorize and format business query language (BQL) statements, and can collapse attributes and parts of BQL queries. You can find related information and download Acuminator at [Visual Studio Marketplace](#).

## Acumatica Customization Platform

Acumatica Customization Platform provides customization tools for the development of applications embedded in Acumatica ERP. Developers that work with Acumatica Customization Platform use the platform API provided by Acumatica Framework.

With Acumatica Customization Platform, you can perform end-customer customizations and create complex solutions for multiple customers. In these customizations, you can modify the user interface, business logic, and database schema without recompilation and re-installation of the application. Customizations are stored separately from the core application code as metadata and can be modified,

exported, or imported. Because customizations are stored separately, they are preserved with the updates and upgrades of the core application.

For details on Acumatica Customization Platform, see [Acumatica Customization Platform](#).

### Web Services APIs

The Acumatica Cloud xRP Platform provides multiple types of web services APIs for development of applications integrated with Acumatica ERP. These applications can perform data migration and data import, integration of Acumatica ERP with external systems, and execution of long-running operations.

You can use the contract-based REST API, contract-based SOAP API, or screen-based SOAP API to access the same business logic as is accessed in the UI. All types of the web services APIs can be used with any customization applied to Acumatica ERP. The contract-based REST API supports the OpenAPI 2.0 (formerly known as Swagger 2.0) specification.

For details on the web services APIs, see [Contract-Based Web Services API](#) and [Screen-Based Web Services API](#).

Acumatica ERP supports the OAuth 2.0 mechanism of authorization for add-on applications that interact with Acumatica ERP through application programming interfaces (APIs). For details on the authorization of applications, see [Authorizing Client Applications to Work with Acumatica ERP](#).

### Mobile API

Acumatica ERP provides the Acumatica mobile application, which allows a user to work with Acumatica ERP through the mobile devices. You can customize the mobile application by using the mobile API. For details on the mobile API, see [Working with Mobile Framework](#).

## Acumatica Framework Development Tools

---

In this topic, you can find a detailed description of the development tools provided with Acumatica Framework.

### Platform API

The Platform API provided with Acumatica Framework is an event-driven programming API, which is traditional in rich GUI applications. This model covers database access, business logic, GUI behavior, and error handling. All coding is done with only C#.

The following code gives an example of the business logic implemented in the business logic controller: The code updates the receipt total when one of the transactions related to the receipt is updated.

```
public virtual void DocTransation_RowUpdated(PXCache cache,
                                             PXRowUpdatedEventArgs e)
{
    DocTransaction old = e.OldRow as DocTransaction;
    DocTransaction trn = e.Row as DocTransaction;
    if ((trn != null) && (trn.TranQty != old.TranQty ||
                        trn.UnitPrice != old.UnitPrice))
    {
        Document doc = Receipts.Current;
        if (doc != null)
        {
            doc.TotalAmt -= old.TranQty * old.UnitPrice;
            doc.TotalAmt += trn.TranQty * trn.UnitPrice;
            Receipts.Update(doc);
        }
    }
}
```

When a user selects a document transaction in the table on a webpage and updates the settings of the transaction, the `RowUpdated` event is triggered, and the code above is executed and updates the receipt total, as shown in the following screenshots.

**Figure: Update of document transaction details**

## Visual Web Designer Support

The Acumatica Framework Integrated Development Environment (IDE) is built on top of Microsoft Visual Studio. However, the Acumatica Framework IDE implements its own set of web controls to generate an advanced GUI in a web browser.

All of Acumatica Framework's web controls have the same rendering and a similar appearance in design mode in the IDE and runtime mode in a web browser. Thus, the developer can utilize all the facilities of the Visual Web Designer component of Visual Studio. The application developer can use the convenient drag-and-drop mechanism to create an application form layout, to perform form visual editing, and to set a control's properties and behavior through an intuitive graphical interface. This approach does not require any knowledge of HTML or Java Script, yet allows the developer to create a professional and appealing web GUI.

The following screenshots illustrate the design (left) and runtime (right) rendering.

**Figure: Webpage in design and runtime mode**

## Master Pages, Templates, and CSS Support

The Visual Studio project and item templates provide reusable and customizable project and item stubs that accelerate the development process, removing the need to create new projects and items from scratch. Project templates provide the basic files needed for a particular project type, include standard assembly references, and set the default project properties and compiler options.

Acumatica Framework distribution includes the following:

- The project template for the creation of a new application
- A set of page templates that automate the creation of typical page layouts

The master pages mechanism in ASP.NET allows for the creation of an application that looks and feels consistent. Master pages define the standard appearance and behavior that is common in all application pages. You create individual content pages that refer to the master page. When a content page is requested, it merges with the master page to produce output that combines the layout and base functionality of the master page with the content of the requested page.

Acumatica Framework fully supports the master pages mechanism and provides you with a set of predefined master pages. You can design your own master pages or modify existing ones.

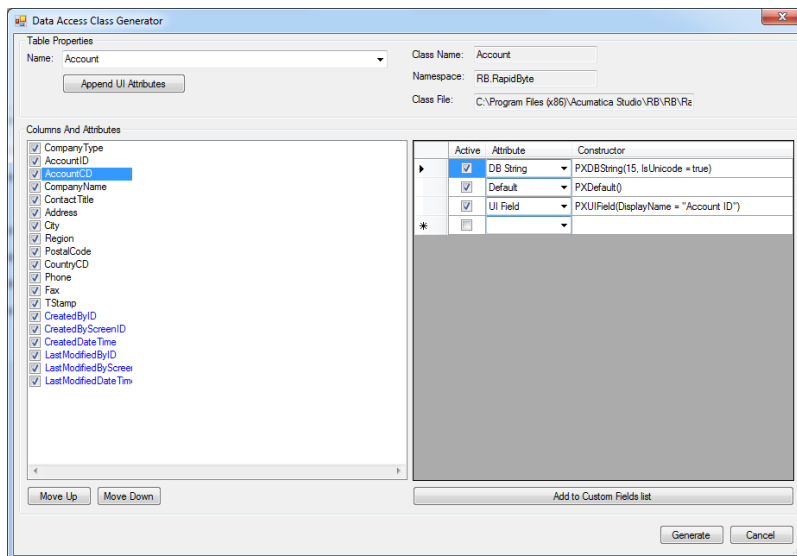
A web application written with Acumatica Framework supports style modification through Cascading Style Sheets (CSS).

## Application Creation Wizards

Acumatica Framework provides a set of wizards for automating the creation of data access classes and webpages. The use of these wizards eliminates the manual steps associated with data access class creation and data binding configuration.

The Data Access Class Generator, which is shown in the following screenshot, provides the application developer with a way to create and modify data access classes. It implements the following functionality:

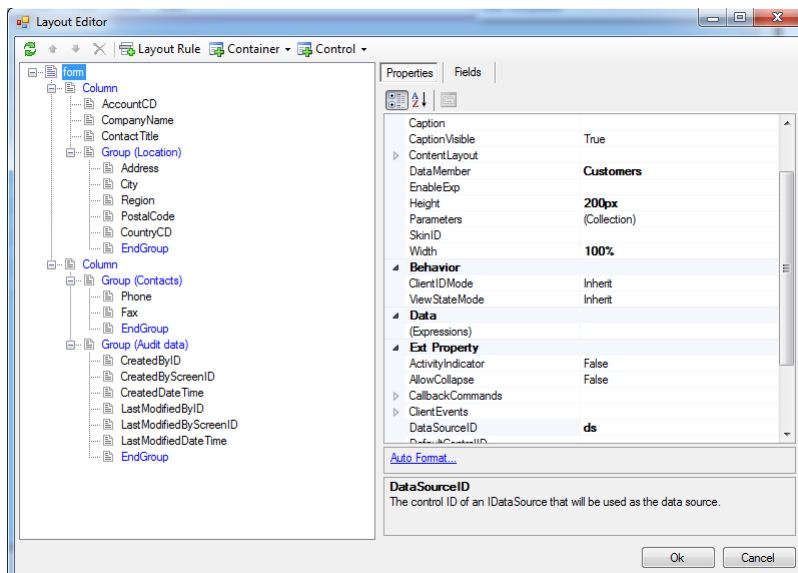
- Reading the data structure from a table, SQL query, or external data source
- Creating a data access class based on the data structure received from external data source
- Reading the data access class structure from its definition and merging this structure with the data structure received from the external data source
- Automatic mapping of application-specific attributes based on the names of the properties of the external data source



**Figure: Data Access Class Generator**

The Layout Editor, which is shown in the following screenshot, automates the creation of new web forms. It uses metadata stored in the business logic controller and data access class to help the application developer create new web forms or to modify existing ones. The **Layout Editor** wizard implements the following features:

- Reading metadata from the business logic controller and the data access class and creating a list of controls that could be added to the webpage
- Adding the controls selected by the programmer to the webpage
- Updating the webpage controls with the changed business logic controller and the data access class metadata

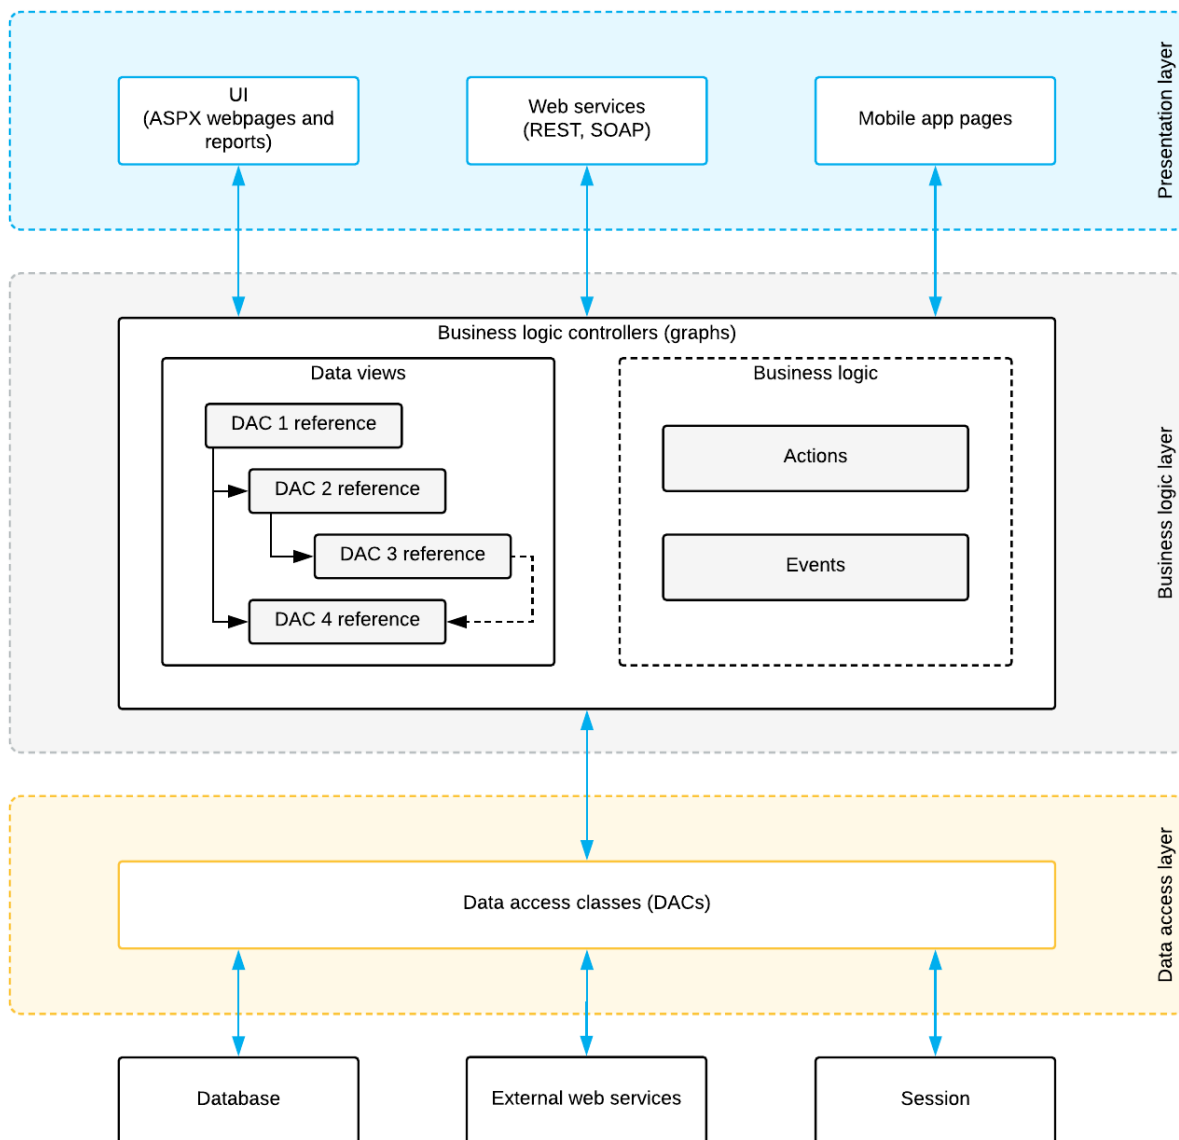


**Figure: Layout Editor**

## Runtime Architecture of an Application Based on Acumatica Framework

In this topic, you can review the architecture of an application created based on Acumatica Framework, such as Acumatica ERP, customizations of Acumatica ERP, and applications based purely on Acumatica Framework.

An application written with Acumatica Framework has  $n$ -tier architecture with a clear separation of the presentation, business, and data access layers, as shown in the following diagram. You can find details about each layer in the sections below.



**Figure: Application architecture**

## Data Access Layer

The data access layer is implemented as a set of data access classes (DACs) that wrap data from database tables or data received through other external sources (such as Amazon Web Services). You can generate a data access class associated with a database table with the help of the Data Access Class Generator, which reads database metadata. When you use it, you select a table and specify the columns that should be reflected in the data access class.

The instances of data access classes are maintained by the business logic layer. Between requests, these instances are stored in the session. On a standalone Acumatica ERP server, session data is stored in the server memory. In a cluster of application servers, session data is serialized and stored in a high-performance remote server through a custom optimized serialization mechanism. For details about data storage in session, see [Session](#).

For details on working with the data access layer, see [Accessing Data](#).

## Business Logic Layer

The business logic is implemented through the business logic controller (also called *graph*). Business logic controllers are classes that you derive from the special API class (`PXGraph`) and that are tied to one or more data access classes.

Each business logic controller conceptually consists of two parts:

- Data views, which include the references to the required data access classes, their relationships, and other meta information
- Business logic, which consists of actions and events associated with the modified data.

Each business logic controller can be accessed from the presentation layer or from the application code that is implemented within another business logic controller. When the business logic controller receives an execution request, it extracts the data required for request execution from the data access classes included in the data views, triggers business logic execution, returns the result of the execution to the requesting party, and updates the data access classes instances with the modified data.

For details on working with the business logic layer, see [Implementing Business Logic](#).

## Presentation Layer

The presentation layer provides access to the application business logic through the UI, web services, and Acumatica mobile application. The presentation layer is completely declarative and contains no business logic.

The UI consists of ASPX webpages (which are based on the ASP.NET Web Forms technology) and reports created with Acumatica Report Designer. The ASPX webpages are bound to particular business logic controllers. You create webpages from the templates provided with Acumatica Framework and modify them with the help of the Layout Editor, which utilizes metadata information extracted from the business logic controller. Reports created with Acumatica Report Designer are loaded and executed by using the predefined business logic controller included in Acumatica Framework. For details about processing of the reports in the system, see [Display of Reports](#).

When the user requests a new webpage, the presentation layer is responsible for processing this request. Webpages are used for generating static HTML page content and providing additional service information required for the dynamic configuration of the web controls. When the user receives the requested page and starts browsing or entering data, the presentation layer is responsible for handling asynchronous HTTP requests. During processing, the presentation layer submits a request to the business logic layer for execution. Once execution is completed, the business logic layer analyzes any changes in the business logic container state and generates the response that is sent back to the browser as an XML document.

For details on the configuration of ASPX webpages, see [Configuring ASPX Webpages and Reports](#).

Web services and mobile app pages provide alternative interfaces to the application business logic. From the side of the business logic controller, a request from a webpage, the web services, or an mobile app page are identical and, thus, cause the execution of exactly the same business logic.



# Getting Started with Acumatica Framework

---

In this part of the guide, you can find the information that you may need to start development with Acumatica Framework.

## System Requirements

You can find the full list of system requirements in [System Requirements for Acumatica Framework 2019 R1](#) in the Installation Guide.

## Installation

For detailed information about the installation of Acumatica Framework, see [Installing Acumatica Framework](#) in the Installation Guide.

## Application Design

For the information about the design of the database structure and user interface of applications based on Acumatica Framework, see [Designing the Application](#) in this guide.

## Development of the Application Code

Before you begin developing application code, we recommend that you complete the following training courses:

- [T100 Development: Introduction to Acumatica Framework](#)
- [T200 Development: Acumatica Framework Fundamentals](#)

For a quick overview of application programming, refer to the topics in this part of the guide.

In the Acumatica Framework Guide, you can find reference information and additional information that is not covered in the training courses. This information is provided in the following parts of the guide:

- [Configuring ASPX Webpages and Reports](#): About the development of ASPX pages
- [Accessing Data](#): About business query language (BQL) and working with data in cache and session
- [Implementing Business Logic](#): About events, attributes, long-running operations, and other topics related to business logic development
- [Troubleshooting Acumatica Framework-Based Applications](#): About debugging the Acumatica Framework-based applications and fixing the common errors

For a detailed description of the Acumatica Framework API, see [API Reference](#).

## Website Management

For the classic UI, if you want to adjust the position of a webpage in the site map, add a new page to the site map, or remove a page from the site map, you configure the site map of the website, as described in [Managing the Site Map](#) in the System Administration Guide.

If you want to modify the position of a webpage in the modern UI, add a page to a workspace, or remove a page in the modern UI, you configure the modern UI as described in [Customizing the Modern User Interface](#) in the System Administration Guide.

You need to grant access rights to each new page. For details on the configuration of access rights, see [Managing User Access Rights](#) in the System Administration Guide.

You can create help topics for any application you have developed with Acumatica Framework by using the built-in wiki-based content management system. For details on creating help topics, see [Managing Wikis](#).

### In This Part

- [Data Querying](#)
- [Business Logic Controller Declaration](#)
- [Data View and Cache](#)
- [Data Modification Scenarios](#)
- [Business Logic Implementation](#)

## Data Querying

This system implements a custom language for writing database queries called BQL (business query language). It is not LINQ and doesn't use it. BQL is written in C# and based on generic classes syntax, but still is very similar to SQL syntax. It has almost the same keywords placed in the order they are used in SQL, as shown in the following example.

```
PXSelect<Product,
    Where<Product.availQty, IsNotNull,
        And<Product.availQty, Greater<Product.bookedQty>>>>>
```

If the database provider is MICROSOFT SQL Server, the framework translates this expression into the following SQL query.

```
SELECT * FROM Product
WHERE Product.AvailQty IS NOT NULL
AND Product.AvailQty > Product.BookedQty
```

BQL gives several benefits to the application developer. It does not depend on database-provider specifics, is object-oriented and extendable. An important benefit is compile-time syntax validation, which helps to prevent SQL syntax errors.

Since BQL is implemented on top of generic classes, you need types that would represent database tables. In the context of Acumatica Framework, they are called data access classes (DACs).

For example, to execute the SQL query from the example above, you should define the `Product` data access class as shown in the following code fragment.

```
using System;
using PX.Data;

// Types used in BQL statements should derive from special interfaces:
// table - IBqlTable, column - IBqlField.
[System.SerializableAttribute()]
public class Product : PX.Data.IBqlTable
{
    // The type used in BQL statements to reference the ProductID column
    public abstract class productID : PX.Data.IBqlField
    {
    }
    // The property holding ProductID value in a record
    [PXDBIdentity(IsKey = true)]
    public virtual int? ProductID { get; set; }

    // The type used in BQL statements to reference the AvailQty column
    public abstract class availQty : PX.Data.IBqlField
    {
    }
}
```

```

// The property holding AvailQty value in a record
[PXDBDecimal(2)]
public virtual decimal? AvailQty { get; set; }

// The type used in BQL statements to reference the BookedQty column
public abstract class bookedQty : PX.Data.IBqlField
{
}
// The property holding BookedQty value in a record
[PXDBDecimal(2)]
public virtual decimal? BookedQty { get; set; }
}

```

Each table field is declared in a data access class twice:

- A public abstract class (which is also referred to as *class field*) to reference a field in the BQL command
- A public virtual property (which is also referred to as *property field*) to hold the table field data

If the DAC is bound to the database, it must have the same class name as the database table.

Fields are bound to the database by means of data mapping attributes (such as `PXDBIdentity` and `PXDBDecimal`), using the same naming convention.

The code below demonstrates an example of how to get data records from the database.

```

// Select Product records
PXResultset<Product> res = PXSelect<Product, Where<Product.availQty, IsNotNull,
    And<Product.availQty, Greater<Product.bookedQty>>>>
    .Select(new PXGraph());
// You can iterate through the result set
foreach(PXResult<Product> rec in res)
{
    // A record from the result set can be cast to the DAC
    Product p = (Product)rec;
    ...
}

```

## Business Logic Controller Declaration

Working with the business data in Acumatica Framework is implemented through the *business logic controller* object also referred as *graph* (graph is a mathematical term for a set of objects where some pairs of objects are connected by links). A graph provides the interface for the presentation logic to operate with the business data and relies on Data Access Layer components to store and retrieve the business data from the database.

The following example shows the declaration of a simple business logic controller.

```

//Declaration of the graph
public class ProductMaint : PXGraph<ProductMaint>
{
    //Declaration of the data view
    public PXSelect<Product> Products;

    //Declaration of the actions
    public PXCancel<Product> Cancel;
    public PXSave<Product> Save;
}

```

In this example, the graph implements the following interfaces:

- **Products** – the *data view* that can be used for querying and modifying the data

- **Cancel** – the *action* that discard all the changes made to the data and reloads it from the database
- **Save** – the *action* that commits the changes made to the data to the database and then reloads the committed data

## Data View and Cache

Data views implement the interfaces for querying the data from the database and submitting modified data to the cache.

Data views are declared in business logic controllers as public fields of `PXSelect` command type, as shown in the following code.

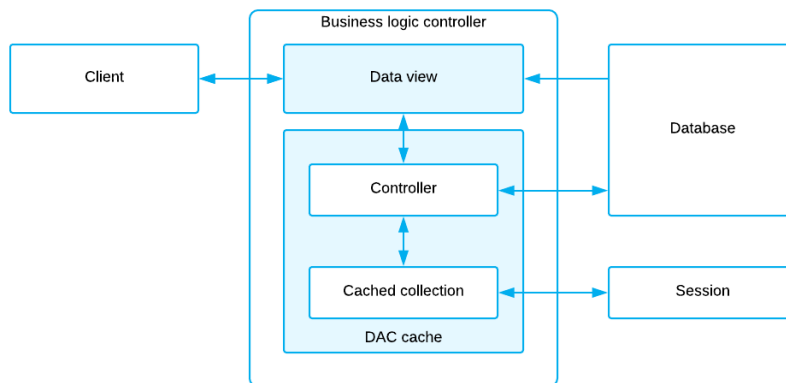
```
public PXSelect<Product> Products;
```

Based on this declaration, the system automatically instantiates the DAC cache.

A *DAC cache* object in the Acumatica Framework is the primary interface for working with individual records from the graph business logic. It has two components and two primary responsibilities:

- The *Cached* collection: In-memory cache that contains modified entity records. The *Cached* collection is instantiated based on the corresponding DAC declaration and managed by the cache.
- The controller: The cache component that implements basic CRUD (create, read, update, delete) operations on the *Cached* collection and triggers a sequence of data manipulation events when modifying or accessing the data in the *Cached* collection. These events can be later subscribed from the graph to implement the business logic associated with the data modification.

The diagram below shows the internal graph structure and responsibilities of the data view and the cache.



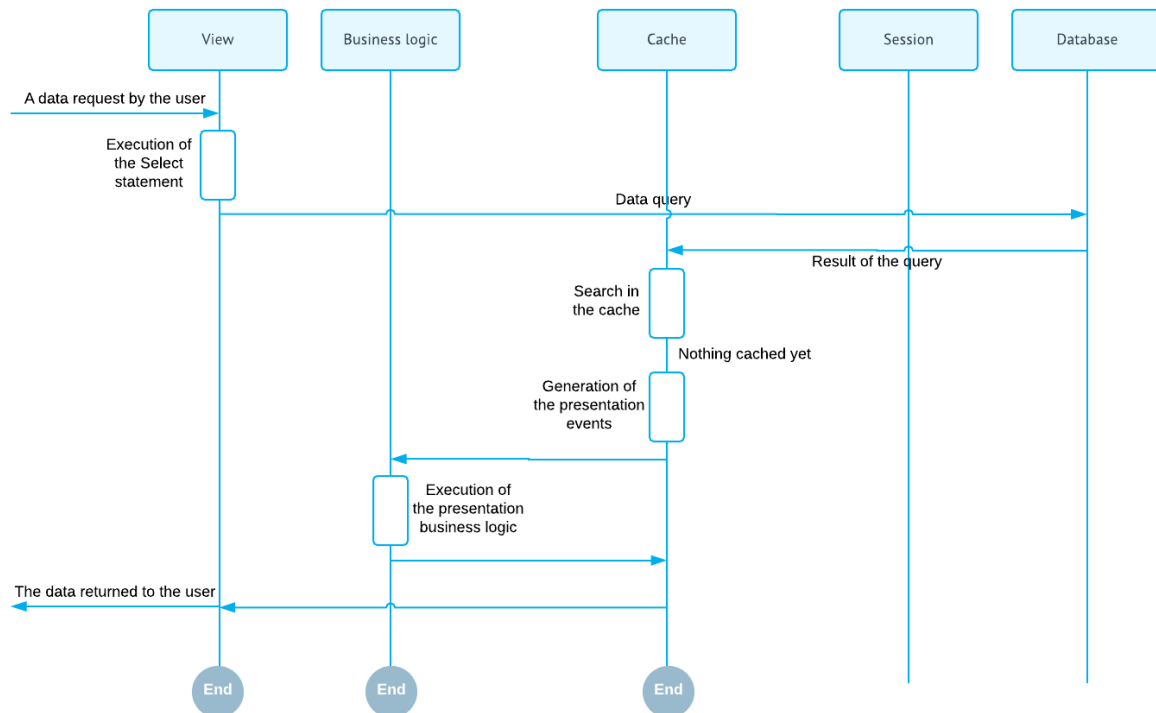
**Figure: The graph structure**

## Data Modification Scenarios

In this topic, you can find the basic data manipulation scenarios that can be executed from the graph business logic or from the user interface. Entity data manipulation through the user interface indirectly invokes the same methods as the direct call from the business logic controller.

## Querying the Data for the First Time

The data can be requested through the `Select` method of the data view. During this operation, the system executes BQL command from the data view declaration. The data returned by the BQL command is passed to the requester. The following diagram illustrates this process.

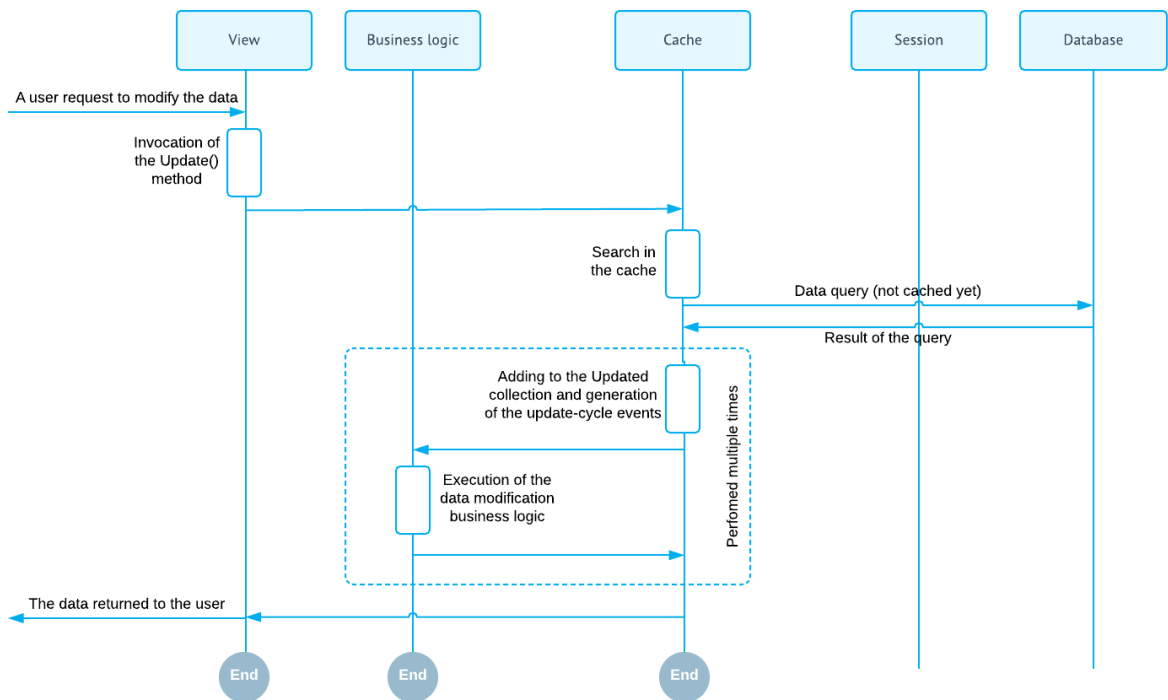


**Figure: Querying the data for the first time**

## Updating an Existing Record

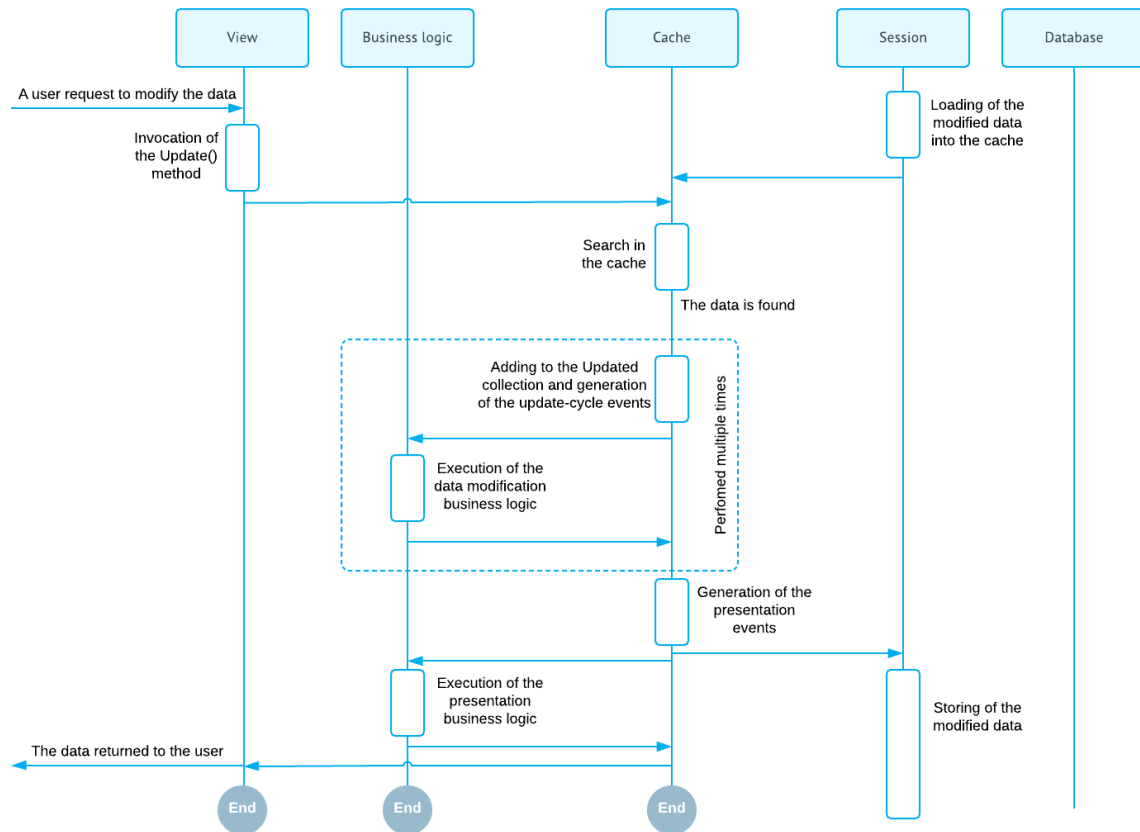
An existing record can be updated through the `Update(record)` method of the data view. This method places the modified record into the cache.

If the data record is not found in the `Cached` collection, the cache controller loads the data record from the database, adds it to the `Cached` collection, marks it as updated, and updates it with the new values. The search of the data record in the `Cached` collection and loading of the data record from the database is based on the DAC key fields. The diagram below illustrates this scenario.



**Figure: Updating the record for the first time**

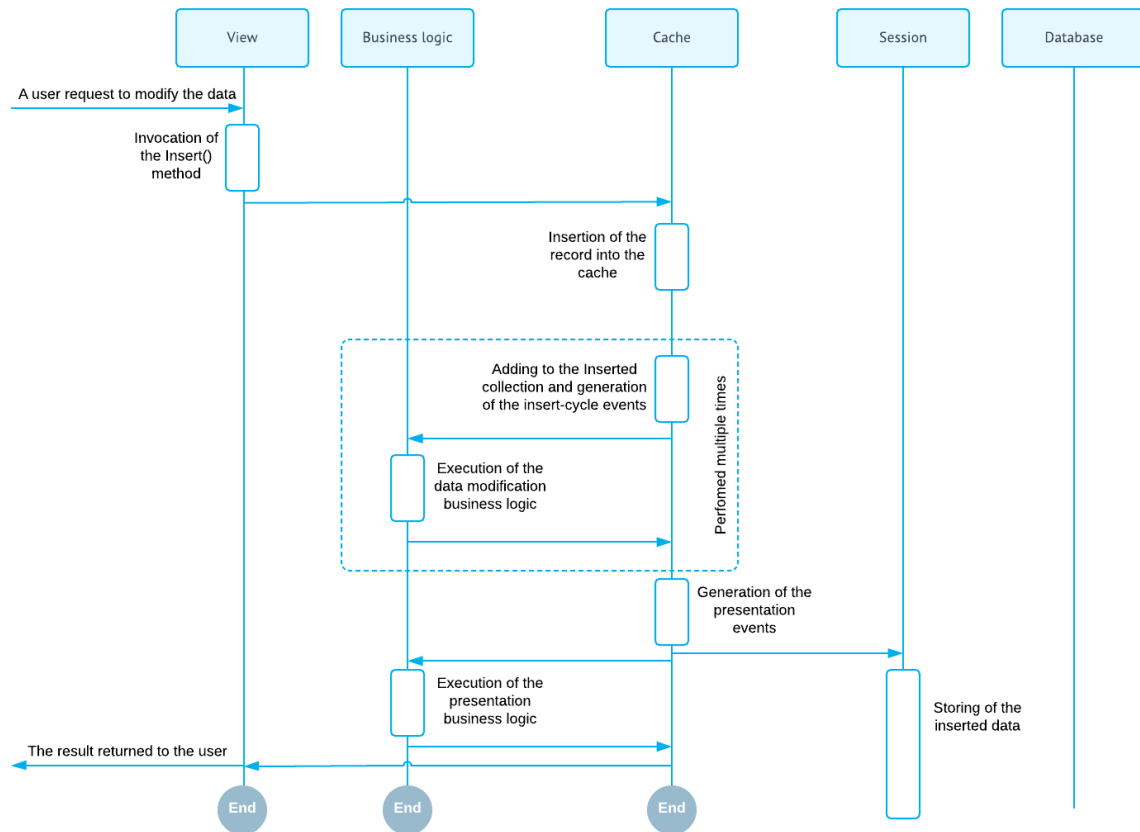
If the updated record exists in the `Cached` collection the cache controller locates it and updates it with the new values. The diagram below illustrates this scenario.



**Figure: Updating the cached (previously modified) record**

### Inserting a New Record

A new record can be inserted into the cache through the `Insert(record)` method of the data view. The new inserted record is added to the `Cached` collection and marked as inserted. The diagram below illustrates this scenario.



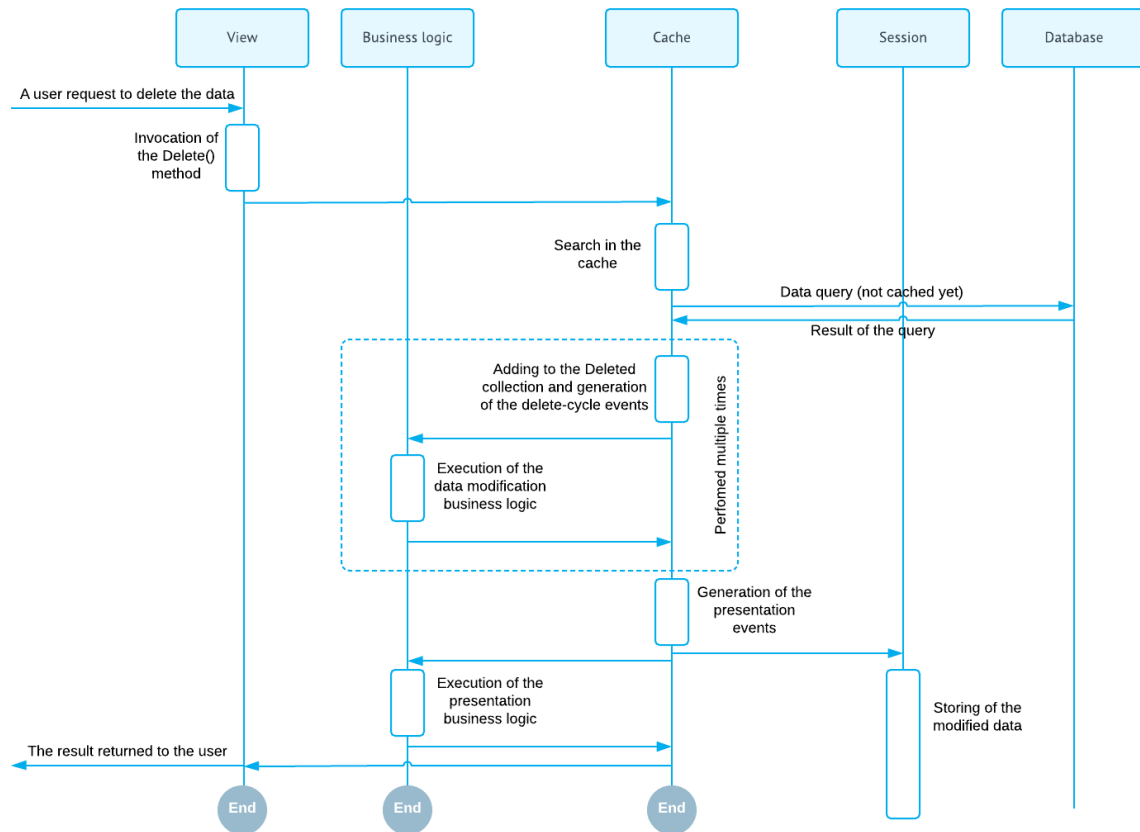
**Figure: Inserting the new record**

### Deleting an Existing Record

An existing record can be deleted from the cache using the `Delete(record)` method, of the data view.

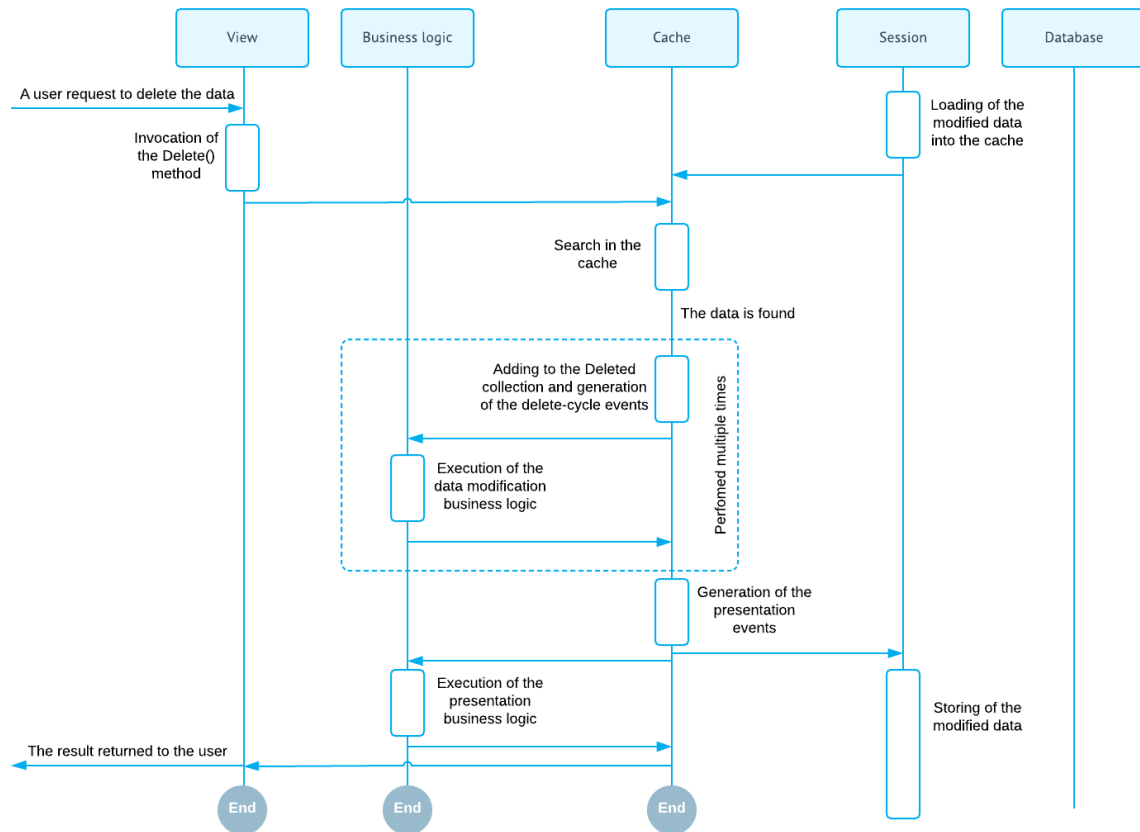
If the data record is not found in the `Cached` collection, the cache controller loads the data record from the database, adds it to the `Cached` collection, and marks it as deleted. The search of the data record in the `Cached` collection and loading of the data record from the database is based on the DAC key fields. The diagram below illustrates this scenario.





**Figure: Deleting the non-cached (unmodified) record**

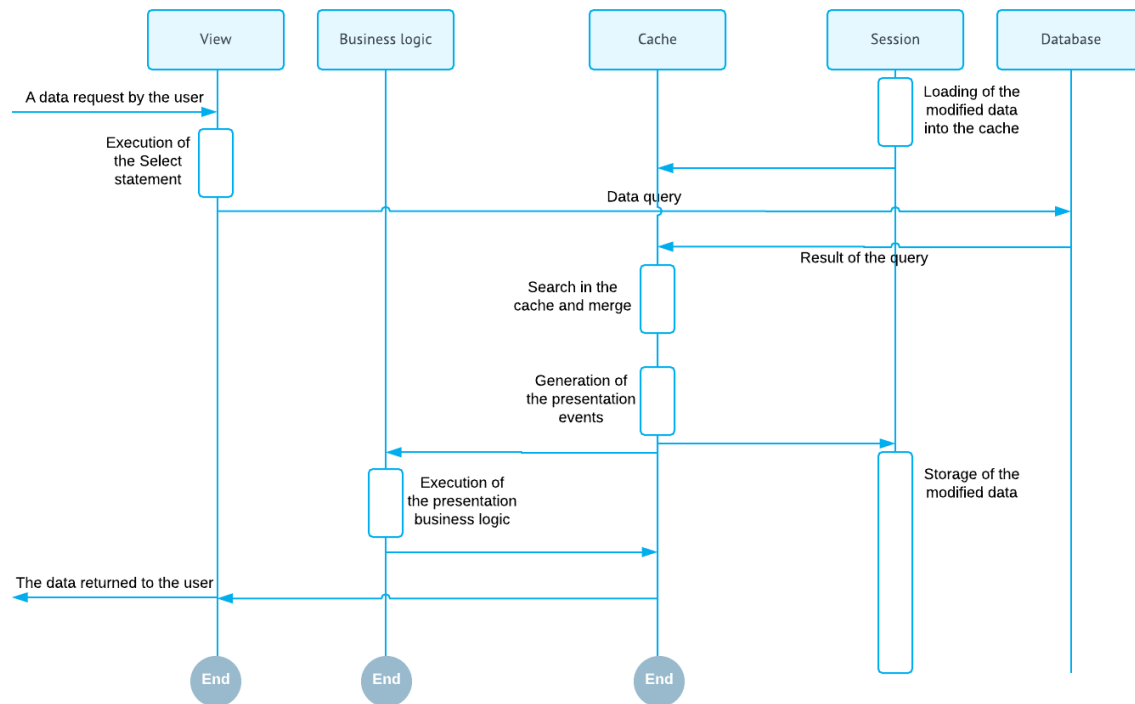
If the deleted record is found in the `Cached` collection, the cache controller locates it and marks as deleted. The diagram below illustrates this scenario.



**Figure: Deleting of the cached (previously modified) record**

### Querying Updated Data

The data can be modified and then queried again. In this scenario, the data records stored in the cache memory are merged with the result of the BQL command execution. Data record merge is based on DAC key fields. The final result of the `Select()` execution incorporates all the earlier record modifications that have not been preserved to the database yet. The diagram below illustrates this scenario.



**Figure: Querying the modified data**

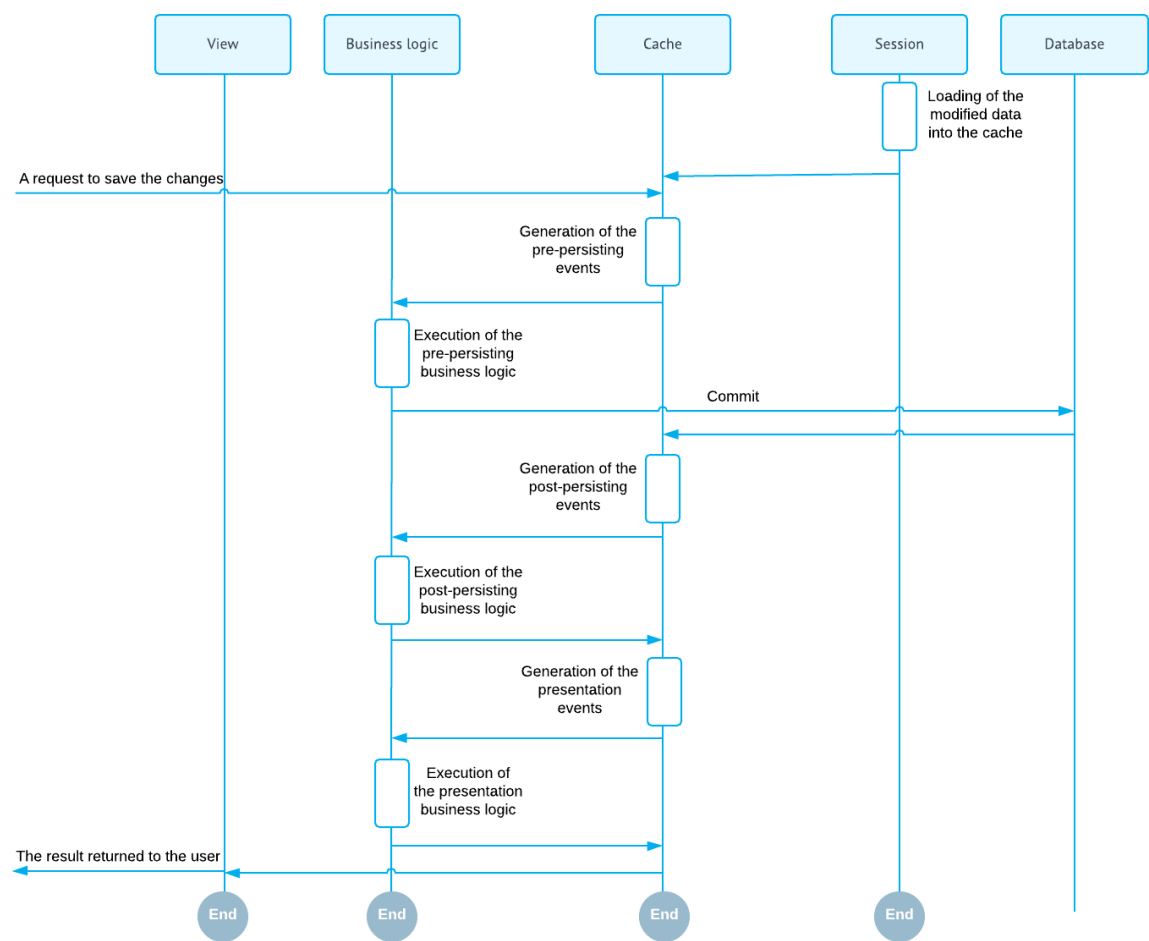
### Persisting Changes to the Database

When the data is modified, the system has two different versions of the data: the new one stored in the caches memory and the original one persisted in the database. At this point you have two options:

- Save the new version of data to the database using the `Persist()` method of the graph
- Discard all in-memory changes and load the original data version using the `Clear()` method of the graph

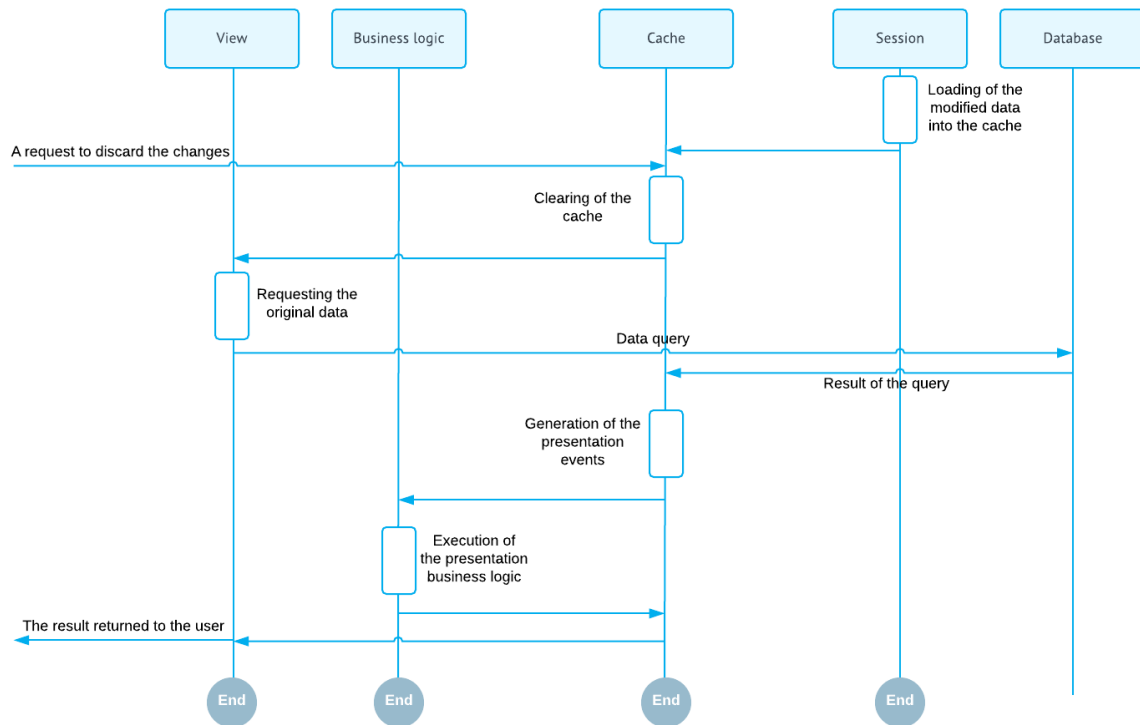
From the user interface these methods are called by invocation of the `Save` and `Cancel` actions. These actions are predefined and mapped to the `Persist()` and `Clear()` methods.

The diagram below illustrated saving of the changes to the database.



**Figure: Saving the changes to the database**

The diagram below illustrates discarding of all in-memory entity changes.



**Figure: Discarding the changes and loading the original data**

### Preserving the Data Version Between the Round Trips and Handling the Subsequent Selects from the Views

It is important to understand that a graph is a stateless object. It is discarded after each data request. In order to preserve the modified data version between the requests, the cache controller serializes the `Cached` collection into the session state and restores it later when the graph is instantiated on the subsequent request. In this scenario, it is very important that the cache contains only the modified entity records, not the complete entity record set.

## Business Logic Implementation

Business logic is implemented by overloading certain methods invoked by the system in the process of manipulating data. For such procedures as inserting a data record or updating a data record, the `PXCache` controllers generate series of events causing invocation of the methods called event handlers. The application is able to interfere in the series of events on different stages. For this purpose, the application implements methods that are executed as event handlers.

Business logic can be divided into common logic relevant to different parts of the application and the logic specific to an application screen (webpage). The common logic is implemented through event handler methods defined in attributes, while the screen-specific logic is implemented as methods in the associated graph.

### Common Business Logic

You implement the common business logic by defining event handlers in attributes. If such attribute is added to the declaration of a data access class, attribute logic is applied to the data records of this type for any graph used to access this table.

There are a number of predefined attributes implemented in the framework. For example, in the following declaration of a data field for a column, the `PXDBDecimal` attribute binds this field to a database column of the decimal type.

```
[PXDBDecimal(2)]
public virtual string AvailQty { get; set; }
```

The attributes that bind a field to a specific data type exist for most database data types.

Another typical example of an attribute is `PXUIField`. It is used to configure the input control for the column in the user interface. This allows having the same visual representation of the column on all application screens (unless a screen redefines it). The following example shows the use of the `PXUIField` attribute.

```
[PXDBDecimal(2)]
[PXUIField(DisplayName = "Available Qty", Enabled = false)]
public virtual string AvailQty { get; set; }
```

You can also defined your own attributes, as shown in the following code.

```
// Application-defined attribute that implements common business logic
public class MyAttribute : PXEventSubscriberAttribute,
                        IPXEventNameSubscriber
{
    // An event handler
    protected virtual void EventName(PXCache sender,
                                     PXRowEventNameEventArgs e)
    {
        ...
    }
    ...
}
```

These custom attributes can also be added to the DAC declaration, as shown in the following example.

```
[PXDBDecimal(2)]
[PXUIField(DisplayName = "Available Qty", Enabled = false)]
[MyAttribute]
public virtual string AvailQty { get; set; }
```

For details about attributes, see [Working with Attributes](#).

## Screen-Specific Business Logic

For a specific screen, the application can redefine the common logic or extend it. For this purpose, you should define event handlers in the graph associated with the screen. Each event handler method is tied to a particular table or a table field via the naming convention.

For example, you can verify a value of a column as shown in the following code.

```
public class ProductRecalc : PXGraph<ProductRecalc>
{
    ...
    // Event handler verifying that the value of the AvailQty column
    // in Product records is greater than 0.
    // It is triggered when, for instance, a Product record is updated.
    protected virtual void Product_AvailQty_FieldVerifying(
        PXCache sender,
        PXFieldVerifyingEventArgs e)
    {
        Product p = (Product)e.Row;
        if (p != null && p.AvailQty != null)
        {
            if (p.AvailQty < 0)
                throw new PXSetPropertyException<Product.availQty>(
```

```
        "Value must be greater than 0.");  
    }  
}
```

For details about events, see [Working with Events](#).

# Designing the Application

---

This part of the guide contains the design guidelines for the database schema and applications built on Acumatica Framework.

## In This Part

- [Designing the Database Structure and DACs](#)
- [Designing the User Interface](#)
- [Naming the Graphs and Event Handlers](#)

## Designing the Database Structure and DACs

---

This chapter covers the main aspects of database design used in Acumatica Framework.

## In This Chapter

- [System and Application Tables](#)
- [Table and Column Naming Conventions](#)
- [Common Columns and Data Types](#)
- [Primary Key](#)
- [Foreign Keys and Nullable Columns](#)
- [Audit Fields](#)
- [Concurrent Update Control](#)
- [Attachment of Additional Objects to Data Records](#)
- [Preservation of Deleted Records](#)
- [Multitenancy Support](#)

## System and Application Tables

The database of your Acumatica Framework-based application consists of the following tables:

- System tables: Those that are created by default for the application template and not used to store your application data
- Application tables: Acumatica ERP tables (which exist if you have implemented customization) and your own tables

Do not add columns to system tables or modify them in any other way. Such modifications could corrupt the application and would be lost during the next database upgrade.

Regarding your own application tables, you have to design and create the needed tables that store your application data. You then map these application tables to data access classes (DACs) that define the object model of the application. In one table, you can keep data records of multiple entities, each of which is defined as a separate data access class in the application object model.

## Table and Column Naming Conventions

In this topic, you can learn how you should name tables and columns in a database that is used by an Acumatica Framework-based application.



## Table Naming Conventions

When you are creating a table, you should consider the following suggestions regarding naming conventions:

- Make sure that table and column names are valid C# identifiers, because these names match the names of the classes and properties you declare in the application. Do not start a table or column name with a digit.
- Do not use the underscore symbol (`_`) in table or column names, because it is a reserved symbol in Acumatica Framework. For example, `TenantType` is a valid column name, while `Tenant_Type` is invalid.
- Use singular nouns for table names. Typically, a table is mapped to a data access class that represents the entity. For instance, the `SOShipment` table contains data records that represent instances of the `SOShipment` entity.



: Acumatica Framework generates SQL statements with table and column names in the same letter case (that is, uppercase or lowercase) as the corresponding data access classes and fields are declared in the application. Also, the DAC Generator tool produces data access class declarations in the same letter case as the tables and columns are defined in the database schema.

- Use two prefixes in table names: a two-letter tenant name and then a two-letter application module prefix. For example, the `MTSVAppointment` table can be used in the Services (SV) module for the MyTenant tenant (which corresponds to the `MT` prefix). These prefixes help to distinguish your application tables from Acumatica ERP tables and tables of other vendors if you create an add-on project or extension library.
- If you add a column to an Acumatica ERP table, start the column name with the `Usr` prefix followed by the two-letter tenant name. For instance, you could use `UsrMTColumn` for the column of the MyTenant tenant. In this case, the column will be preserved during upgrades. In your own application tables, there are no strict requirements to start column names with any prefixes.
- Be sure that custom indexes on Acumatica ERP tables start with the `Usr` prefix followed by the two-letter tenant name, so that the indexes will be preserved during upgrades.

## Column Naming Conventions

We recommend that you use the following suffixes in column names:

- *ID* for surrogate keys, including database identity columns, such as `CustomerID`
- *CD* for natural keys, such as `CustomerCD`
- *Nbr* for numbering identifiers, such as `OrderNbr`
- *Price* for prices, such as `UnitPrice`
- *Cost* for costs, such as `UnitCost`
- *Amt* for amounts, such as `FreightAmt`
- *Total* for totals, such as `OrderTotal`
- *Qty*, *QtyMin*, and *QtyMax* for quantities, such as `OrderQty`
- *Date* for dates, such as `OrderDate`
- *Time* for time points and time spans, such as `BillableTime`
- *Pct* for percents, such as `DiscountPct`

## Common Columns and Data Types

You should use the following data types for columns. In the **Type Attribute on the Data Field** column in the table below, you can find the most common type attributes that are added to the corresponding data fields in the data access class declaration.

### Common Data Types

Value	Data Type (SQL Server)	Data Type (MySQL)	Type Attribute on the Data Field
Database identity	int	INT	[PXDBIdentity]
Natural key (for example, document number)	nvarchar(15)	VARCHAR(15) with utf8mb4 character set	[PXDBString(15, IsKey = true, IsUnicode = true)]
Line number	int	INT	[PXDBInt]
Short string (for example, a name or unit of measure)	nvarchar(20), nvarchar(50)	VARCHAR(20), VARCHAR(50) with utf8mb4 character set	[PXDBString(20, IsUnicode = true)]
Long string (such as a description)	nvarchar(255)	VARCHAR(255) with utf8mb4 character set	[PXDBString(255, IsUnicode = true)]
Type or status identifier (for instance, a document type)	int or char(1)	INT or CHAR(1)	[PXDBInt] or [PXDBString(1, IsFixed = true)] respectively
Boolean flag (for example, active/inactive)	bit	TINYINT(1)	[PXDBBool]
Price or cost, monetary units	decimal(19, 6)	DECIMAL(19, 6)	[PXDBDecimal(6)]
Amount or total, monetary units	decimal(19, 4)	DECIMAL(19, 4)	[PXDBDecimal(4)]
Quantity, pieces	decimal(25, 6)	DECIMAL(25, 6)	[PXDBDecimal(6)]
Maximum, minimum, or threshold quantity, pieces	decimal(9, 6)	DECIMAL(9, 6)	[PXDBDecimal(2)]
Percent, rate (for example, discount percent)	decimal(9, 6)	DECIMAL(9, 6)	[PXDBDecimal(2)]
Weight or volume	decimal(25, 6)	DECIMAL(25, 6)	[PXDBDecimal(6)]
Date	smalldatetime	DATETIME	[PXDBDate]
Time span	int	INT	[PXDBTimeSpan(DisplayMask = "t", InputMask = "t")]
Coefficient (such as a conversion factor)	decimal(9, 6)	DECIMAL(9, 6)	[PXDBDecimal(1)]

## Primary Key

You have to define the primary key in each application table that you create. The primary key may consist of one column or multiple columns. The primary key must include the `CompanyID` column if one is defined in the table. For details on the `CompanyID` column, see [Multitenancy Support](#).

For each table, you can use one of the following typical variants of primary keys:

- One key column included in the primary key in the table and set as the key in the data access class
- A pair of columns, with one column included in the primary key in the table and the other column set as the key in the data access class
- Multiple columns that are included in the primary key and set as the compound key in the data access class



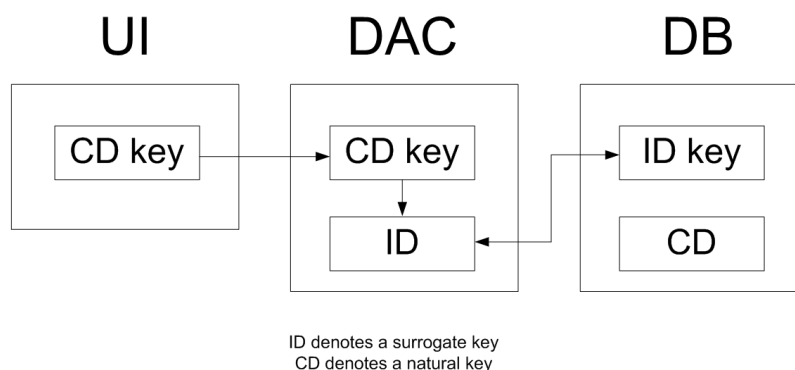
: In a setup table, only the `CompanyID` column must be included in the primary key.

### One Key Column

You may use one key column for rather short tables. For instance, you can use the two-letter country code from ISO 3166 as the key in the `Country` table.

### A Pair of Columns with Key Substitution in the UI

If you want to represent a user-friendly key in the user interface (UI) that corresponds to a surrogate key in the database, you can use a pair of columns and the key substitution mechanism provided by Acumatica Framework. You can define two columns in a table, one for the surrogate key (typically the database identity column) and one for the natural key, and set only the surrogate key as primary in the table. In the application object model, you set the key to only the data field that is a natural key. In this case, Acumatica Framework provides the ability to transparently work with different keys at the database and application levels. In the UI, users work with only the natural key while the database operates with the surrogate key (see the graphic below, which illustrates key substitution).



**Figure: Key substitution in Acumatica Framework**

For instance, you can define two columns in the `Product` table, `ProductID` and `ProductCD`. `ProductID` is the identity column that is the only column included in the primary key of the table. `ProductCD` is the string key of a product instance, which is entered by the user through the UI. The `ProductCD` column isn't included in the primary key and is handled as the unique key column by Acumatica Framework.

## Multiple Column Key

A compound key consisting of multiple columns may be used for complex entities. For instance, you can include two columns, `DocType` and `DocNbr`, in the primary key for the `Document` table. In the `DocDetail` table, you may use `DocNbr` and `DocDetailNbr` as the compound primary key. The corresponding data fields should be also set as the key fields in the data access class.

## Foreign Keys and Nullable Columns

In the database, you have to define the primary key in each application table that you create. The primary key defines the unique data record identifier, which provides table-level integrity of data.

There are no strict requirements to define column-level constraints and foreign keys in application tables. Whether you define the constraints at the database level depends on the design approach you use. At a higher level of the application object model, which is represented by data access classes, you can flexibly define any level of constraints, including default values, nullable fields, and parent-child relationships between data access classes. If you aren't sure whether a column should allow a null value, you can allow null values for it in the database. Later, in the data access class, you can make the data field either required or nullable; you can even make the field required on one page and optional on another.



: For Boolean and decimal columns, we recommend that you define default values either in the database or in data access classes. This simplifies the application code by helping to avoid checking of values for nulls multiple times.

## Audit Fields

Audit fields keep meta information on the creation and the last change of a database record. Audit fields are updated automatically by the framework.

To enable the tracking of audit data for a particular table, you should add the columns listed below to the table and declare the corresponding audit data fields in the data access class. You have to add the corresponding type attribute to each audit field. If the audit columns are properly created in the database table and the corresponding data fields are declared in the data access class, Acumatica Framework automatically updates audit data in these fields every time a data record is modified from the application. The audit column parameters and DAC attributes are given below.

### Audit Columns

Database Column Name	Data Type (SQL Server)	Data Type (MySQL)	Type Attribute on the Data Field
CreatedByID	uniqueidentifier; not null	CHAR(36) with ASCII character set; not null	[PXDBCreatedByID]
CreatedByScreenID	char(8); not null	CHAR(8) with ASCII character set; not null	[PXDBCreatedByScreenID]
CreatedDateTime	smalldatetime; not null	DATETIME; not null	[PXDBCreatedDateTime]
LastModifiedByID	uniqueidentifier; not null	CHAR(36) with ASCII character set; not null	[PXDBLastModifiedByID]
LastModifiedByScreenID	char(8); not null	CHAR(8) with ASCII character set; not null	[PXDBLastModifiedByScreenID]

Database Column Name	Data Type (SQL Server)	Data Type (MySQL)	Type Attribute on the Data Field
LastModifiedDateTimes	smalldatetime; not null	DATETIME; not null	[PXDBLastModifiedDateTime]

## Concurrent Update Control

You can add the SQL Server time stamp column to a table to make Acumatica Framework able to handle concurrent updates. The corresponding time stamp data field should be declared in the data access class. If the time stamp data field is declared, Acumatica Framework handles the time stamp column automatically. Acumatica Framework checks the row version every time the row is modified. We recommend that you add the time stamp column, with the parameters shown in the following table, to all tables of your application.

### *The Time Stamp Column*

Database Column Name	Data Type (SQL Server)	Data Type (MySQL)	Type Attribute on the Data Field
TStamp	timestamp; not null	TIMESTAMP (6); not null	[PXDBTimestamp]

## Attachment of Additional Objects to Data Records

You can attach additional objects to a data record—for instance, attach a text note or an uploaded file or multiple uploaded files to a data record.

You turn on or off support for data record attachments for each particular table individually. To turn on support for data record attachments, add a column that stores the global data record identifier (typically, `NoteID`) to the table and declare the corresponding field in the data access class. For more information on uploading files through an application page, see [To Display an Attached Image on the Webpage](#). See below for the parameters of the global identifier column and the attribute that should be added to the corresponding DAC field.

### *The Global Data Record Identifier Column (NoteID)*

Database Column Name	Data Type (SQL Server)	Data Type (MySQL)	Type Attribute on the Data Field
Global data record identifier (typically named <code>NoteID</code> )	uniqueidentifier; not null	CHAR (36)	[PXNote]

## Preservation of Deleted Records

Acumatica Framework provides a low-level mechanism (which is performed on the database level) for preserving deleted data records in the database. With this mechanism, when an application initiates the deletion of a data record, the data access layer generates an SQL query that marks the data record as deleted but does not permanently remove the data record from the table. When data records are selected from the table, the data access layer generates the SQL query, which returns only data records that are not marked as deleted. The data records that are preserved in this way can be restored.

You can turn on or off the preservation of deleted data records for each table individually. To preserve data records in a particular table, add the `DeletedDatabaseRecord` column to the table and do not declare the data field in the data access class. When a data record is deleted in the table, the framework automatically preserves the deleted data record transparently to the application developer.

### The DeletedDatabaseRecord Column


Database Column	Data Type (SQL Server)	Data Type (MySQL)	Type Attribute on the Data Field
DeletedDatabaseRecord	bit; not null	TINYINT(1); not null	Not declared in DAC

## Multitenancy Support

Multiple tenants can work on the same instance of an Acumatica Framework-based application with completely isolated data. The application looks identical to all tenants, but each tenant has exclusive access to its data only. Data is isolated at the lowest level of the application, in the data access layer that executes SQL queries for the tenant of the user who is currently signed in.

### Multitenancy Support

The following graphic illustrates how different logical tenants work with the Acumatica Framework-based application in a multitenant configuration. They work with the same application but have isolated data access, as if they are working with different database instances.



**Figure: Multitenant Acumatica Framework-based application**

Multitenancy support is turned on or off for each particular table individually. To turn on multitenancy support for a table, add the `CompanyID` column to it and include the column in the primary key (see the column parameters in the table below) and all indexes. The `CompanyID` column is handled automatically by the framework and should not be declared in data access classes. If a table doesn't have the `CompanyID` column, all data from the table is fully accessible to all tenants that exist in the database. For more information, see [Support of Multiple Tenants](#).

### The CompanyID Column


Database Column Name	Data Type (SQL Server)	Data Type (MySQL)	Type Attribute on the Data Field
CompanyID	int; not null; included in primary key and all indexes	INT; not null; included in primary key and all indexes	Not declared in DAC

### Support for Shared Data Access Between Tenants

Acumatica Framework provides shared data access in a multitenant configuration. Acumatica Framework supports a hierarchy of logical tenants that may work with a combination of shared and individual data. In shared access mode, every tenant may work with its individual copy of a data record; copies differ by `CompanyID`. All copies represent the same logical object in the application but different data records in the database. For instance, each tenant may use the individual settings of the application.

The graphic below shows a possible multitenant configuration with shared data access between Tenant 1, Tenant 2, and Tenant 3. The users of Tenant 2 have access to the data of all three tenants. The users

from each of the other two tenants have access to their company's individual data only. Physically, the data of all three tenants is stored in a single database instance.



**Figure: Shared data access in a multitenant Acumatica Framework-based application**

Support for shared data access is turned on or off for each particular table individually. To turn on support for shared data access for a table, add the `CompanyMask` column to the table (see the column parameters in the table below). The `CompanyMask` column is handled automatically by the framework and should not be declared in data access classes. If a table doesn't have the `CompanyMask` column, shared data access is not available for this table.

#### **The `CompanyMask` Column**

Database Column Name	Data Type (SQL Server)	Data Type (MySQL)	Type Attribute on the Data Field
<code>CompanyMask</code>	<code>varbinary(32)</code> , not null, default 0xAA	<code>VARBINARY(32)</code> , not null, default 0xAA	Not declared in DAC

`CompanyMask` is a 32-bit mask. In this mask, each two bits correspond to each tenant. The first of these two bits specifies whether the record may be read by this tenant, and the second bit specifies whether the record may be written to by this tenant. For example, suppose that `CompanyMask` is set to 0xBE02 for a record. That is, it specifies the following mask: 10 11 11 10 00 00 00 10, which designates that the record may be both read and written to by the tenants with company IDs 2 and 3, the record may be read by the tenants with IDs 4 and 5 and the system tenant (which has ID 1), and the record may not be read or written to by other tenants.

```
CompanyMask: 10 11 11 10 00 00 00 10
CompanyID:   4  3  2  1  8  7  6  5
```

The default value of `CompanyMask` is 0xAA, which means that the record may be read by all tenants.

## Designing the User Interface

This chapter summarizes the webpage design and style conventions used in Acumatica Framework.

### **In This Chapter**

- [Development Environment Options](#)
- [Form and Report Numbering](#)
- [Item Grouping on the Form Toolbar](#)

## Development Environment Options

To create stand-alone applications with Acumatica Framework or develop customizations for Acumatica ERP, the environment where you install and use Acumatica Framework, should meet particular requirements that are described in [System Requirements for Acumatica Framework 2019 R1](#).

## Web Forms Designer Settings

We recommend the following settings for the Microsoft Visual Studio environment to ensure a uniform webpage appearance:

1. Under the **Tools > Options > Web Forms Designer > CSS** section, set the following options:
  - **Font and text:** *CSS (classes)*
  - **Padding and borders:** *CSS (classes)*
  - **Floating, positioning, and sizing:** *CSS (inline styles)*
  - **Bullets and numbering:** *CSS (classes)*
  - **Background:** *CSS (classes)*
  - **Margins:** *CSS (classes)*
2. Under the **Tools > Options > Web Forms Designer > CSS Styling** section, select **Auto Style Application**, and specify the following settings:
  - **Only reuse classes with the prefix "auto-style":** Selected
  - **Use width and height attributes for image instead of CSS:** Selected
  - **Use <strong> and <em> for bold and italic text:** Cleared
  - **Use shorthand properties when generating styles:** Selected
  - **Change positioning to absolute for controls added using Toolbox, paste, or drag and drop:** Selected

## Design Mode Settings

We also recommended that you use the following settings of the Design mode of ASPX pages in Visual Studio:

- **View > Visual Aids > CSS Display:none Elements:** Cleared
- **View > Visual Aids > CSS Visibility:hidden Elements:** Cleared

## Form and Report Numbering

In this topic, you can find the guidelines for form and report numbering in Acumatica ERP.

### Form Numbering

When you are numbering forms in Acumatica ERP, use the following conventions:

```

XX999999
| | | | _ Subscreen Sequential Number
| | | _ Screen Sequential Number
| | _ Screen Type:
| |         10: Setup
| |         20: Maintenance
| |         30: Data Entry
| |         40: Inquiry
| |         50: Processing
| |         60: Reports
| |         Two-Letter Module Code

```

### Report Numbering

When you are numbering reports in Acumatica ERP, use the following conventions in addition to those outlined above:

```

XX6X9999

```



```
|
| _____ Report Type:
|         61: Review Reports (Reports for document review prior to release)
|         62: Register Reports (Reports used to print audit information
|                               on processed documents or entities)
|         63: Balance Reports (Reports reflecting current or historical
|                               balance information)
|         64: Forms (Printed webpages)
|         65: Inquiry Reports (Reports that provide status information
|                               required for operational management)
|         66: Statistical Reports (Reports that provide statistical or
|                               historical information)
```

## Item Grouping on the Form Toolbar

Menu items can be grouped on the form toolbar to keep a reasonable number of buttons on the toolbar. When you are building the menu structure, use the menus described below.

### Data Entry Forms

- **Actions:** Use this menu to group the operations that give the user the ability to process the document, including the actions that navigate to related data entry forms (with the system filling in appropriate settings) so users can quickly create linked documents. For example, see the **Enter Payment/Apply Memo** action on the [Invoices and Memos](#) (AR301000) form. The most frequently used operations can be placed on the toolbar outside any groups as separate buttons that provide quick access to the actions. For example, notice the **Release** action on the [Invoices and Memos](#) form.
- **Reports:** Use this menu to group the actions that open related Report Designer reports and printable forms of documents.
- **Inquiries:** Use this menu to group the actions that navigate to related inquiry forms.

### Inquiry Forms

- **Actions:** Use this menu to group the operations that give the user the ability to navigate to related data entry forms.
- **Reports:** Use this menu to group the actions that open the related Report Designer reports.

### Maintenance Forms

- **Actions:** Use this menu to group the operations that update the settings of the master record and navigate to related data entry forms.
- **Reports:** Use this menu to group the actions that open related Report Designer reports.
- **Inquiries:** Use this menu to group the actions that navigate to related inquiry forms.

## Naming the Graphs and Event Handlers

---

In this chapter, you can find the naming conventions for the graphs and event handlers.

### In This Chapter

- [Graph Naming](#)
- [Naming Conventions for Event Handlers Defined in Graphs](#)

## Graph Naming

When you are creating business logic controllers (graphs), use the following suffixes in the names of the graphs, depending on the types of the webpages they are used for:

- **Maint:** For the graphs for maintenance pages that are helper pages used for the input of data on the data entry and processing pages, and for the graphs for the setup pages that provide the configuration parameters for the application. For example, `CountryMaint` can be the name of the graph for the Countries maintenance page, which provides editing of the list of countries.
- **Entry:** For the graphs for data entry pages that are used for the input of business documents. For example, `SalesOrderEntry` can be the name of the graph for the Sales Order data entry page, which provides the basic functionality for working with sales orders.
- **Inq:** For the graphs for inquiry pages, which display a list of data records selected by the specified filter. For example, `SalesOrderInq` can be the name of the inquiry page named Sales Order Inquiry, which provides the list of documents selected by the specified customer.
- **Process:** For the graphs for processing pages that provide mass processing operations. For example, `SalesOrderProcess` can be the name of the Approve Sales Orders processing page, which provides mass approval of sales orders.

## Naming Conventions for Event Handlers Defined in Graphs

In Acumatica Framework, you must adhere to the naming conventions for an event handler to be implemented in a graph or graph extension. The name of the event handler must include the event type and the object to be processed by the handler.

The name of a data record event handler must have the following segments, which are separated by the `_` symbol:

1. The name of the DAC declared in the server
2. The name of the record event supported by the server

Therefore, the name of a data record event handler must be in the following format:

*DACName\_EventName* (such as `SOOrder_RowSelected`).

The name of a data field event handler must have the following segments, which are separated by the `_` symbol:

1. The name of the DAC declared in the server
2. The name of the data field declared within the DAC whose name is specified in the first segment
3. The name of the field event supported by the server

Therefore, for a field event handler, the name must be in the following format:

*DACName\_FieldName\_EventName* (such as `SOOrder_CustomerID_FieldUpdated`).

# Configuring ASPX Webpages and Reports

---

In this part of the guide, you can find information about how the webpages of Acumatica ERP or an Acumatica Framework-based application work and how to configure the ASPX code of these webpages. You can also find a brief overview of reports.

## In This Part

- [Overview of ASPX Webpages in Acumatica Framework](#)
- [Configuring Containers](#)
- [Configuring Tables](#)
- [Configuring Tabs](#)
- [Configuring Boxes](#)
- [Configuring Layout and Size](#)
- [Maintaining Reports](#)

## Overview of ASPX Webpages in Acumatica Framework

---

In Acumatica Framework-based applications, you can configure the appearance of webpages from the front end (that is, by configuring ASPX pages) and from the back end (that is, by using the attributes and events provided by the Acumatica platform).

In this chapter of the guide, you can find information about how to create the ASPX code of the webpages. This chapter also includes a technical overview of the user interface of an Acumatica Framework-based application. For information about the configuration of the UI from the back end, see [Configuring the UI from the Back End](#).

### Technical Overview of the User Interface

The user interface of Acumatica ERP or an Acumatica Framework-based application uses the same core technologies. To learn how the user interface works, see [Processing of a Button Click](#) and [Technical Overview of the Modern UI](#). You can find a detailed description of the user interface in [Acumatica ERP Interface Guide](#).

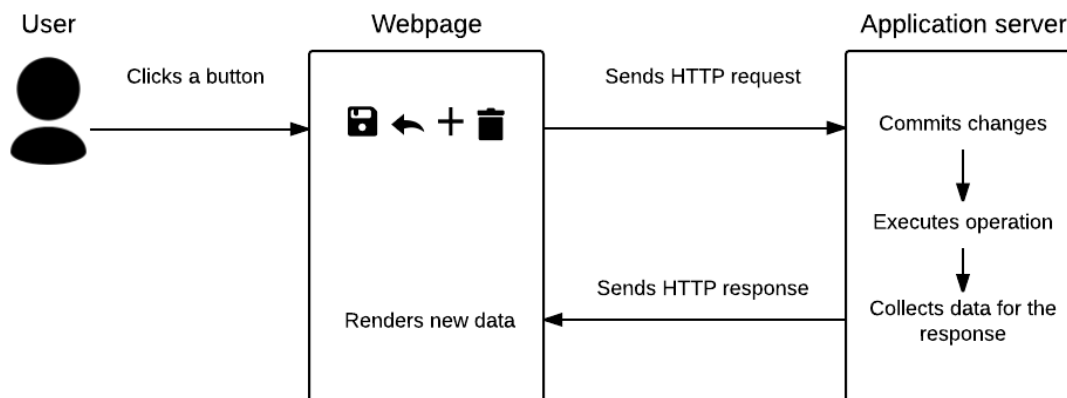
### Configuration of ASPX Webpages

You can configure the appearance and behavior of webpages by using the properties of the ASPX objects that can be used on the Acumatica ERP or Acumatica Framework webpages. For details on how to work with the ASPX code of the webpages, see the following topics:

- To configure containers, such as `PXFormView`, `PXGrid`, `PXTab`, `PXTreeView`, and `PXPanel`: [Configuring Containers](#)
- To configure tables (`PXGrid` and `PXGridColumn`): [Configuring Tables](#)
- To configure tabs (`PXTab` and `PXTabItem`): [Configuring Tabs](#)
- To configure boxes, such as `PXTextEdit`, `PXCheckBox`, and `PXGroupBox`): [Configuring Boxes](#)
- To configure the layout (`PXLayoutRule`) and size of controls: [Configuring Layout and Size](#)

## Processing of a Button Click

When a user clicks a button on an ASPX page, such as the **Save** button on the toolbar of a form, the webpage creates a postback HTTP request to the server. When processing the request, the application server does the following: commits changes (if required by the operation initiated by the user), executes the operation, and collects data for the response. Then the application server sends the response to the webpage, which renders the new data. The following diagram illustrates this process, which is described in more detail in the sections of this topic.



**Figure: Processing a button click**

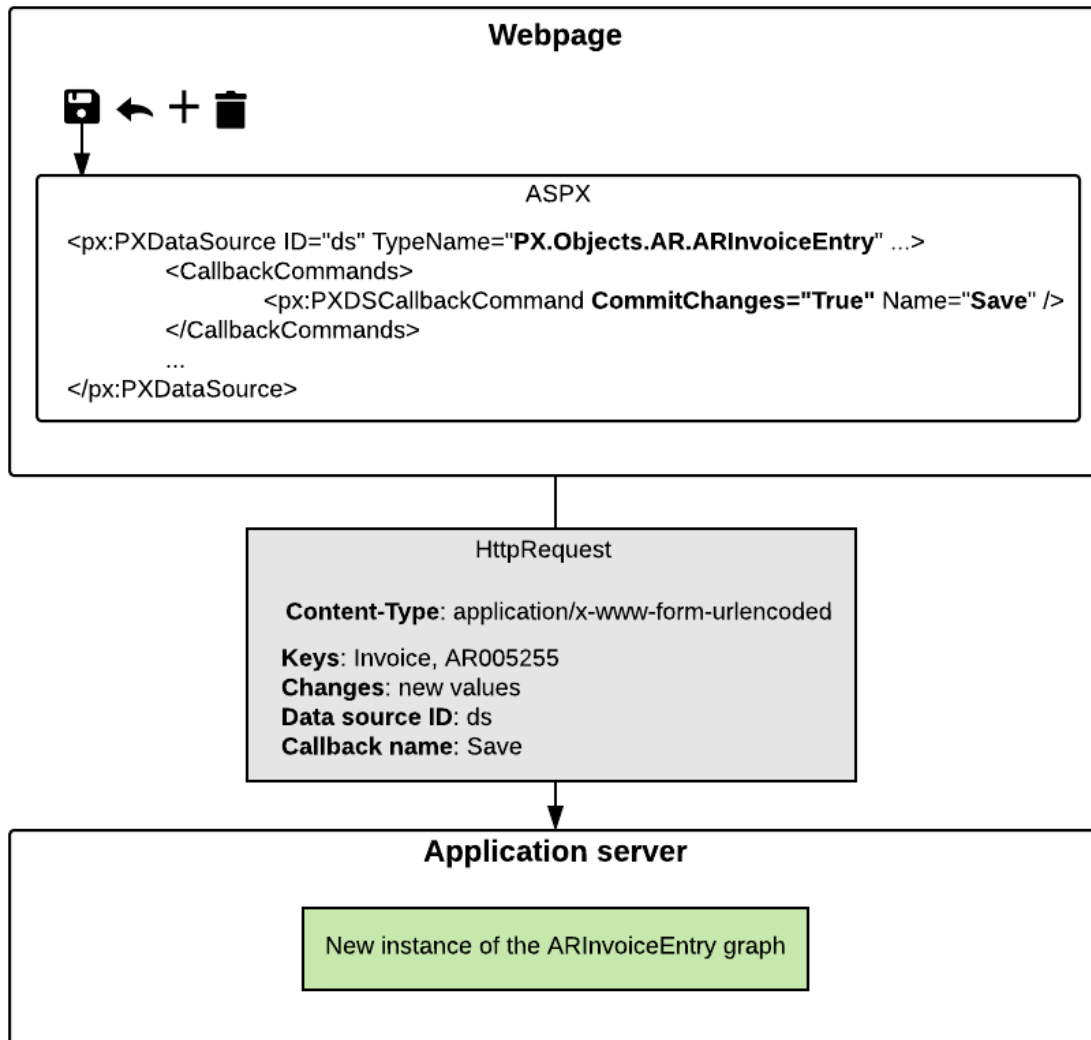
### Sending of the HTTP Request

When the user clicks a button on a webpage, the page creates an HTTP request to the server. The request includes the following information:

- The values of the key fields of the record currently displayed on the page
- The changes that have been made to the data on the page
- The information on the command that was initiated by the user—the data source ID and the callback name

The `PXDataSource.TypeName` property defines the graph that processes data for the page. When the application server receives the request, the server creates a new instance of the graph to process the data of the request. The properties of `PXDSCallbackCommand`, such as `CommitChanges` and `RepaintControls`, indicate which operations should be performed on the server in addition to the operation initiated by the user.

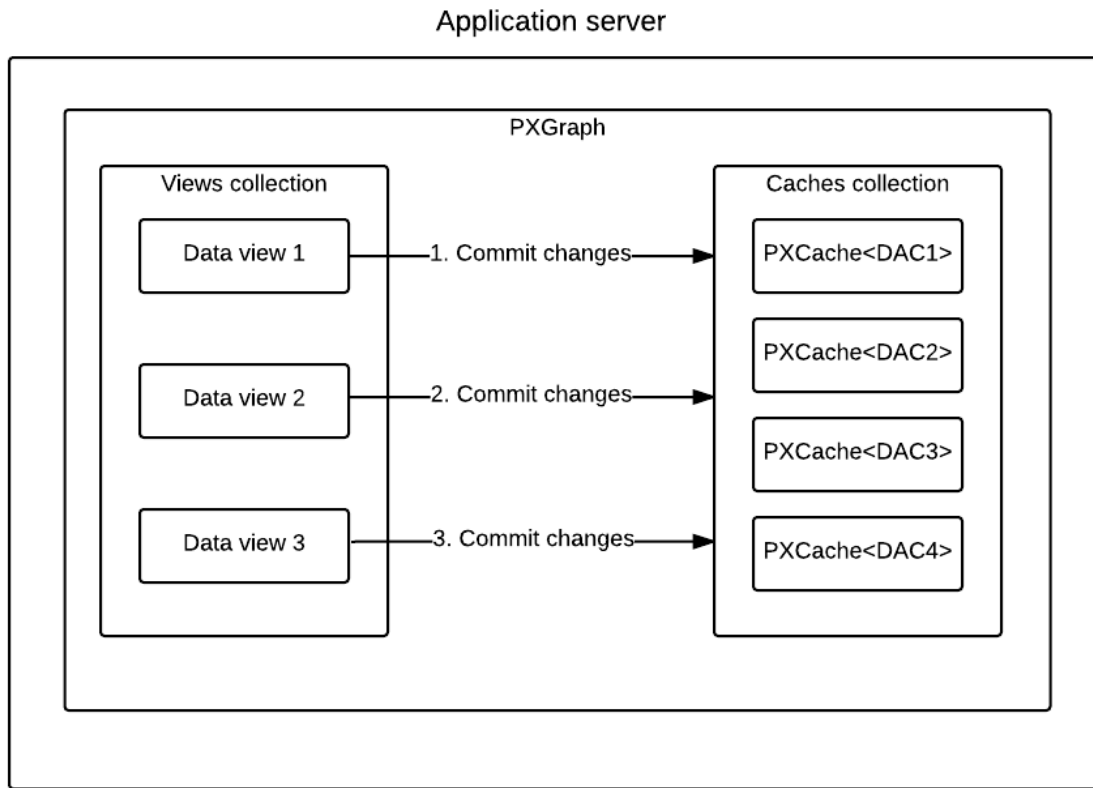
The diagram below shows how the HTTP request is sent to the server.



**Figure: Sending the HTTP request**

### Commitment of Changes to the Cache

If the callback command initiated by the user has the `CommitChanges` property set to `true` or the `CommitChangesIDs` property specified, the server commits the changes before executing the command. The graph instance commits the changes to the cache in the order in which the data views are defined in the graph, as shown in the following diagram.



**Figure: Committing changes**

### Execution of the Command

After the changes have been committed, the graph instance executes the operation initiated by the user, such as saving data to the database. You can find details on the sequence of events raised when data is inserted, updated, deleted, or saved to the database in [Data Manipulation Scenarios](#).

### Collection of the Data for the Response

When the command execution is completed, the application server does the following:

1. If the `RepaintControls` or `RepaintControlIDs` property of `PXDSCallbackCommand` specifies any controls to be repainted after the command is executed, the application server includes in the response all information that is necessary to repaint these controls on the webpage. (By default, the value of the `PXDSCallbackCommand.RepaintControls` property is `All`, which means that all controls on the page are repainted.)
2. The application server executes the `Select` method for each data view of the graph.

### Sending of the HTTP Response and Rendering of the Controls

The application server sends the response to the page. The response includes data in XML format; the parameters that are necessary for the controls to be repainted are specified in JSON format, as shown in the following fragment of the response.

```

<Controls>
  <Control ID="ctl100_phF_form_edDocType"
    Props="{items:&quot;INV|Invoice|1;DRM|Debit Memo|1;CRM|
      Credit Memo|1;FCH|Overdue Charge|1;SMC|Credit WO|1&quot;;,

```

```

        value:&quot;FCH&quot;}}" />
    <Control ID="ctl100_phF_form_edRefNbr" Props="{value:&quot;AR005254&quot;}" />
    ...
</Controls>

```

The scripts in the browser (the scripts from `PX.Web.UI.Scripts`) find the controls to be repainted by IDs and repaint these controls by using the data provided in the response. Most of the scripts in `PX.Web.UI.Scripts` contain a class that works with one control. For example, `px_textEdit.js` includes the `PXTextEdit` class, which works with the `PXTextEdit` control.

### Exceptions to the Process

For the buttons not found on the main toolbar, the process described in this topic may slightly differ.

For example, for the table toolbar buttons, which perform actions on particular rows of grids, the scripts translate the data in XML format, which is returned in the response, to HTML format by using XSLT.

For the buttons in dialog boxes (the `PXSmartPanel` control), no commit of changes is performed. Whether data is selected from the database depends on the particular dialog box. The data that is returned from the server is in HTML format.

## Technical Overview of the Modern UI

Acumatica ERP provides two versions of its user interface: the modern UI, which is introduced in Acumatica ERP 2017R1, and the classic UI, which was available in previous versions. In this topic, you can find a technical overview of the modern UI.

For more details on the elements of the modern UI, see [Modern Acumatica ERP User Interface](#) in the Interface Guide.

### Technologies in the Modern UI

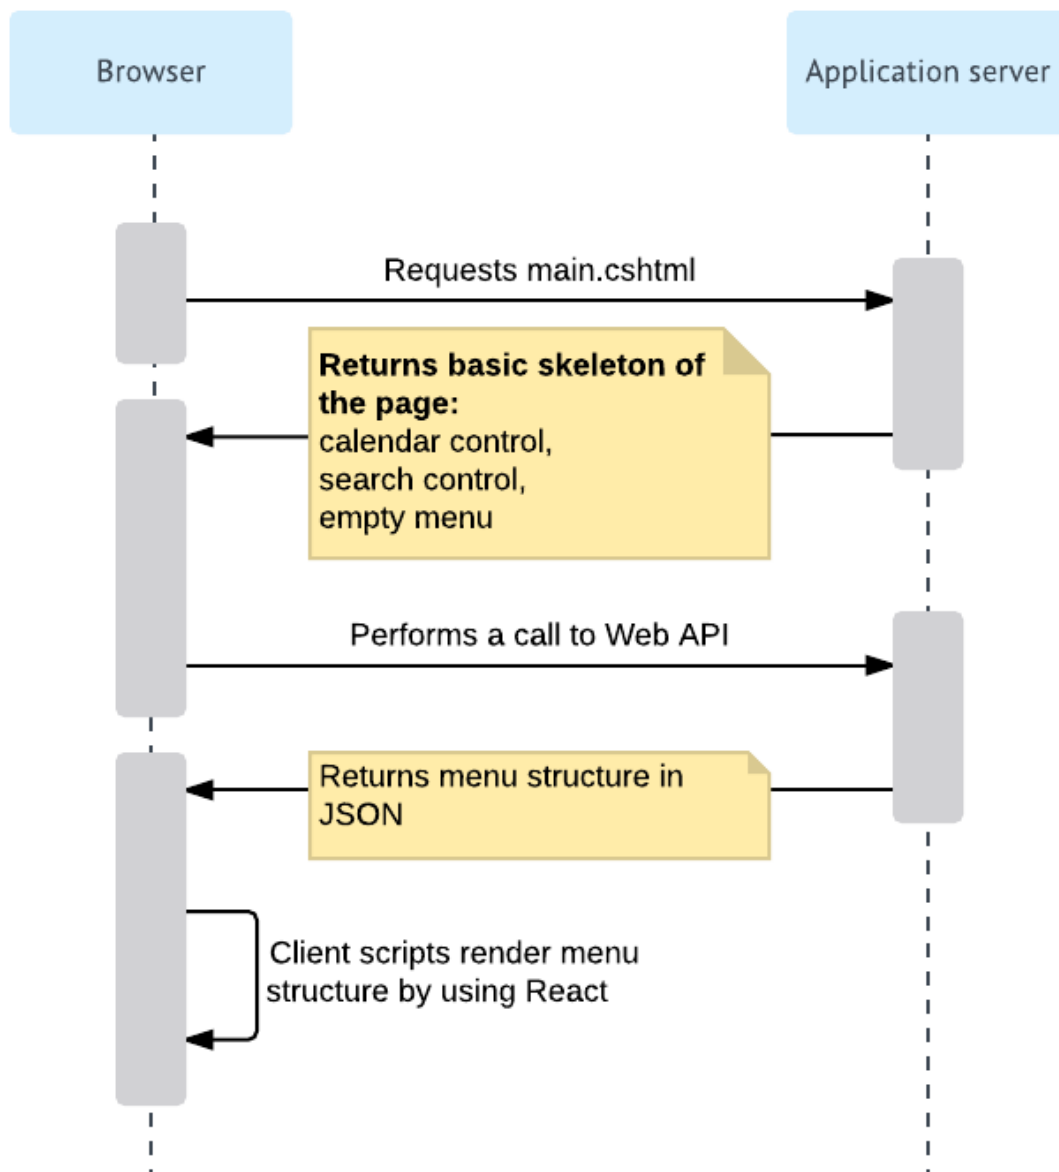
The modern UI is a new webpage frame that can be used for navigation between Acumatica ERP forms instead of the classic UI navigation, which provides the site map that users use to navigate between pages.

With the modern UI, the Acumatica ERP form is located in the inner frame, which is completely independent of the modern UI frame. The webpage renders the navigation elements of the modern UI and the forms separately by using different technologies. The work of Acumatica ERP forms is based on the ASP.NET Web Forms technology, while the modern UI uses the ASP.NET MVC technology with the Razor view engine. Therefore, regardless of whether a user selects the modern UI or the classic UI for navigation between Acumatica ERP forms, the forms themselves work in the same way and use the same business logic.

The server side of the modern UI uses the ASP.NET Web API framework. The client side uses the React library, which is a JavaScript library, to render main menu items, workspaces, tiles, and other navigation elements of the modern UI.

### Work of the Modern UI

The following diagram shows how the browser renders the modern UI elements. This process is described in more detail in the remaining sections of this topic.



**Figure: Rendering of the Modern UI**

### Request of main.cshtml

The browser retrieves `main.cshtml`, which is an ASP.NET MVC view, by sending the HTTP GET request. On the server side, this request is processed by the `MainController.Main()` method (`PX.Web.UI.Frameset.Controllers`), which creates a `System.Web.Mvc.ViewResult` object that renders a view to the response. The returned view contains the basic skeleton of the webpage, which includes the calendar control, the search control, and the empty menu.

### Request of the Menu Structure

The `getSiteMap` function in `site.js` uses jQuery to send an AJAX request to ASP.NET Web API of the application server. On the server side, this request is processed by the `SiteMapController` class (`PX.Web.UI.Frameset.WebApi.Controllers`).





: To match the incoming request to the appropriate processing classes, the system uses the ASP.NET MVC attribute routing. For example, the `SiteMapController` class is annotated with the `[FramesetRoutePrefix("sitemap")]` attribute, which defines the "frameset/sitemap" route.

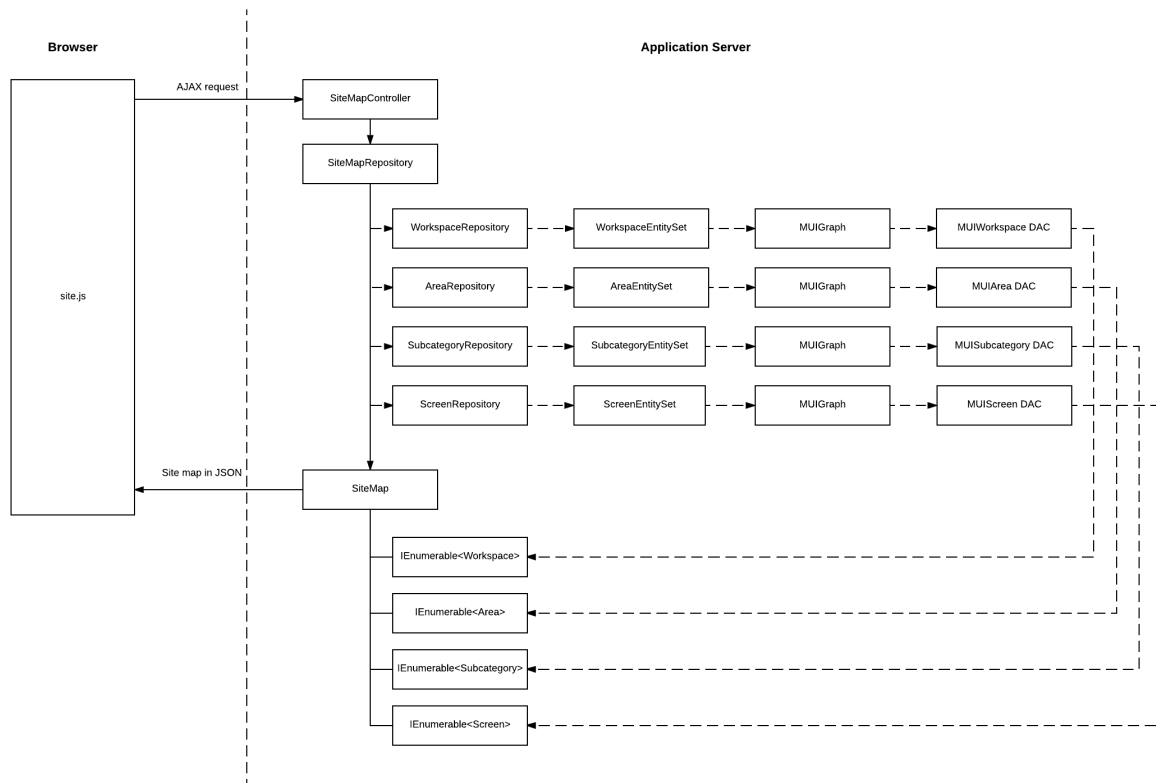
To get the site map structure, the `SiteMapController` class uses the `SiteMapRepository` class, which implements the `ISiteMapRepository` interface. The `SiteMapRepository` class fetches different entities of the modern UI and assembles them in one structure, which is then passed to the browser. The system serializes the structure to JSON format by using the standard ASP.NET Web API classes.

The `SiteMapRepository` class uses other classes that have the `Repository` suffix in their names, such as `TileRepository` and `WorkspaceRepository`, to retrieve the entities that are used in the modern UI. These classes are completely independent from the database. To fetch the entities from the database, the `Repository` classes use the classes that implement the `IEntitySet` interface (`PX.Web.UI.Frameset.Model`), such as `ScreenEntitySet`. The classes use the `MUIGraph` graph to fetch data from the database. (The graph performs only simple data operations, and does not contain any complicated business logic). For each entity, there is a DAC that is used to access data in the database. The DACs correspond to the following database tables, which are used to store data for the elements of the modern UI.

#### Database Tables

Table	Description
MUIWorkspace	Stores information about the workspaces in the application. For more information on the workspaces, see <a href="#">Workspaces in the Modern UI</a> in the Interface Guide.
MUIFavoriteWorkspaces	Stores information about the workspaces that have been pinned to the main menu. The workspaces that are not included in this list are displayed when a user clicks the <b>More Items</b> menu item. For details about the main menu, see <a href="#">Main Menu in the Modern UI</a> in the Interface Guide.
MUIArea	Stores information about the areas to which workspaces belong. Areas are used to group workspaces in the <b>More Items</b> menu by types.
MUISubcategory	Stores information about the categories of Acumatica ERP forms. Categories are used to group forms in a workspace by types. For details on the categories, see <a href="#">Categories</a> .
MUIScreen	Stores information about the locations of the Acumatica ERP forms in the modern UI. The table is connected to the <code>SiteMap</code> table by the <code>NoteID</code> column.
MUIPinnedScreen	Stores information about the Acumatica ERP forms pinned to workspaces.
MUIFavoriteScreen	Stores information about the Acumatica ERP forms that have been added to favorites.
MUITile	Stores information about the tiles in workspaces. A tile is a special button on a workspace that you click to open a form or report with predefined settings. For details on the tiles, see <a href="#">Tiles</a> .
MUIFavoriteTile	Stores information about the tiles that have been added to favorites.
MUIUserPreferences	Stores information about the position of the main menu, which can be on the left of the browser page (default) or on the top of the browser page.

The following diagram illustrates the process of retrieving data for the modern UI.



**Figure: Retrieval of the site map**

## Rendering of the Modern UI Elements

The main script that is used to render the modern UI is `site.js`. It contains classes that use the React library to render modern UI elements. Each such class includes the `render` method, which uses React classes to render the element. The following tables lists the main classes and their methods.

### The Classes for Rendering the Modern UI

Class	Description
MenuModules	<p>Renders the main menu (which contains the list of workspaces). For more information on the main menu, see <a href="#">Main Menu in the Modern UI</a> in the Interface Guide.</p> <p>In addition to the <code>render</code> method, the class has the following methods:</p> <ul style="list-style-type: none"> <li><code>onClick</code>: Processes the clicking of the Edit and Delete buttons for the items of the main menu in Menu Editing mode.</li> <li><code>onClickFav</code>: Processes the clicking of the Pin button in a workspace.</li> <li><code>onDragStart</code>, <code>onDragOver</code>, <code>onDragLeave</code>, and <code>onDrop</code>: Process operations related to dragging the items of the main menu in Menu Editing mode.</li> </ul> <p>For details on menu editing, see <a href="#">Menu Editing Mode in the Modern UI</a>.</p>
TopLinks	<p>Renders the tiles in the workspaces. For details about the tiles, see <a href="#">Tiles</a> in the Interface Guide.</p> <p>In addition to the <code>render</code> method, the class has the following methods:</p>

Class	Description
	<ul style="list-style-type: none"> <li><code>onClick</code>: Processes the clicking of the Edit and Delete buttons for the tiles in Menu Editing mode and clicking of the Favorite button.</li> <li><code>onDragStart</code>, <code>onDragOver</code>, <code>onDragLeave</code>, and <code>onDrop</code>: Process operations related to dragging the tiles in Menu Editing mode.</li> </ul> <p>For details on menu editing, see <a href="#">Menu Editing Mode in the Modern UI</a>.</p>
MenuColumn	<p>Renders a list of forms in a workspace. For the information about workspaces, see <a href="#">Workspaces in the Modern UI</a> in the Interface Guide.</p> <p>In addition to the <code>render</code> method, the class has the following methods:</p> <ul style="list-style-type: none"> <li><code>onClick</code>: Processes the clicking of the Edit and Delete buttons for a form in a workspace in Menu Editing mode.</li> <li><code>onClickFav</code>: Processes the clicking of the Favorite icon for a form in a workspace.</li> <li><code>onClickPin</code>: Processes the clicking of a check box when a user selects a form in a workspace in Menu Editing mode.</li> </ul> <p>For more information on menu editing, see <a href="#">Menu Editing Mode in the Modern UI</a>.</p>
ModuleMenu	<p>Renders all lists of forms in a workspace. For details about workspaces, see <a href="#">Workspaces in the Modern UI</a> in the Interface Guide.</p> <p>In addition to the <code>render</code> method, the class has the following methods:</p> <ul style="list-style-type: none"> <li><code>getItemsInCol</code>, <code>arrangeLinks</code>, and <code>arrangeLinks2</code>: Arrange links to forms in lists.</li> <li><code>onDragStart</code>, <code>onDragOver</code>, <code>onDragLeave</code>, and <code>onDrop</code>: Process operations related to dragging the links to forms in a workspace in Menu Editing mode.</li> </ul>

The `site.js` script also contains webpage event handlers, such as handlers for button-clicking events, which use jQuery to handle the events.

### Customization of the Modern UI

An administrator can configure the modern UI to fit the work purposes of the organization, as described in [Customizing the Modern User Interface](#) in the System Administration Guide.

To change the styles of the elements of the modern UI, the developer can change the CSS related to these elements.

If a developer has added a new form or report to the Acumatica ERP site in a customization project, the location of the form in the modern UI is included in the customization project along with the *SiteMap* customization project item, which is created either automatically or manually for the new item. For details, see [To Add a New Custom Form to a Project](#) and [To Add a Custom Analytical Report to a Project](#) in the Customization Guide. For the custom generic inquiries and dashboards, the information about the location in the modern UI is included in the *GenericInquiryScreen* and *Dashboard* customization project items, respectively.

## Configuring Containers

In this chapter, you can find information about the configuration of different types of containers, such as `PXFormView`, `PXGrid`, `PXTab`, `PXTreeView`, and `PXPanel`.

## In This Chapter

- [Use of the DataMember Property of Containers](#)
- [Use of the SkinID Property of Containers](#)
- [Use of the Caption Property of Containers](#)
- [Use of Multiple Data Views for Boxes in Containers](#)

## Use of the DataMember Property of Containers

If you need to find out which data view provides data for a control container on a form, perform a search to find the *DataMember* string in the appropriate ASPX code. The `DataMember` property is used to bind a control container of a form to a data view defined in the business logic controller (BLC, also referred as *graph*) of the form. The property value is the name of the data view.

### Container Types That Have the DataMember Property

In Acumatica Framework, `DataMember` is used to specify a data view for the following container types:

- `PXFormView`
- `PXGrid`
- `PXTab`
- `PXTreeView`



: The `PXTreeView` container is not supported by the tools of the Acumatica Customization Platform.

By default, a nested container inherits the `DataMember` property from the parent container. If a nested container is `PXFormView`, `PXGrid`, or `PXTab`, it can be bound to another data view.

If the `DataMember` property is available for other ASPX objects, it has a special purpose. For example, you can specify the `DataMember` property for a `PXSelector` lookup control to define the appropriate data view for the grid of the lookup window.

### Property Value

Each `DataMember` property value can correspond to any data view name of the BLC. Any data view except for the main data view can be used by an unlimited number of containers. The main data view must be bound to a single container.

For a container to contain a box for a data field, the container must be bound to a data view declared within a BLC for the following reasons:

- A data field is declared in a data access class (DAC). An instance of the DAC record can exist in the cache of a BLC that contains the declaration of a data view with the DAC reference in the BQL statement.
- Each time a data record is selected in the container, the container creates a callback to the `PXDataSource` control that is specified in the `DataSourceID` property of the container. The data source control creates a remote procedure call to the application server to execute the *Display* operation on the data view that is specified as the `DataMember` for the container. The data view checks the existence of the record in the cache; if the check fails, the data view executes the BQL request and stores the obtained record in the cache.
- The data view provides all data exchange operations with the database, cache, and `PXDataSource` control.

## Use of the SkinID Property of Containers

In the code of Acumatica ERP, predefined skins are used to assign a style and a set of toolbar buttons to a container. The `SkinID` property of a container specifies which of these skins the system should apply to the container. A skin is specific to a particular container; you cannot share a skin setting between containers of different types. If you do not set the `SkinID` property, the container uses the default skin, if one is defined.

### Predefined Skins

The following table lists and describes the predefined skins that are recommended to use for the `PXFormView`, `PXGrid`, and `PXPanel` containers.

UI Element	SkinID	Description	Example
<code>PXFormView</code>	<i>Transparent</i>	Is used to display a simple form container that has no caption and cannot be collapsed.	The form container on the <b>Financial Details</b> tab item of the <a href="#">Bills and Adjustments</a> (AP301000) form
<code>PXGrid</code>	<i>Attributes</i>	Is used to display a simple grid without a toolbar. The grid contains a predefined set of rows, which can be edited.	The <b>Attributes</b> grid on the Attributes tab item of the <a href="#">Non-Stock Items</a> (IN202000) form
	<i>Details</i>	Is used to render a detail grid in a master-detail data entry page. The grid has a toolbar that holds the default actions, such as <b>Refresh</b> , <b>Add</b> , <b>Remove</b> , <b>Fit to Screen</b> , and <b>Export to Excel</b> ; it can also hold custom actions. The grid has no caption and paging is allowed.	The grid on the <b>1099 Settings</b> tab item of the <a href="#">Accounts Payable Preferences</a> (AP101000) form
	<i>Inquire</i>	Is used to display data without rows being added or removed. The grid has a toolbar that contains the <b>Refresh</b> , <b>Fit to Screen</b> , and <b>Export to Excel</b> default actions and can contain custom actions. The grid has no caption, and paging is allowed.	The grid on the <b>Attributes</b> tab item of the <a href="#">Customers</a> (AR303000) form
	<i>Primary</i>	Is used to display an editable primary grid that does not contain its own toolbar. To work with the grid, the user applies the action buttons of the form toolbar. The grid has no caption, and paging is allowed.	The grid on the <a href="#">Entry Types</a> (CA203000) form
	<i>PrimaryInquire</i>	Is used to display a primary grid without the availability to edit data. The grid does not contain its own toolbar. To work with the grid, the user applies the action buttons of the form toolbar, which does not contain the <b>Add</b> , <b>Delete</b> , and <b>Switch Between Grid and Form</b> buttons. The grid has no caption, and paging and filtering are allowed.	The grid on the <a href="#">Release AP Documents</a> (AP501000) form
	<i>ShortList</i>	Is used to display a small grid with a few records inside a form view. The grid has a toolbar that contains the <b>Refresh</b> , <b>Add</b> , and <b>Remove</b> default actions.	The <b>Sales Categories</b> grid on the Attributes tab item of the <a href="#">Non-Stock Items</a> (IN202000) form

UI Element	SkinID	Description	Example
PXPanel	<i>Buttons</i>	Is used in dialog boxes to display a horizontal row of buttons with right alignment.	The group of buttons in the <b>Add PO Receipt</b> dialog box, which opens if you click the <b>Add PO Receipt</b> button in the toolbar of the <b>Document Details</b> tab item of the <a href="#">Bills and Adjustments</a> (AP301000) form
	<i>Transparent</i>	Is used to group controls in a form container. The panel has no caption.	The group of controls on the <b>Template Settings</b> tab item of the <a href="#">Order Types</a> (SO201000) form

## Use of the Caption Property of Containers

If you plan to use a container in the mobile site map, we recommend that you specify a unique name for the `Caption` property of the container. Then in the mobile site map, you can refer to the container by the specified caption in the `Name` attribute of the `<sm:Container>` tag, as the following code shows.

```
<sm:Container Name="ValueOfTheCaptionProperty">
...
</sm:Container>
```

Otherwise, in the WSDL schema, the API web service assigns to the container the name of the first child element. If you use this name in the mobile site map, an error may occur after the container content is reordered because the container name might be changed in the WSDL schema.

## Use of Multiple Data Views for Boxes in Containers

The Acumatica Customization Platform supports the use of multiple data views for controls in the same container of an ASPX page.

For example, you can create a container and set the `DataMember` property to the name of the data view that provides most of the fields for boxes in the container. If you also want to create a control for a data field that cannot be accessible through that data view but can be accessible through another data view of the same graph specified in the `TypeName` property of the `PXDataSource` control, you can specify the required data view immediately in the `DataField` property, as follows.

```
<px:PXField ... DataField="DataViewName.FieldName" />
```

The following code snippet shows how to use the `MyDataView`, `AnotherDataView`, and `OnceMoreDataView` data views declared in the same graph or in extensions for the graph to define boxes for data fields in a `PXFormView` container on an Acumatica ERP form.

```
<px:PXFormView ... DataMember="MyDataView" ...>
...
<px:PXNumberEdit ... DataField="MyField_01" />
<px:PXSegmentMask ... DataField="MyField_02" />
<px:PXDateTimeEdit ... DataField="AnotherDataView.FieldName" />
<px:PXTextEdit ... DataField="MyField_05" />
<px:PXSelector ... DataField="OnceMoreDataView.OtherFieldName" />
...
</px:PXFormView>
```

## Configuring Tables

In this chapter, you can learn how to configure tables (grids) on webpages by using the `PXGrid` and `PXGridColumn` ASPX objects.

### In This Chapter

- [Use of the `SyncPosition` Property of `PXGrid`](#)
- [Use of the `DisplayMode` Property of `PXGridColumn`](#)
- [Use of the `Type` Property of `PXGridColumn`](#)

### Use of the `SyncPosition` Property of `PXGrid`

If a form contains a grid and the form toolbar includes an action to process a single record that is highlighted in the grid, the action delegate method must have a reference to the highlighted record in the cache.

To use the `Current` property of a `PXCache` object to access the record highlighted in a grid, the `Current` property must be synchronized with record highlighting in the grid. To force the system to provide this synchronization, you have to set the `SyncPosition` property of the `PXGrid` container to `True`.



: If you need to make an action button on the toolbar unavailable, when a grid is empty, you should set the `DependOnGrid` property of the appropriate `PXDSCallbackCommand` object in the `PXDataSource` control to the value that is specified in the `ID` property of the `PXGrid` element.

### Use of the `DisplayMode` Property of `PXGridColumn`

The Acumatica Customization Platform supports the following values for the `DisplayMode` property of a column in a grid.

Value	Description
<i>Value</i>	Default value. An indicator that the column cell contains the value of the field.
<i>Text</i>	If there is a description defined for the field, an indicator that the column cell contains the description of the field.
<i>Hint</i>	If there is a description defined for the field, an indicator that the column cell contains both the value and the description of the field.



: The priority of the `Type` property is higher than the priority of the `DisplayMode` property. If the `Type` property is set, for example, to `CheckBox`, the `DisplayMode` property is ignored.

### Use of the `Type` Property of `PXGridColumn`

The Acumatica Customization Platform supports the following values for the `Type` property of a column in a grid.

Value	Description
<i>NotSet</i>	The default value. An indicator that the field value is displayed in the column as a plain string that is formed based on the field data format.
<i>CheckBox</i>	An indicator that the field value is displayed in the column as a check box, which is selected if the field value is <code>True</code> .
<i>HyperLink</i>	An indicator that the field value is displayed in the column as a hyperlink.
<i>DropDownList</i>	An indicator that the column cell is rendered as a drop-down list that contains all the values specified for the referred data field.

Value	Description
<i>Icon</i>	An indicator that the field value contains an image URL and is displayed in the column as the referred image.

### Example

The following code fragment defines the grid columns on the [Automation Schedule Statuses](#) (SM205030) form.

```
...
<px:PXGridColumn AllowUpdate="False" DataField="LastRunStatus" Width="40px"
  Type="Icon" TextAlign="Center" />
<px:PXGridColumn DataField="ScreenID" DisplayFormat="CC.CC.CC.CC"
  Label="Screen ID" LinkCommand="AUScheduleExt_View" />
<px:PXGridColumn DataField="Description" Label="Description" Width="200px" />
<px:PXGridColumn AllowNull="False" DataField="IsActive" Label="Active"
  TextAlign="Center" Type="CheckBox" Width="60px" />
...
```

In the code, the `Type` property for the `LastRunStatus` data field (which corresponds to the **Status** column shown in the screenshot below) is set to `Icon`. Because the field value contains the image URL, the column cell displays the referred image.

For the `IsActive` data field (which corresponds to the **Active** column), the `Type` property is set to `CheckBox`. As you can see in the screenshot, the column cells are rendered as check boxes.

Revision Two HQ - Automation Schedules					
VIEW SCREEN INITIALIZE SCHEDULER					
	Status	* Screen ID	* Description	Active	* Starts On
		<a href="#">SM 20 50 60</a>	Send Reports	<input type="checkbox"/>	8/24/2014
		<a href="#">SM 50 70 00</a>	Process pending email	<input type="checkbox"/>	8/19/2014
>		<a href="#">SM 50 70 10</a>	Send/Receive Emails	<input type="checkbox"/>	3/19/2013
		<a href="#">SM 50 70 10</a>	Send and receive email	<input type="checkbox"/>	3/20/2014
		<a href="#">SM 50 70 10</a>	Send and receive email	<input type="checkbox"/>	8/19/2014

Figure: Viewing different types of columns on the Automation Schedules form

## Configuring Tabs

In this chapter, you can learn how to configure tables on webpages by using the `PXTab` and `PXTabItem` ASPX objects.

### In This Chapter

- [Conditional Hiding of a Tab Item](#)

## Conditional Hiding of a Tab Item

You can use the `Visible` property of the `PXTabItem` element to set the visibility of the tab item. However, if you need to set the dependency of a tab item's visibility from a condition, you should use the `VisibleExp` and `BindingContext` properties of the `PXTabItem` element.

### VisibleExp Property

The `VisibleExp` property contains a condition expression that defines a `Boolean` value used to set the visibility of the tab item. The expression must consist of two parts and an operator to compare these



parts. The expression can contain the values of controls that belong to the container specified in the `BindingContext` property.

### BindingContext Property

The `BindingContext` property defines the ID of the container of the controls whose values can be used in the expression of the `VisibleExp` property.

### Example

As an example of the conditional hiding of a tab item, on a form with form and tab containers, if you need to set the visibility of a tab item to depend on a check box of the form container, you can define the `VisibleExp` and `BindingContext` properties of the `PXTabItem` element, as illustrated in the following ASPX code snippet.

```
<px:PXFormView ID="form" ...>
...
  <px:PXCheckBox ... ID="myControlID" ... />
...
</px:PXFormView>
...
<px:PXTab ...>
...
  <px:PXTabItem ... BindingContext="form" ...
    VisibleExp="DataControls[&quot;myControlID&quot;].Value == true">
...
</px:PXTab>
```

In the code above, the expression uses the `DataControls` .NET property of the `form` object as a dictionary to find the needed control by the specified ID.

## Configuring Boxes

---

In this chapter, you can find information about the configuration of different types of boxes, such as `PXTextEdit`, `PXCheckBox`, and `PXGroupBox`.

### In This Chapter

- [Use of the `CommitChanges` Property of Boxes](#)
- [Use of the `DataField` Property of `PXGroupBox`](#)
- [Use of the `Caption` Property of `PXGroupBox`](#)
- [Use of the `RenderStyle` Property of `PXGroupBox`](#)

## Use of the `CommitChanges` Property of Boxes

If you need to process the value in a box every time the user changes this value, you need to set the `CommitChanges` property of the box to `True` to enable callbacks for the box.

If callback is enabled for a box in a container on a page, the user has changed the box value, and focus is no longer on the box on the page, the container immediately collects all the modified data and a callback is created to pass the data to the `PXDataSource` control of the page (see the diagram below).

The `PXDataSource` control creates a remote procedure call to the application server to execute the *Update* operation with the modified data on the data view that is specified as the `DataMember` property for the container. The data view executes the sequence of events for update of a data record (for details, see [Sequence of Events: Update of a Data Record](#)) on the data in the cache object of the business logic controller. The cache object raises the events that you can handle to process the modified data.

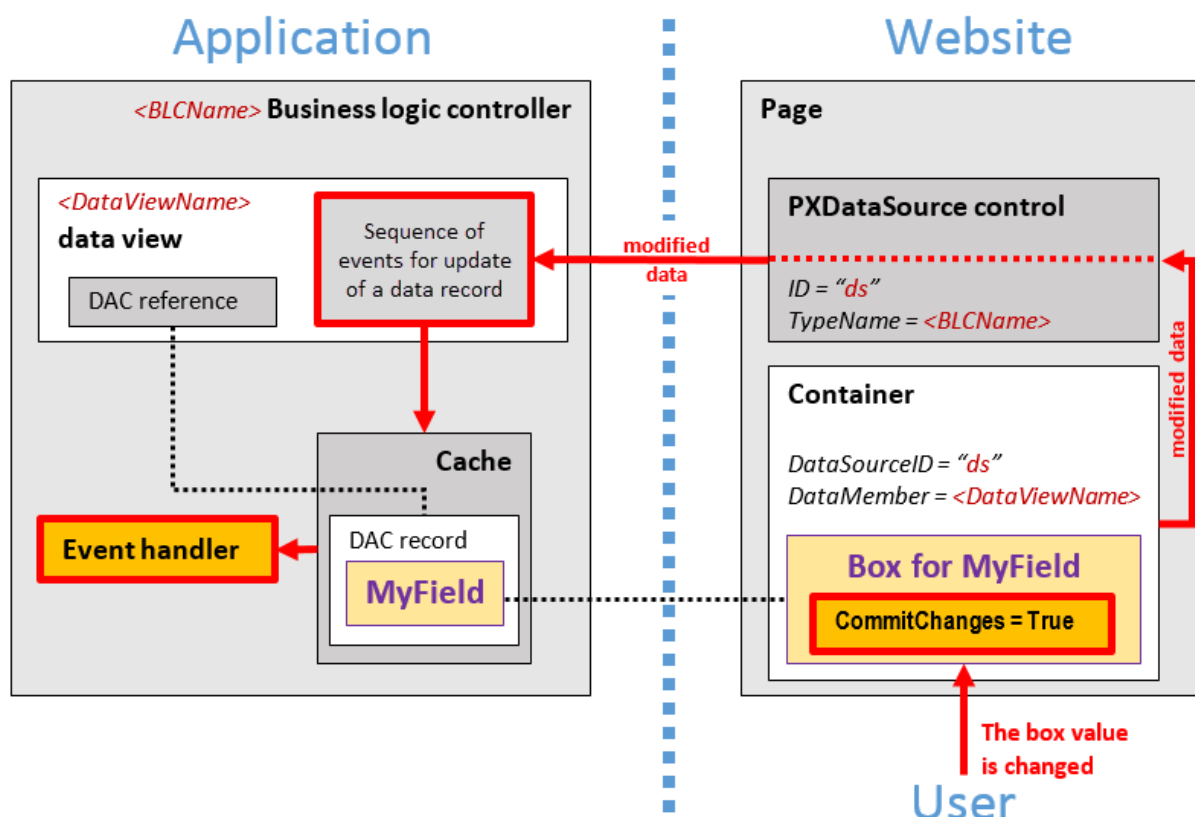


Figure: Usage of the CommitChanges property to process a modified data

## Use of the DataField Property of PXGroupBox

You use a group box to display a data field with a list attribute as a set of radio buttons, where one radio button is used to display and select each single constant value of the field. To bind a group box to a data field, you have to specify the name of the data field in the `DataField` property of the `PXGroupBox` element in the ASPX code, as follows.

```
<px:PXGroupBox ... DataField="<Field Name>" ...>
```



**Attention:** The group box must contain a radio button for each value defined in the list of the field.

In the `DataField` property of the `PXGroupBox` element, you can specify the name of a data field that is accessible through another data view of the same graph. See [Use of Multiple Data Views for Boxes in Containers](#) for details.

## Use of the Caption Property of PXGroupBox

You can define a caption for a group box by using the `Caption` property of the `PXGroupBox` element in the ASPX code as follows.

```
<px:PXGroupBox ... Caption="Example of Group Box Caption" ...>
```

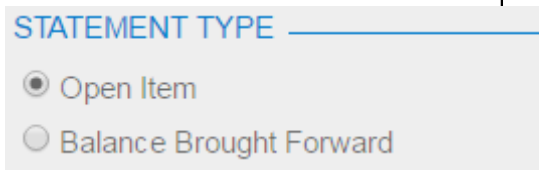
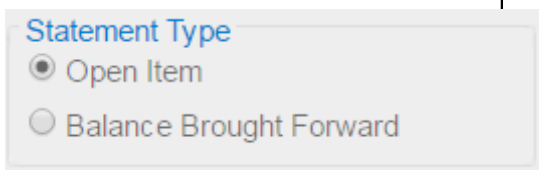

If the `RenderStyle` property of a `PXGroupBox` element is set to *Simple*, the `Caption` property is ignored. See [Use of the RenderStyle Property of PXGroupBox](#) for details.

## Use of the RenderStyle Property of PXGroupBox

To define the style of a group box on the form, you have to select a value of the `RenderStyle` property of the `PXGroupBox` element in the ASPX code, as follows.

```
<px:PXGroupBox ... RenderStyle="StyleName" ...>
```

The Acumatica Framework supports the following `RenderStyle` values for the `PXGroupBox` element.

Name	Description	Example
Fieldset	Indicates that the group of radio buttons can be displayed with a caption in the same style as in a grouping layout rule.	
RoundBorder	The default value. Indicates that the group of radio buttons can be displayed with a caption in a rounded border.	
Simple	Indicates that the group of radio buttons can be displayed without a caption and border.	

## Configuring Layout and Size

In this chapter, you can learn how to configure the layout of webpages by using the `PXLayoutRule` ASPX object and how to configure the size of ASPX controls.

### In This Chapter

- [Predefined Size Values](#)
- [Use of the StartRow and StartColumn Properties of PXLayoutRule](#)
- [Use of the ColumnWidth, ControlSize, and LabelsWidth Properties of PXLayoutRule](#)
- [Use of the ColumnSpan Property of PXLayoutRule](#)
- [Use of the Merge Property of PXLayoutRule](#)
- [Use of the GroupCaption, StartGroup, and EndGroup Properties of PXLayoutRule](#)
- [Use of the SuppressLabel Property of PXLayoutRule](#)

## Predefined Size Values

You can use the predefined values described in the table below for the following properties:

- `ColumnWidth`, `LabelsWidth`, and `ControlSize` of the `PXLayoutRule` component
- `LabelsWidth` and `Size` of a control

### Predefined Values

The following table shows the values in pixels that correspond to the predefined constants.

Predefined Value	ColumnWidth	LabelsWidth and ControlSize of a Layout Rule; LabelsWidth and Size Properties of a Control
XXS	100px	40px
XS	150px	70px
S	200px	100px
SM	-	150px
M	250px	200px
XM	300px	250px
L	350px	300px
XL	400px	350px
XXL	450px	400px

### Setting of Predefined Values

Note the following points about setting the predefined sizes of controls and their labels:

- For any property for which there are predefined values, you can specify a value in pixels, such as *55px*. (This format is mandatory if you don't use abbreviations, because the property value can be defined only in pixels.)
- There is no predefined value for the `Width` property of a control. Therefore, you can specify a value for this property by typing any value in pixels, such as *55px*. Before specifying the `Width` property value for a control, you must define the `Size` property value for the control as *Empty*.




: The `Width` property is declared in ASP.NET. The `Size` property is declared in Acumatica Framework, so you can use the predefined values.

## Use of the `StartRow` and `StartColumn` Properties of `PXLayoutRule`

In this topic, you can find information how to organize controls on a webpage into rows and columns.

### Default Layout

By default, the system places all the controls of a container into a column within the first row, as shown in the diagram below. To do this, the system initially sets to *True* the `StartRow` property value for the uppermost `PXLayoutRule` component in a container.



**Figure:** The default layout of controls of a container on a form

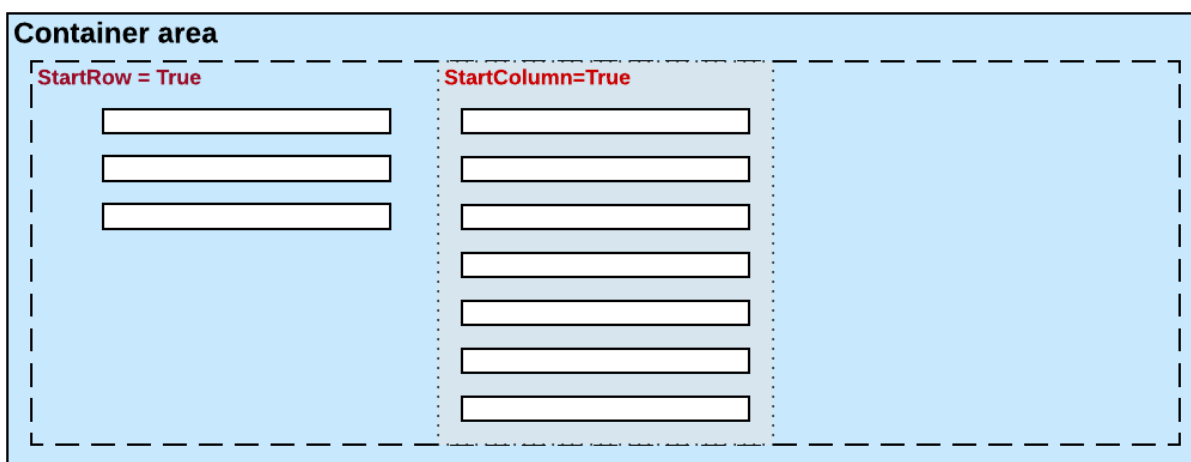
The controls continue to be placed within a single column until you add a layout rule with the `StartColumn` or `Merge` property value set to *True*.



**Important:** For the proper layout, the `StartRow` property value must be set to *True* for the uppermost `PXLayoutRule` component of a container.

### Splitting of Controls into Columns

You can place controls in multiple columns within a row by adding `PXLayoutRule` components with the `StartColumn` property value set to *True*. This property creates a new column of controls within the current layout row, as the following diagram shows.



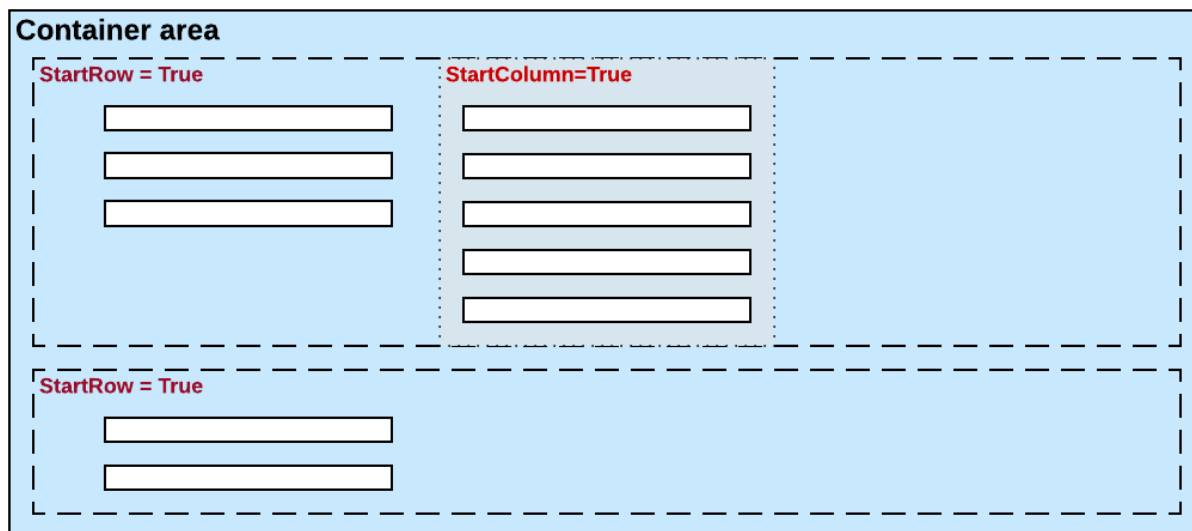
**Figure:** Creation of a new column

The first control under this rule is the highest control in the column.

### Splitting of Controls into Rows

Every new `PXLayoutRule` component that has the `StartRow` property value set to *True* initializes a new independent placeholder of controls, which are placed in a single column by default. To place controls in

multiple columns within the new row, you should include in the placeholder a new layout rule with the `StartColumn` property value set to `True`, as shown in the following diagram.



**Figure: Creation of a new row**

### Sizes of Rows and Columns

Because the values of the `ColumnWidth`, `ControlSize`, and `LabelsWidth` properties are never inherited from the previously declared `PXLayoutRule` component, you might need to define these properties exclusively for every new row and column.

## Use of the `ColumnWidth`, `ControlSize`, and `LabelsWidth` Properties of `PXLayoutRule`


You can use the `PXLayoutRule` components to define the sizes for every control (that is, its input area) and its label within a column, or merged set of controls.

### Required Properties

Every `PXLayoutRule` component that has the `StartRow` or `StartColumn` property value set to `True` must have one of the following sets of properties defined:

- `LabelsWidth` and `ControlSize`
- `LabelsWidth` and `ColumnWidth`

The following diagram illustrates the meaning of the `LabelsWidth`, `ControlSize`, and `ColumnWidth` properties.



**Figure: Use of the layout rule properties to define control sizes**



: You should not set property values for both `ColumnWidth` and `ControlSize` for the same `PXLayoutRule` component; in this case, the system would use the value of the `ControlSize` property.

## Setting of the Size

Please note the following points about setting the sizes of controls and their labels:

- The values of the `ColumnWidth`, `ControlSize`, and `LabelsWidth` properties must be defined exclusively for every `PXLayoutRule` component; they are never inherited from the previously declared one.
- You can change the size of a single control or its label by defining values for the `Size`, `Width`, and `LabelsWidth` properties of the control. Property values that are set for a control have a higher priority than the property values of the `PXLayoutRule` component.
- You can assign a predefined size abbreviation (such as `XXS`, `L`, or `XL`) for the `ColumnWidth`, `LabelsWidth`, and `ControlSize` properties of a layout rule and the `LabelsWidth` and `Size` properties of a control. (See [Predefined Size Values](#) for details.)
- The `PXDateTimeEdit` and `PXNumberEdit` control types have a predefined `Width` property value, which you cannot change by setting the `ColumnWidth` or `ControlSize` property values for the appropriate `PXLayoutRule` component. To change the width of this control, set a value for the `Size` or `Width` property of the control.

## Use of the ColumnSpan Property of PXLayoutRule

You specify the `ColumnSpan` property value for a `PXLayoutRule` component by manually typing the number of columns spanned by the first control placed below the rule.

### Example

As an example of the use of the `ColumnSpan` property, the form container on the [Customers](#) (AR303000) form has three columns of boxes, and there is a layout rule with the `ColumnSpan` property set to 2 in the first column. This property forces the system to make the box span two columns, as shown in the following screenshot.

The screenshot shows the 'Customers' form (AR303000) with a layout rule that has `ColumnSpan` set to 2. The form is divided into three columns. The first column contains a 'Customer ID' field and a 'Customer Name' field. The second column contains a 'Status' dropdown menu. The third column contains 'Balance' and 'Prepayments Balance' fields. The 'Customer Name' field is highlighted with a red box, indicating it spans two columns. The form also has a 'General Info' tab selected, and other tabs like 'Billing Settings', 'Delivery Settings', 'Payment Methods', 'Contacts', 'Salespersons', 'Attributes', 'Activities', 'GL Accounts', and 'Mailing Settings' are visible.

**Figure: The box that spans two columns on the form**

### Dependencies

A `PXLayoutRule` component with the `ColumnSpan` property value specified is handled as follows:

- The `LabelsWidth` property value is always inherited from the previously declared `PXLayoutRule` component that has the `StartRow` or `StartColumn` property value set to `True`.
- If a value for the `ColumnWidth` or `ControlSize` property is specified for the component, this value is ignored.

## Use of the Merge Property of PXLayoutRule

Merging means placing controls so that they are horizontally aligned. Horizontal alignment is performed for the controls that are placed between a layout rule with the `Merge` property set to `True` and any other subsequent layout rule. Therefore, to cancel merging for all of the following controls, you have to add a `PXLayoutRule` component with or without the `Merge` property specified.

## Example

As an example of the use of the `Merge` property, the **Billing Settings** tab item on the [Customers](#) (AR303000) form has three pairs of merged check boxes in the **Print and Email Settings** group. This property forces the system to render the boxes in one column, as shown in the following screenshot.

The screenshot shows the 'Billing Settings' tab for a customer. The 'PRINT AND EMAIL SETTINGS' section contains three pairs of checkboxes, each pair enclosed in a red box. The pairs are: 'Send Invoices by Email' and 'Print Invoices', 'Send Dunning Letters by Email' and 'Print Dunning Letters', and 'Send Statements by Email' and 'Print Statements'. The checkboxes are arranged in a single column, demonstrating the effect of the `Merge` property.

**Figure:** The boxes merged into a single column on the form

## Dependencies

A `PXLayoutRule` component with the `Merge` property value set to `True` is handled as follows:

- If the `ColumnWidth` property value is set for the same `PXLayoutRule` component, the value is ignored.
- The default values for the `ControlSize` and `LabelsWidth` properties are inherited from the previously declared `PXLayoutRule` component with the `StartRow` or `StartColumn` property value set to `True`. You can override these property values, if necessary, by specifying the `ControlSize` and `LabelsWidth` property values from the predefined list of options. (See [Predefined Size Values](#) for details.)

## Use of the GroupCaption, StartGroup, and EndGroup Properties of PXLayoutRule

You can organize controls in a container within groups to make users' work more logical.

### Grouping of Controls

To group multiple controls within a column, generally you have to add two `PXLayoutRule` components that have the following properties set to define the first and the last controls in the group, respectively:

- `GroupCaption` and `EndGroup`: To create a group with the caption specified in the `GroupCaption` property
- `StartGroup` and `EndGroup`: To create a group without a caption

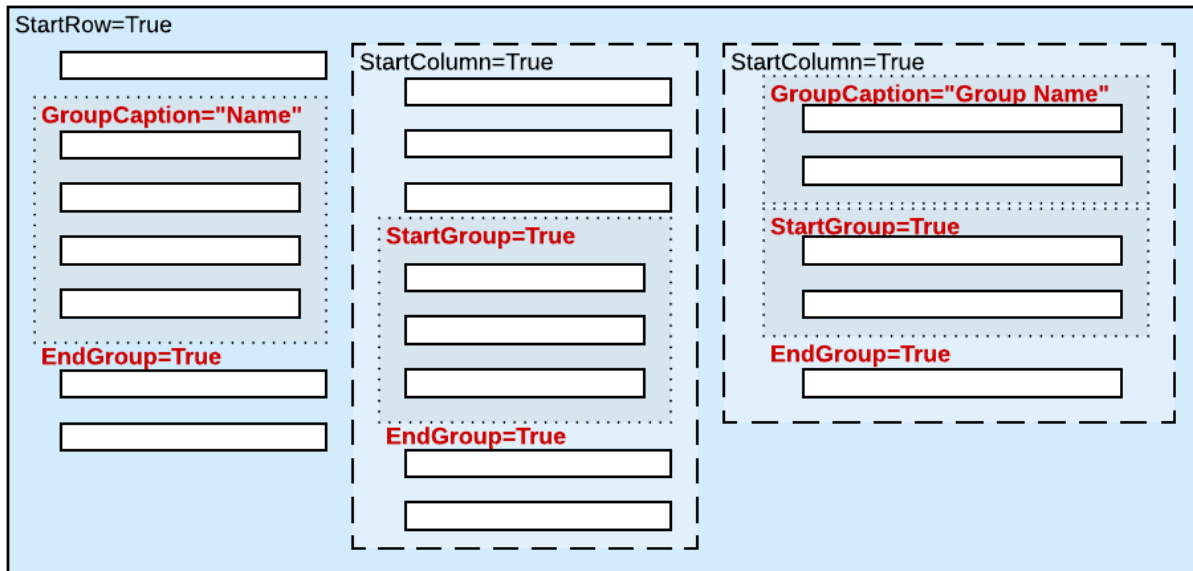




: You can specify both the `GroupCaption` property and the `StartGroup` property for a `PXLayoutRule` component that starts a group.

For example, by specifying the `GroupCaption` property value for the corresponding `PXLayoutRule` components placed above a control, you start the group of controls and set up the header for the group. You should also add a `PXLayoutRule` component with the `EndGroup` property value set to `True` below (in the code) the last control that is included in the group.

You end a group by using a `PXLayoutRule` component with a `GroupCaption`, `StartGroup`, or `EndGroup` property specified. Therefore, if there is another group that starts immediately below a group, you can omit the layout rule that ends the upper group, as shown in the third column of the row displayed in the example in following diagram.



**Figure:** Possible use of layout rules with grouping properties

## Dependencies

The system works as follows for all `PXLayoutRule` components with the `GroupCaption` or `StartGroup` property value specified:

- If the `GroupCaption`, `StartGroup`, or `EndGroup` property is set for a `PXLayoutRule` component, the system ignores the `ColumnWidth` property value specified for the component.
- The default values for the `ControlSize` and `LabelsWidth` properties are inherited from the previously declared `PXLayoutRule` component with the `StartRow` or `StartColumn` property value set to `True`. You can override these property values, if necessary, by specifying the `ControlSize` and `LabelsWidth` property values in the layout rule that starts a group. (See [Predefined Size Values](#) for details.)

## Use of the SuppressLabel Property of PXLayoutRule

Every control for a data field contains both a label and the input area of the control. The label is displayed left of the input area, except with check boxes; the label of a check box is displayed right of the input area of the check box. When you add a check box to a form, the check box control is automatically aligned both left and right with other input controls in the appropriate column. As a result, the area of the form left of a check box is empty.

## SuppressLabel Property

To hide the labels of the controls placed within a column, you should set the `SuppressLabel` property value of the `PXLayoutRule` component of the column to `True`. Then within the column, all check boxes are placed without any space to the left of the input control.



: If needed, you can left-align a check box in the column by setting to `True` the `AlignLeft` property of the control. Also, you can set the `SuppressLabel` property value to `True` for any other control to hide its label.

The `SuppressLabel` property affects all of the controls of the group that are placed under the `PXLayoutRule` component with the `True` value of this property. The `SuppressLabel` property value must be defined for every `PXLayoutRule` component for the controls placed beneath the component and included in the same column; this property is never inherited from the previously declared property.



: The `SuppressLabel` property value is never applied to `PXLayoutRule` components that have the `ColumnSpan` property value specified.

## Example

The **Parent Info** group on the **Billing Settings** tab item on the [Customers](#) (AR303000) form is initially displayed with **Parent Account** displayed, as shown in the following screenshot.

**Figure: A group of controls with labels**

If you set the `SuppressLabel` property of the group layout rule to `True`, the label of the **Parent Account** box is hidden and all check boxes are displayed without any space to the left of the check boxes, as shown in the following screenshot.

**PARENT INFO**

☐ Consolidate Balance

☐ Consolidate Statements

☐ Share Credit Policy

**PRINT AND EMAIL SETTINGS**

☐ Send Invoices by Email ☐ Print Invoices

☒ Send Dunning Letters by Email ☐ Print Dunning Letters

☐ Send Statements by Email ☐ Print Statements

Statement Type:

☐ Multi-Currency Statements

**Figure:** The same group of controls after applying the `SuppressLabel` property to the group

## Maintaining Reports

In Acumatica Report Designer, you can create custom reports or modify existing reports and then use these reports in Acumatica ERP or an Acumatica Framework-based application. For more information on the creation of the custom reports with the Report Designer, see [Acumatica Report Designer Guide](#).

In this chapter, you can find details about how the system renders reports in Acumatica ERP.

### In This Chapter

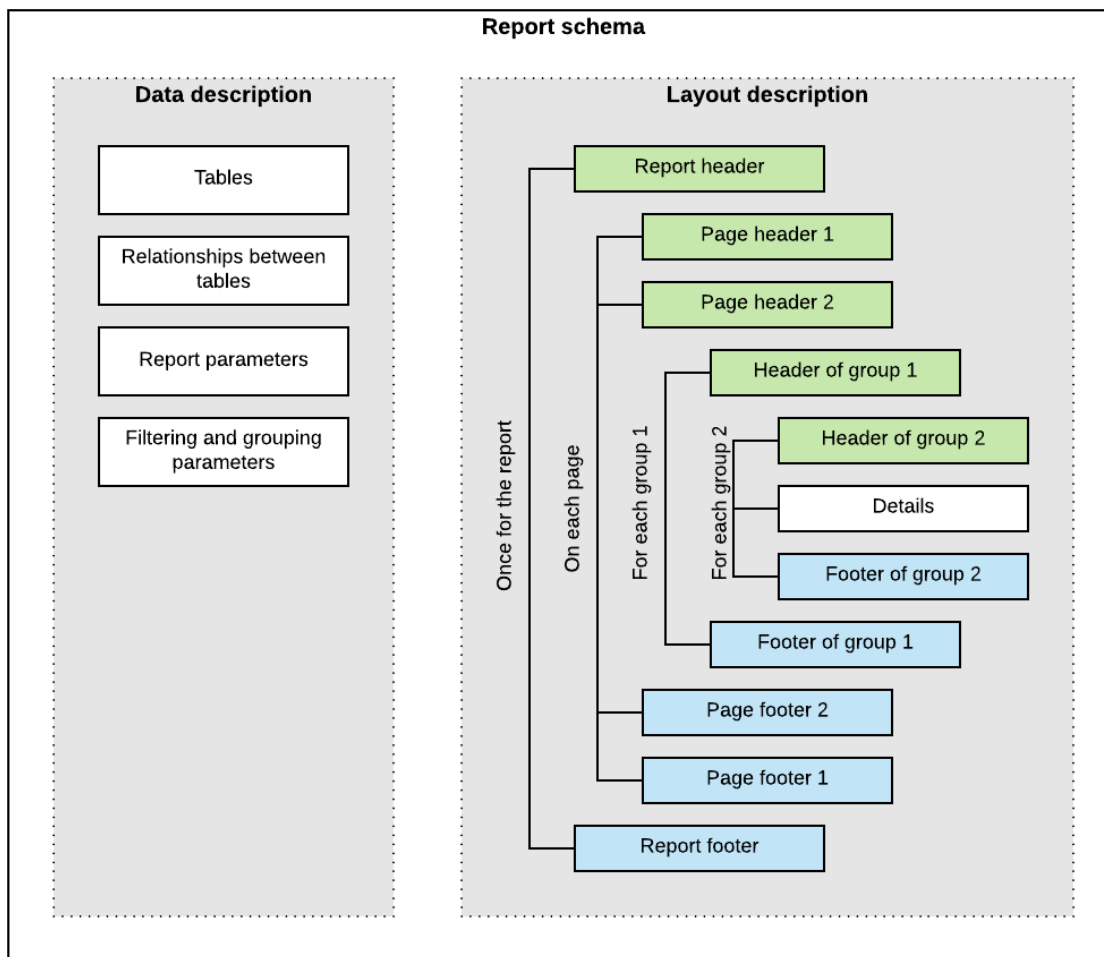
- [Display of Reports](#)
- [Display of Analytical Reports](#)

## Display of Reports

In this topic, you can find information about how Acumatica ERP displays reports that are created with the Acumatica Report Designer.

### What a Report Is

A report is an RPX file (which is created with the Acumatica Report Designer) that contains the report schema in XML format—that is, the description of the data that should be displayed in the report and the description of the report layout. (See the following diagram.) The description of the data of the report includes the following: the database tables that provide data for the report, the relationships between these tables, the parameters that can be specified before the report is run, and the filtering and grouping parameters. The report layout is a tree of headers, details, and footers.



**Figure: Report schema**

Reports can be saved in files on disk, or in the `UserReport` table of the database of an Acumatica ERP instance. (The `UserReport` table uses the file name of the report as a key and stores the schema of the report in XML format in the `Xml` column.)

The saved report can be published on an Acumatica ERP site—that is, added to the site map and to any applicable workspaces so that the users can work with the report.

For details on creating reports with the Acumatica Report Designer, see the [S130 Reporting: Inquiry, Report Writing, Dashboards](#) training course.

Acumatica ERP provides the following ways to run a report that you have created by using the Acumatica Report Designer:

- From a report form, which is added to the site map, when the user clicks the **Run Report** button on the form toolbar
- From a maintenance or entry form, when the user clicks the action button whose name is associated with the report name

## How the Report Is Launched from the Report Form

When a user opens the report form, the webpage performs the POST HTTP request to the `ReportLauncher.aspx` page, passing the name of the report file as the `ID` query string parameter of the request, as shown in the following example.

```
http://localhost/AcumaticaDB/frames/ReportLauncher.aspx?
id=YF123456.rpx&HideScript=On
```

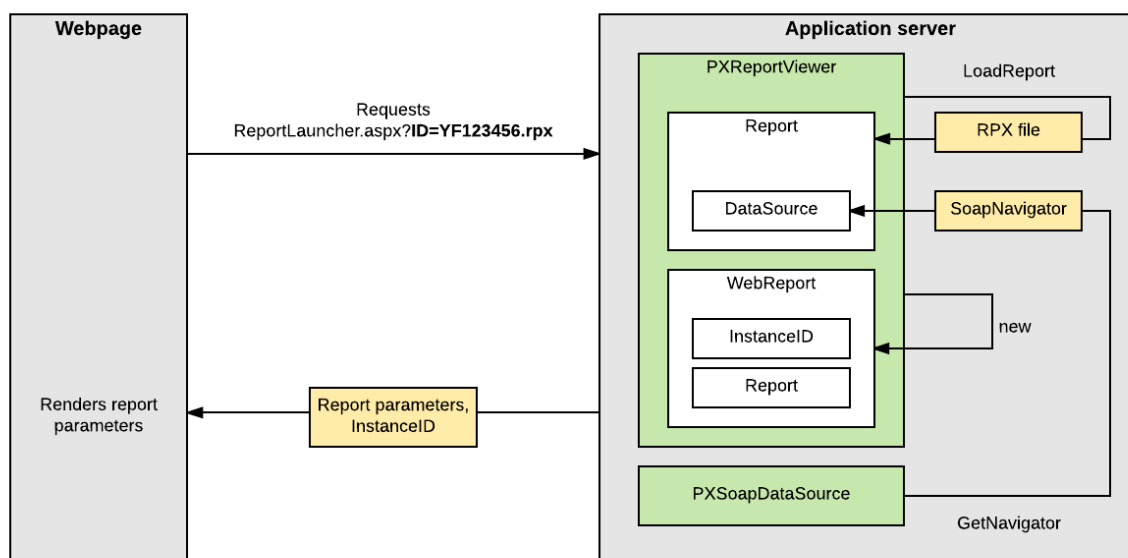
The `ReportLauncher.aspx` page contains the `PXReportViewer` control, whose JavaScript objects and functions are designed to obtain the report data and display the data on the form, and the `PXSoapDataSource` control, which is used to retrieve data for the report.

On the server side, an instance of the `PX.Web.UI.PXReportViewer` class processes the request as follows:

1. Loads the report schema from the file on disk or from the database (by using the `LoadReport` method) to a `PX.Reports.Controls.Report` object (which stores the report schema in memory and provides the methods for working with this schema).
2. If the report schema is loaded successfully, performs the following:
  - a. Instantiates a `PX.Reports.Web.WebReport` object that will store data of the launched report and assigns an instance ID to `WebReport`.
  - b. Binds the `Report` object to the data source that is specified by the `PXSoapDataSource` control of the ASPX page. The `PX.Web.UI.PXSoapDataSource` class instantiates a `SoapNavigator` object, which will be then used to retrieve data for the report from the database.

The server returns an XML response with the report parameters to display and with the ID of the instance of `WebReport` in the session. The browser displays the report parameters and other options on the report form.

The following diagram illustrates how the report form is launched.



**Figure: Launch of the report form**

## How the Report Data Is Retrieved

After the user has selected the values of the parameters of the report and clicked the **Run Report** button on the form toolbar of the report form, the webpage sends the GET request to `PX.ReportViewer.axd` on the server with the ID of the `WebReport` instance in the session (which was created when the report was launched), as shown in the following example.

```
http://localhost/AcumaticaDB/PX.ReportViewer.axd?
  InstanceID=10bf4a13a38c4af39cafc80926e407f2&OpType=Report
  &PageIndex=0&Refresh=True
```

To process the request, the server invokes the `Render` method of the `WebReport` class, which launches the generation of the report as a long-running operation in a separate thread. For details on how the long-running operation is processed for a report, see [Processing a Report as a Long-Running Operation](#).

To retrieve the data of the report from the database, the system uses the `PX.Data.Reports.SoapNavigator` object (to which a reference is stored in the `Report` object). `SoapNavigator` instantiates a `PXGraph` object (without a type parameter) and composes a BQL command as an instance of the `PX.Data.Reports.BqlSoapCommand` class. `BqlSoapCommand` inherits from the `PX.Data.BqlCommand` class and is optimized for retrieving data for the reports. `BqlSoapCommand` has the `IndexReportFields` method, which uses the `PX.Data.PXDependsOnFieldsAttribute` attribute to get the dependent fields of the report recursively. For details on how BQL commands are used to retrieve data from the database, see [Translation of a BQL Command to SQL](#).

The system processes the data and creates a `ReportNode` object. That is, the system creates the sections of the report based on the data retrieved from the database and on the report schema from the `Report` object, and calculates all formulas inside the sections. Then the system uses the resulting `ReportNode` object, which contains all sections with all needed values, to render data in the needed format.

## How the Report Data Is Displayed

After the long-running operation has completed, the `PXReportViewer` control displays the report on the report form.

If a report is displayed in HTML format and the user turns the pages of the report, the webpage sends the GET request to `PX.ReportViewer.axd`. The query string parameters of the request are the ID of the report in the session and the number of the page, as shown in the following request URL.

```
http://localhost/AcumaticaERP/PX.ReportViewer.axd?
  InstanceID=008f2a8afb1a4998ade98bc64fc30ad9&OpType=Report
  &PageIndex=0
```

The format in which the report is displayed (either PDF or HTML) is specified in the `OpType` query string parameter of the request, as shown in the following request URL.

```
http://localhost/AcumaticaERP/PX.ReportViewer.axd?
  InstanceID=008f2a8afb1a4998ade98bc64fc30ad9&OpType=PdfReport&Refresh=True
```

If a user turns the pages of the report or changes the format of the report, the system creates the results of the report from the `ReportNode` object stored in the session by using the renderer for the needed format. That is, the system does not retrieve the data of the report from the database and does not process this data to create a `ReportNode` object once again.

## How the Report Is Launched from the Maintenance or Entry Form

On a form, when a user clicks the action button to generate a report, the data source control of the form creates a request to the Acumatica ERP server to execute the action delegate defined for the button. The server creates an instance of the graph, which provides the business logic for the form and invokes the action delegate method. The action delegate obtains from the form the data required to

define the report parameters and throws an exception of the `PX.Data.PXReportRequiredException` type with the report ID and these parameters. The system processes the exception, saves the report parameters to the session, and redirects the user to the `ReportLauncher.aspx` page.

The `ReportLauncher.aspx` page loads the report schema and instantiates a `WebReport` object, as described in [How the Report Is Launched from the Report Form](#). Instead of retrieving the values of report parameters from the webpage, the system opens the report with the parameters stored in the session. The system retrieves data for the report, as described in [How the Report Data Is Retrieved](#).

## Display of Analytical Reports

In this topic, you can find information about how Acumatica ERP displays analytical reports.

### What an Analytical Report Is

An analytical report is a specific type of report that you can construct in Acumatica ERP by using the Analytical Report Manager (ARM) toolkit. You will likely want to use the ARM toolkit rather than the Acumatica Report Designer to create the following types of reports:

- Financial reports that display data that is posted to the general ledger accounts and accumulated in the General Ledger module. The system gets the data of the general ledger accounts from the `GLHistory` table of the database.
- Project accounting reports that display data that is accumulated in the Project module. The system gets the data for these reports from the `PMHistory` table of the database.

For details on analytical reports, see the [F350 Reporting: Analytical Reports](#) training course.

The schema of analytical reports are stored in a set of database tables that have `RM` prefix in their names, such as `RMReport`, `RMSRowSet`, and `RMColumnSet`. The analytical report is identified in the system by its code, which you specify in the **Code** box on the [Report Definitions](#) (CS206000) form. The report code is stored in the `ReportCode` column of the `RMReport` table.

The schema of an analytical report can include the position of the report in the site map so that the users can work with the report. To run the report, a user clicks the **Run Report** button on the form toolbar of the report form.

### How the Analytical Report Is Launched from the Report Form

When a user opens the report form of an analytical report, the webpage performs the POST HTTP request to the `RMLauncher.aspx` page, passing the report code with `.rpx` appended as the `ID` query string parameter of the request, as shown in the following example.

```
http://localhost/AcumaticaDB/frames/rmlauncher.aspx?id=dbsp.rpx&HideScript=On
```

The `RMLauncher.aspx` page contains the `PXReportViewer` control, whose JavaScript objects and functions are designed to obtain the report data and display the data on the form, and the `ARmDataSource` control, which is used to retrieve data for the analytical report.

On the server side, the system (by using an instance of the `PX.Web.UI.PXReportViewer` class) processes the request as follows:

1. Loads the parameters of the analytical report from the database (by using the `LoadReport` method) to a `PX.Reports.Controls.Report` object as follows:
  - a. The system retrieves the data source of the report viewer as an instance of the `PX.CS.RMReportReader` class, which is a graph (derived from `PXGraph<RMReportMaint, RMReport>`) that implements the `PX.Report.ARM.Data.IARmDataSource` interface.
  - b. By using the `GetReport` method of `IARmDataSource`, the system retrieves the parameters of the analytical report from the database to the `PX.Reports.ARM.ARMReport` object.

- c. By using the `CreateReport` method of the `PX.Reports.ARM.Data.ARMProcessor` class, the system creates a `Report` object with the parameters retrieved from the `ARMReport` object.
- 2. If the report parameters are loaded successfully, performs the following:
  - a. The system instantiates a `PX.Reports.Web.WebReport` object that will store data of the launched report and assigns an instance ID to `WebReport`.
  - b. The system initializes the `ProcessMethod` field of the `WebReport` object with the processing function for analytical reports.

The server returns an XML response with the report parameters to display and with the ID of the instance of `WebReport` in the session. The browser displays the report parameters on the report form.

### How the Data of the Analytical Report Is Retrieved

After the user has selected the values of the parameters of the analytical report and clicked the **Run Report** button on the form toolbar of the report form, the webpage sends the GET request to `PX.ReportViewer.axd` on the server with the ID of the `WebReport` instance in the session (which was created when the report was launched).

To process the request, the server invokes the `Render` method of the `WebReport` class, which launches the generation of the report as a long-running operation in a separate thread. For details on how the long-running operation is processed for a report, see [Processing a Report as a Long-Running Operation](#).

To retrieve the data of the analytical report from the database, the system uses the `PX.Objects.CS.RMReportReaderGL` (for financial reports) and `PX.Objects.CS.RMReportReaderPM` (for project accounting reports) extensions of the `PX.CS.RMReportReader` graph.

### How the Report Data Is Displayed

After the long-running operation has completed, the `PXReportViewer` control displays the analytical report on the report form.



# Accessing Data

---

The topics in this part of the guide explain how an application based on Acumatica Framework can access data from the application database and the data stored in the session.

## In This Part

- [Querying Data in Acumatica Framework](#)
- [Creating Fluent BQL Queries](#)
- [Creating Traditional BQL Queries](#)
- [Creating LINQ Queries](#)
- [Defining Relationships Between DACs](#)
- [Working with Data in Cache and Session](#)

## Querying Data in Acumatica Framework

---

In Acumatica Framework, you generally use business query language (BQL) to query data from the database. BQL statements represent specific SQL queries and are translated into SQL by Acumatica Framework, which helps you to avoid the specifics of the database provider and validate the queries at the time of compilation. Acumatica Framework provides two dialects of BQL: traditional BQL and fluent BQL. For details about these dialects, see [Creating Traditional BQL Queries](#) and [Creating Fluent BQL Queries](#).

To query data from the database, you can also use language-integrated query (LINQ), which is a part of .NET Framework. In the code of Acumatica Framework-based applications, you can use both the standard query operators (provided by LINQ libraries) and the Acumatica Framework-specific operators that are designed to query database data.

This chapter explains the aspects that are common to traditional BQL, fluent BQL, and LINQ and provides a high-level comparison of the approaches for querying data in Acumatica Framework.

## In This Chapter

- [BQL and LINQ](#)
- [Data Access Classes](#)
- [PXView and PXCache of the Data View](#)
- [Data Query Execution](#)
- [Translation of a BQL Command to SQL](#)
- [Merge of the Records with PXCache](#)
- [Comparison of Fluent BQL, Traditional BQL, and LINQ](#)
- [Fluent BQL and Traditional BQL Equivalents](#)
- [To Execute BQL Statements](#)
- [To Process the Result of the Execution of the BQL Statement](#)

## BQL and LINQ

When a data request occurs, the system creates an instance of a business logic controller (also referred as a *graph*). The graph contains the data views that you define in code. In these data views, you define

the queries to be executed to retrieve the requested data by using business query language (BQL), which is provided by Acumatica Framework. You also use BQL to define the data queries directly in code and in attributes.

BQL is written in C#; it is based on generic class syntax, which is similar to SQL syntax. Thus, BQL has almost the same keywords as SQL does, placed in the order in which they are used in SQL. BQL offers several benefits to the application developer. BQL does not depend on the specifics of the database provider, and it is object-oriented and extendable. Also, BQL provides compile-time syntax validation, which helps to prevent SQL syntax errors.

You can also use language-integrated query (LINQ) provided by the `System.Linq` library when you need to select records from the database in the code of Acumatica Framework-based applications or if you want to apply additional filtering to the data of a BQL query. However, you still have to use BQL to define the data views in graphs and to specify the data queries in the attributes of data fields.

### Fluent BQL and Traditional BQL

Acumatica Framework provides two dialects of BQL: fluent BQL and traditional BQL. Traditional BQL was the initial language for data queries in Acumatica Framework; it provides the benefits described above. Fluent BQL provides the following advantages as compared to traditional BQL:

- It is easier to read and edit fluent BQL queries than traditional BQL queries because each section of a fluent BQL query does not depend on the others and can appear in only specific places of the query. Also, fluent BQL queries contain fewer commas and angle brackets and do not use numbered classes (such as `Select2` or `Select6`).
- You do not need to select a suitable class for a fluent BQL query (such as `PXSelectOrderBy<,>` or `PXSelectJoinOrderBy<,>`); instead, you simply start typing the command, and IntelliSense in Visual Studio offers continuations that are relevant for the current query state.

For a detailed list of differences between the dialects, see [Comparison of Fluent BQL, Traditional BQL, and LINQ](#).

The following code shows an example of a data view written in fluent BQL.

```
SelectFrom<Product>.
    Where<Product.availQty.IsNotNull.
        And<Product.availQty.IsGreater<Product.bookedQty>>>.View products;
```

The following code shows the same data view written in traditional BQL.

```
PXSelect<Product,
    Where<Product.availQty, IsNotNull,
        And<Product.availQty, Greater<Product.bookedQty>>>> products;
```

Suppose the database provider is Microsoft SQL Server. Acumatica Framework translates the fluent and traditional BQL queries shown above into the following SQL query, where [list of columns] is the list of columns of the `Product` table.

```
SELECT [list of columns] FROM Product
WHERE Product.AvailQty IS NOT NULL
AND Product.AvailQty > Product.BookedQty
```

Acumatica Framework explicitly enumerates the columns of the database table in the SQL query. For details on which columns are enumerated, see [Translation of a BQL Command to SQL](#).

### LINQ

To configure a LINQ query, you can use the following variants of syntax:

- Query expressions, which use standard query operators from the `System.Linq` namespace (such as `where` or `orderby`) or Acumatica Framework-specific operators from the

PX.Data.SQLite namespace (such as `SQL.BinaryLen`, which is shown in the following example of this syntax).

```
ProductMaint graph = PXGraph.CreateInstance<ProductMaint>();
var goods = from p in graph.Select<Product>()
where
    p.ProductCD.Length == 5 &&
    p.GroupMask.Length == 4 &&
    (p.WorkGroupID & 0b10) != 0
select new
{
    p.ProductID,
    p.ProductCD,
    p.ProductName,
    Len = p.ProductName.Length,
    BLen = SQL.BinaryLen( p.ProductName) + 1,
    p.GroupMask,
    p.WorkGroupID
};
```

- Explicit (method-based) syntax. The arguments of the methods used in this syntax are lambda expressions. In these expressions, you can use the standard C# operators and Acumatica Framework-specific operators from the `PX.Data.SQLite` namespace (such as `SQL.BinaryLen`, which is shown in the following code). The code below is equivalent to the query expression shown above.

```
ProductMaint graph = PXGraph.CreateInstance<ProductMaint>();
var goods = graph.Select<Product>()
    .Where( p =>
        p.ProductCD.Length == 5 &&
        p.GroupMask.Length == 4 &&
        (p.WorkGroupID & 0b10) != 0)
    .Select( p => new
    {
        p.ProductID,
        p.ProductCD,
        p.ProductName,
        Len = p.ProductName.Length,
        BLen = SQL.BinaryLen(p.ProductName) + 1,
        p.GroupMask, p.WorkGroupID
    });
```

## Data Access Classes

The classes that represent database tables in Acumatica Framework are called *data access classes* (DACs). You derive these classes from the `PX.Data.IBqlTable` interface. The name of a class is usually the same as the name of the database table to which it provides access (except with the DACs that have the `PXTable` or `PXProjection` attributes, which change the default binding of DACs to database tables). We recommend that you add to each DAC the `PXCacheName` or `PXHidden` attribute to specify whether and how the DAC is available in generic inquiries and reports.

For each table column, you add a data field to the corresponding data access class by declaring the following members:

- A `public abstract` class (which is also referred to as a *class field*)

You use this class to reference the table column in a business query language (BQL) statement. The declaration of the class field is different in the fluent BQL dialect than it is in the traditional BQL dialect. For details about the declaration, see [Data Access Classes in Fluent BQL](#) and [Data Access Classes in Traditional BQL](#). We recommend that you use the fluent BQL style of DAC declaration because it can be used both in fluent BQL and in traditional BQL. The style of class field declaration is not important for queries defined with language-integrated query (LINQ).

- A `public virtual` property (which is also referred to as *property field*)

You bind the data field to the table column by specifying the type attribute that is derived from the `PXDBFieldAttribute` class, such as `PXDBString`, and specifying the name of the column as the name of the property. If you don't need to bind the property to a database column (for example, if you want the value of the property to be calculated from the database fields), you specify an unbound type attribute, such as `PXDBCalced`. You assign the property a name that starts with an uppercase letter. For the lists of bound and unbound type attributes, see [Bound Field Data Types](#) and [Unbound Field Data Types](#).

You use the property, which, in the system, holds the column data of the table, in the queries defined with LINQ. In the SQL command generated from BQL, the framework explicitly lists columns for all bound data fields defined in the DAC. For the unbound data fields whose property attribute defines a BQL command, if this data field is used in a BQL query, the system translates the BQL command of the property to SQL when the BQL query is translated to SQL. For more information on the translation of BQL to SQL, see [Translation of a BQL Command to SQL](#).

The following code shows an example of the `Product` DAC declaration in the fluent BQL style.

```
using System;
using PX.Data;

[Serializable]
public class Product : PX.Data.IBqlTable
{
    // The class used in BQL statements to refer to the ProductID column
    public abstract class productID : PX.Data.BQL.BqlInt.Field<productID>
    {
    }
    // The property holding the ProductID value in a record
    [PXDBIdentity(IsKey = true)]
    public virtual int? ProductID { get; set; }

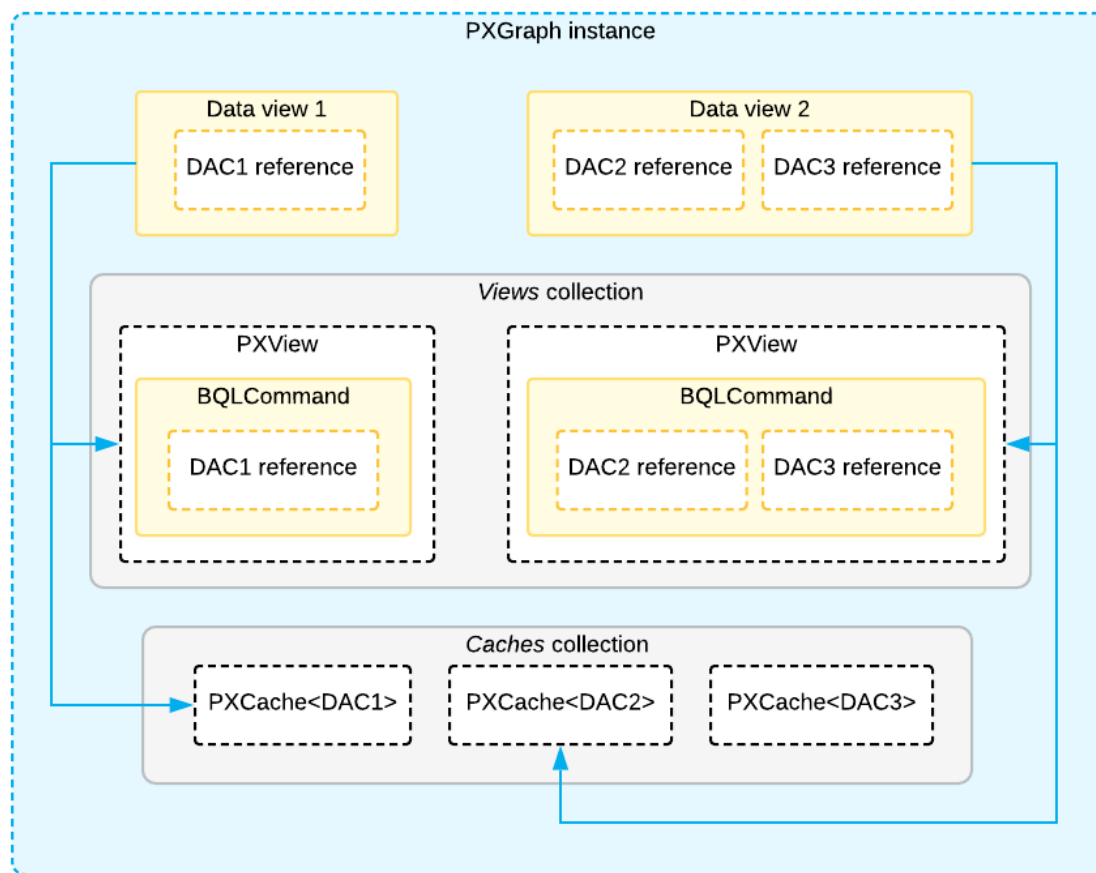
    // The class used in BQL statements to refer to the AvailQty column
    public abstract class availQty : PX.Data.BQL.BqlDecimal.Field<availQty>
    {
    }
    // The property holding the AvailQty value in a record
    [PXDBDecimal(2)]
    public virtual decimal? AvailQty { get; set; }
}
```

## PXView and PXCache of the Data View

You define a data view with a class derived from the `PXSelectBase` class, such as `SelectFrom<>.View` in fluent BQL and `PXSelect<>` in traditional BQL. When a graph executes a data view, the graph creates the following objects:

- The `PXView` object, which contains the BQL command that corresponds to the data view
- The `PXCache<Node>` objects whose type parameter is defined by the data access classes (DACs) that are used in the BQL command

The `PXView` object uses the BQL command to retrieve data from the database and stores the retrieved data in the `PXCache` object. The data view stores references to the corresponding `PXView` object and the `PXCache` object of the main DAC of the data view, as shown in the following diagram.



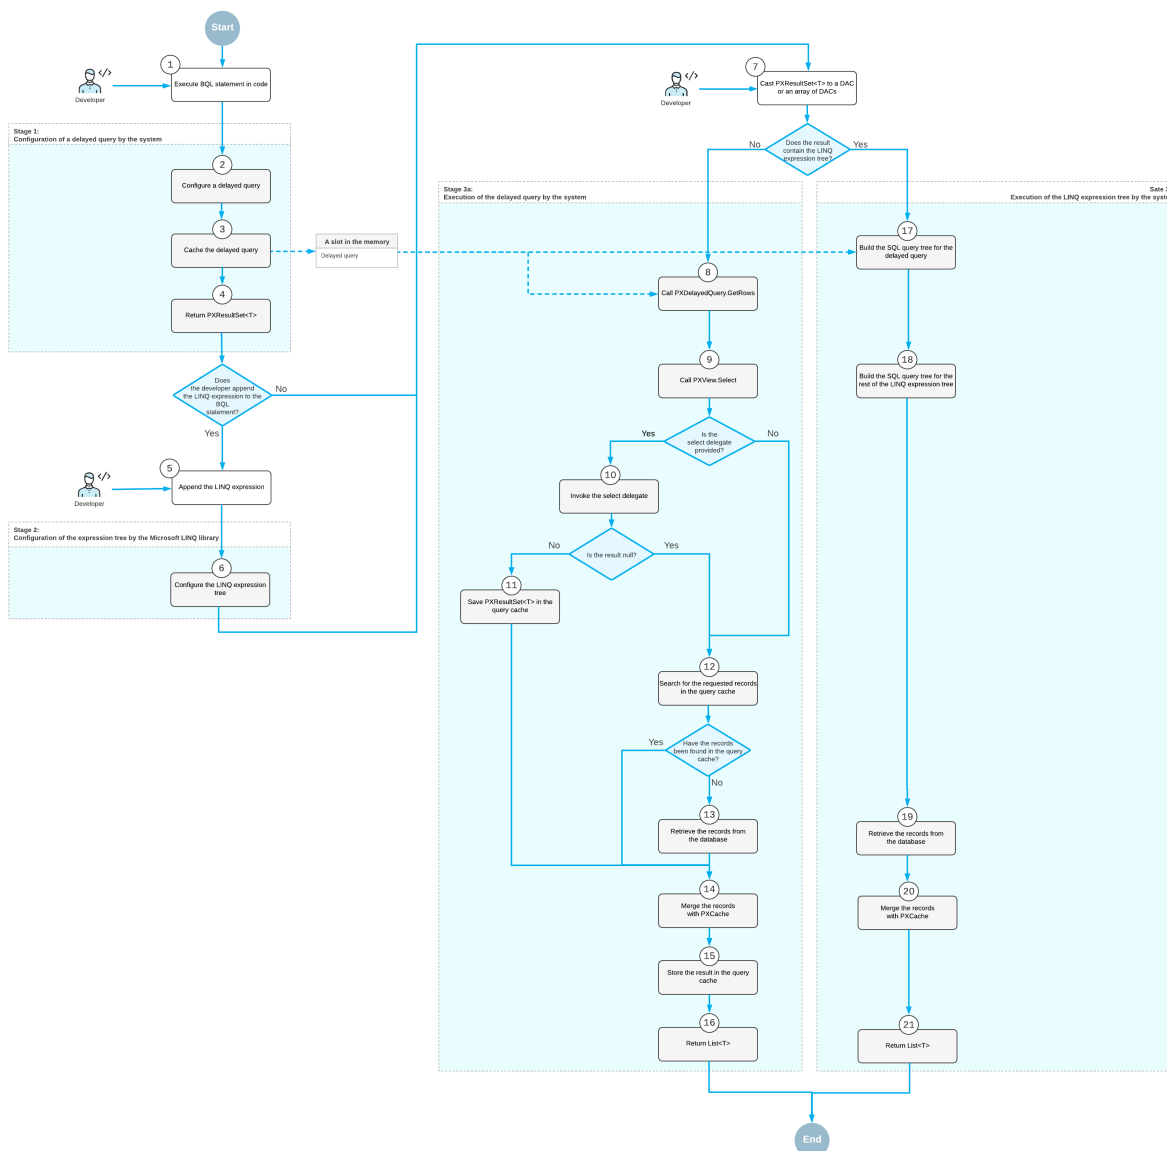
**Figure: Example of relationships between classes**

## Data Query Execution

The system executes a data query in the following stages, which are described in detail below:

- Stage 1: When a developer executes a BQL statement in code, Acumatica Framework configures a delayed query.
- Stage 2: If a language-integrated query (LINQ) statement is appended to the BQL statement, Microsoft LINQ configures the expression tree, which includes the delayed query.
- Stage 3: When the developer casts the result of the query to a data access class (DAC) or an array of DACs, the system does the following:
  - 3a: If the result of the query contains the expression tree created by LINQ, the system configures the SQL query tree that corresponds to the LINQ expression tree, and executes the SQL query tree.
  - 3b: If the result of the query is created only by BQL, the system configures the SQL query tree for the delayed query and executes this query tree.

The whole process is illustrated in the following diagram.



**Figure: Data query execution**

## Configuration of a Delayed Query

In code, you execute a business query language (BQL) statement in one of the following ways:

- You declare a data view (a `PXSelectBase`-derived class) as a member in a graph, and you specify this data view as the data member of the webpage control. The system uses this data view for basic data manipulation (inserting a data record, updating a data record, and deleting a data record) and executes the data view by calling the `Select()` method.
- You use the `static Select()` method of a `PXSelectBase`-derived class with a graph object as the parameter.
- You dynamically instantiate a `PXSelectBase`-derived class in code and execute it by using its `Select()` method. (You provide the graph object as a parameter to the class constructor.)
- You instantiate a class derived from the `BqlCommand` class (such as a `Select` class in traditional BQL or `FromSelect` in fluent BQL), create a `PXView` object that uses this `BqlCommand` class, create a graph object, and call one of the view's `Select()` methods.

When the `Select()` method is executed, Acumatica Framework does the following:

1. Configures a delayed query by creating a `PXDelayedQuery` instance. The `PXDelayedQuery` instance contains a reference to a `PXView` object, which contains references to `PXGraph` and the `BqlCommand` object to be executed.
2. Caches the delayed query by using the `PXContext.SetSlot` method. (For details on the slots, see [Use of Slots to Cache Data Objects](#).)
3. Returns a `PXResultset<T>` object whose type parameter is set to the DAC specified as the type parameter of the `SelectFrom` class (in fluent BQL) or as the first type parameter of the `PXSelect` class (in traditional BQL). This result set contains information about the delayed query.

You can iterate through the result set in a `foreach` loop, obtaining either DAC instances or `PXResult<>` instances. A `PXResult<>` instance represents a tuple of joined records from the result set. `PXResult<>` can be cast to any of the DAC types joined in the BQL statement. For more information on the use of the `PXResultset<T>` class, see [To Process the Result of the Execution of the BQL Statement](#).

### Configuration of the LINQ Expression Tree

Because the `PXResultset<T>` class implements the `IQueryable<T>` interface, developers can modify `PXResultset<T>` by using LINQ. If the developer appends LINQ statements to a result set, Microsoft LINQ incorporates the result set as an instance of the `SQLQueryable<T>` class in the LINQ expression tree. The resulting expression tree is an instance of the `SQLQueryable<T>` class, which contains references to an instance of `PXGraph`, Microsoft LINQ expression tree, the base `PXResultset<T>`, and an instance of `PX.Data.SQLTree.SQLQueryProvider`.

### Execution of the Delayed Query

Once you cast the result of the execution of the `Select()` method to a DAC or an array of DACs, or if you iterate through the DACs in the result by using the `foreach` statement, the system performs the following steps:

1. The system calls the `PXDelayedQuery.GetRows` method for the delayed query of the result set. This method internally calls the `PXView.Select()` method for the data view referred to in the delayed query.
2. If the select delegate is provided, inside the `PXView.Select()` method, the system invokes the select delegate by using the `PXView.InvokeDelegate` method and saves the result in the query cache of the graph. (The query cache stores the result set obtained by the execution of a specific BQL command.)
3. Inside the `PXView.Select()` method, the system searches for the requested records in the query cache by using the `PXView.LookupCache` method. If no records are found, the system requests data from the database by using the `PXView.GetResult` method. For details on the retrieval of records from the database, see [Translation of a BQL Command to SQL](#).
4. The system merges the records retrieved from the database or from the query cache with the modified records stored in `PXCache` by using the `PXView.MergeCache` method. For details about the merge, see [Merge of the Records with PXCache](#).
5. The system saves the result of the query in the query cache by using the `PXView.StoreCached` method.
6. The system returns the result as a `List<T>` type.

### Execution of the LINQ Expression Tree

Once you iterate over the LINQ expression tree, the system performs the following steps:

1. The system calls the `SQLQueryProvider.Execute()` method, which builds the `Remotion.Linq` expression tree based on the Microsoft LINQ expression tree and calls

`Remotion.Linq.QueryModel.Execute()` method with the `PX.Data.SQLTree.SQLInqExecutor` instance as a parameter.

2. The system builds the SQL query tree from the `Remotion.Linq.QueryModel` by calling the `SQLInqExecutor.ExecuteCollection<T>()` method. In this method, the system executes the `SQLInqQueryModelVisitor.VisitQueryModel()` method, which does the following:
  - a. Calls the `SQLInqQueryModelVisitor.VisitMainFromClause()` method, which builds the SQL query tree for the BQL statement that corresponds to the base `PXResultset<T>` of the query. This method internally calls the `BqlCommand.GetQueryInternal` method, which is described in [Translation of a BQL Command to an SQL Query Tree](#).
  - b. Builds the SQL query tree for the rest of the `Remotion.Linq` expression tree by calling the methods of `SQLInqQueryModelVisitor` for particular clauses and the columns included in the result of the query. If the system cannot build the SQL query tree for particular elements of the `Remotion.Linq` expression tree, the system falls back to the execution of the delayed query for the base BQL statement. For details about the fallback, see [Fallback to the LINQ to Objects Mode](#).
3. Within the `SQLInqExecutor.ExecuteCollection<T>()` method, the system uses the built SQL query tree to request the data from the database. For details about how the SQL query tree is translated to the SQL text that is passed to the database, see [Translation of the SQL Query Tree to SQL Text](#).
4. The system merges the records retrieved from the database with the modified records stored in `PXCache`. For details about the merge, see [Merge of the Records with PXCache](#).
5. The system returns the result as a `List<T>` type.

## Translation of a BQL Command to SQL

When the system executes a delayed query and calls the `PXView.GetResult` method to retrieve the data from the database, the system converts the business query language (BQL) command (`PX.Data.BqlCommand`) to the SQL query tree (`PX.Data.SQLTree.Query`), applies the needed restrictions on the SQL query tree (such as company and branch restrictions), and then converts the SQL query tree to the text of the SQL command for the target database type. This process is described in detail in the following sections.

### Translation of a BQL Command to an SQL Query Tree

To request the SQL query tree of the command, the system recursively calls the following methods:

1. `PXView.GetResult`
2. `PXGraph.ProviderSelect`
3. `PXDatabaseProviderBase.Select`
4. `BqlCommand.GetQuery`

The `BqlCommand.GetQuery` method calls the `BqlCommand.GetQueryInternal` method, which uses other methods as follows to prepare the SQL query tree:

1. If the BQL command contains aggregation, the `BqlCommand.AppendAggregatedFields` method appends to a new `Query` instance the SQL expressions (`PX.Data.SQL.SQLExpression`) that correspond to the fields that are surrounded with appropriate aggregation functions. If the BQL command does not contain aggregation, the `BqlCommand.AppendFields` method appends to a new `Query` instance the SQL expressions that correspond to the fields to be selected. The fields to be selected are the DAC fields that subscribe to the `OnCommandPreparing` event and are not restricted by `PXFieldScope`.
2. For each `Join` clause, the `IBqlJoin.AppendQuery` method adds to the `Query` instance the `Joiner` instance that corresponds to the `Join` clause.

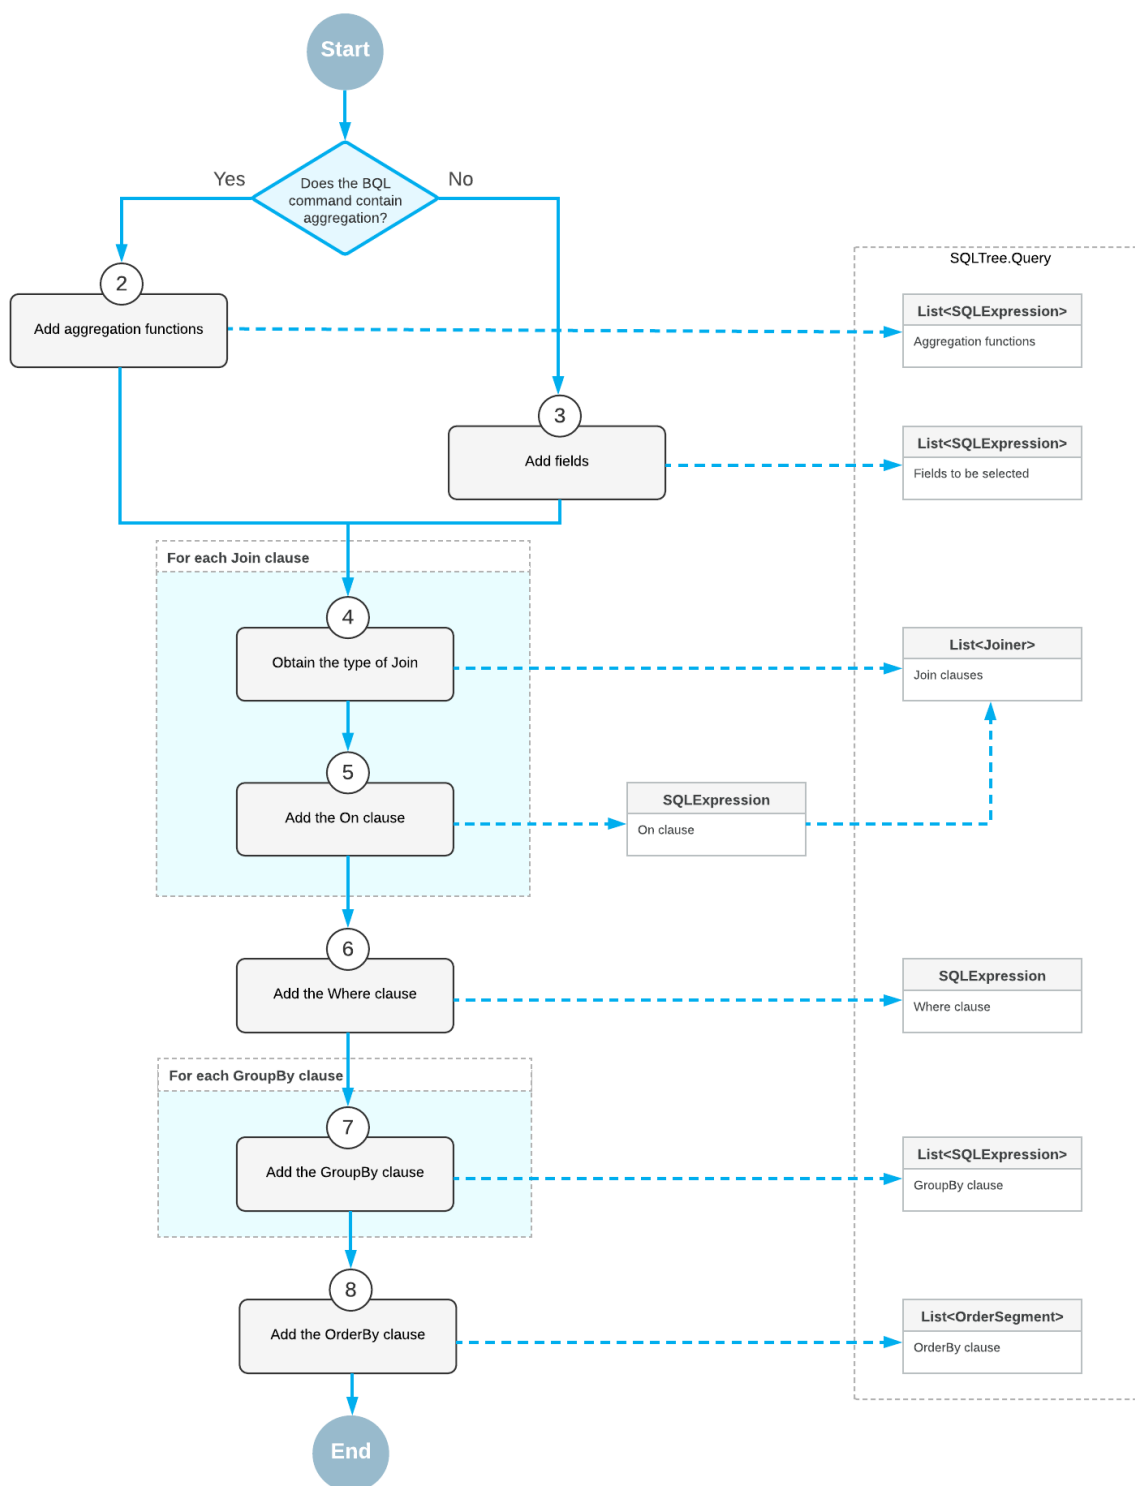


The `IBqlJoin.AppendQuery` method obtains the type of `Join` and, for all classes in the `On` clause that implement the `IBqlCreator` interface, successively executes the `IBqlCreator.AppendExpression` method, starting from the `On` class and then proceeding with enclosed classes, such as the `Where` classes and comparison classes. For the DAC fields (`IBqlField`-derived classes), the `BqlCommand.GetSingleExpression` method obtains the SQL expression.

3. For all classes in the `Where` and `GroupBy` clauses that implement the `IBqlCreator` interface, the system successively executes the `IBqlCreator.AppendExpression` method, which appends to the `Query` instance the SQL expression that corresponds to the classes. For the DAC fields (`IBqlField`-derived classes), the `BqlCommand.GetSingleExpression` method obtains the SQL expression.
4. The `IBqlOrderBy.AppendQuery` method adds to the `Query` instance the list of `OrderSegment` instances that corresponds to the `OrderBy` clause.

For each sorting class (a `IBqlSortColumn`-derived class) in the `OrderBy` clause, the `IBqlSortColumn.AppendQuery` method adds to the `Query` instance the `OrderSegment` instance that corresponds to the sorting column. For the DAC fields (`IBqlField`-derived classes), the `BqlCommand.GetSingleExpression` method obtains the SQL expression. If the original BQL statement does not specify ordering, the system adds to the `Query` instance sorting by the DAC key fields (in ascending order).

The following diagram shows the conversion of a BQL command to an SQL query tree.



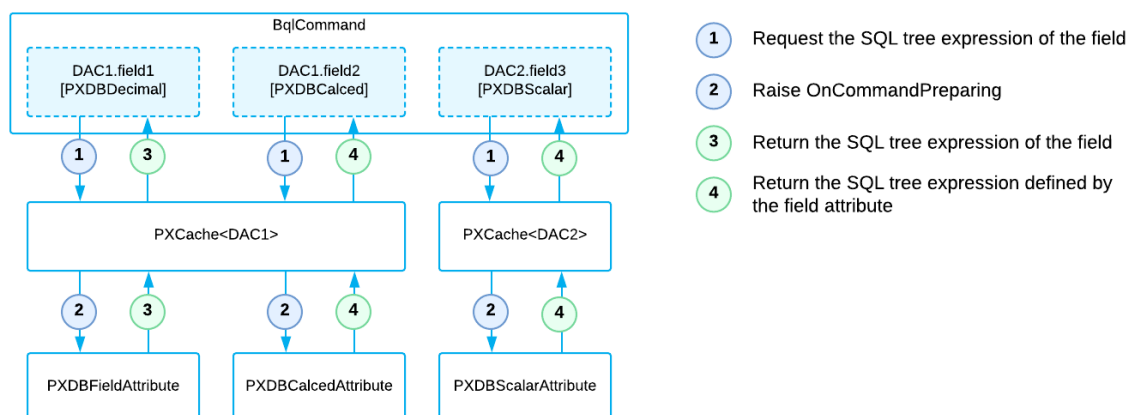
**Figure: Conversion of a BQL command to an SQL query tree**

### SQL Tree Expression of a Field

To obtain the SQL tree expression of each field of the BQL command, the `BqlCommand` instance creates a `PXCache` instance that corresponds to the data access class (DAC) to which the field belongs. The `PXCache` instance generates the `OnCommandPreparing` event with the specified `PXDBOperation` type

(which specifies the type of the database operation). The attribute assigned to the DAC field (that is, the attribute that implements the `IPXCommandPreparingSubscriber` interface) handles the event and returns the `PX.Data.SQLTree.SQLExpression` instance that corresponds to the field or to the BQL statement that is defined by the field attribute. (For example, the `PXDBCalced` and `PXDBScalar` attributes define BQL statements.)

The following diagram shows how `BqlCommand` obtains the SQL tree expression for the fields of a BQL command.



**Figure: Retrieval of the SQL tree expression for the fields of a BQL command**

### Translation of a BQL Command with Parameters to an SQL Query Tree

Before the system requests the SQL query tree of a BQL command, the `PXView` object retrieves the values of the parameters used in the query as follows:

- For a field with `FromCurrent` appended (in fluent BQL) or specified in the `Current` parameter (in traditional BQL), the `PXView` object retrieves the field value from the `Current` object of the `PXCACHE` object. If the current field value is null, the `PXView` object triggers the `FieldDefaulting` event handlers and retrieves the default value from the `PXDefault` attribute value (if any).



: The default value is not retrieved if `FromCurrent.NoDefault` is appended to the field (in fluent BQL) or if the `Current2` parameter is used (in traditional BQL).

- For a field with `AsOptional` appended (in fluent BQL) or specified in the `Optional` parameter, if the explicit field value is specified, the `PXView` object triggers the `FieldUpdating` event, whose handlers can transform the external presentation of the field value to an internal value (for example, transform `ProductCD` to `ProductID`). If the field value is not specified, the `PXView` object retrieves the field value from the `Current` object of the `PXCACHE` object. If the current field value is null, the `PXView` object triggers the `FieldDefaulting` event and retrieves the default value from the `PXDefault` attribute value (if any).



: The default value is not retrieved if `AsOptional.NoDefault` is appended to the field (in fluent BQL) or if the `Optional2` parameter is used.

When `BqlCommand` creates the SQL query tree that corresponds to the BQL command, `IBqlCreator.AppendExpression` (which is implemented by the `ParameterBase<Field>` class) includes the parameters in the SQL query tree. After `BqlCommand` has created the SQL query tree that corresponds to the BQL command, the system inserts into the SQL query tree the actual values of the parameters retrieved by `PXView`.

### Translation of the SQL Query Tree to SQL Text

A `PXGraph` instance stores information about the target database type in its `SqlDialect` property. `SQLTree.Query` has the `Connection` property, which is responsible for the conversion of the SQL query tree to the SQL text in the format of the target database. To convert the SQL query tree to text, the system does the following:

1. Calls the `SQLDialect.GetConnection` method of the graph instance to retrieve the needed `SQLTree.Connection`.
2. Passes this `Connection` instance to the `SQLTree.Query.SQLQuery` method, which converts the SQL query tree to the text for Microsoft SQL or MySQL, depending on the passed `Connection`.

### Merge of the Records with `PXCache`

For the queries defined with business query language (BQL) or with BQL and LINQ, the system merges the records retrieved from the database with the modified records stored in `PXCache` as follows:

1. If the query is read-only, the result set is not merged with any `PXCache` object. The system returns the data records as they are currently stored in the database.



: A query is read-only if the `IsReadOnly` property of the underlying `PXView` object is `true`. For example, the traditional BQL statements that use aggregation or are based on one of the `PXSelectReadOnly` classes are read-only. The fluent BQL statements that have `.ReadOnly` appended are read-only.

2. If the query is not read-only and contains filtering by data access class (DAC) fields by using LINQ (that is, only the values in specific columns of the database tables are returned in the results of the query), the system checks whether the primary key fields of the primary DAC (the first DAC in the BQL statement) are returned in the results of the query. If the primary key fields are not in the results, no merge with any `PXCache` object is performed. If all primary key fields are in the results, the result set is merged with the `PXCache` object that corresponds to the primary DAC.
3. If the query is not read-only, does not contain filtering of DAC fields by using LINQ, and does not contain joins, the result set is merged with the contents of the appropriate `PXCache` object, and the system returns the result set updated with the modifications stored in `PXCache`.
4. If the query is not read-only, does not contain filtering by DAC fields by using LINQ, and joins data from multiple tables, the result set is merged with only the `PXCache` object that corresponds to the first table of the BQL statement. The `PXResultset<>` object, which represents the result set, contains objects of the generic `PXResult<>` type. This type can be cast to the data access classes (DACs) that represent the joined tables. The instance of the primary DAC to which `PXResult<>` is cast contains the records from the database that are updated with the modifications stored in `PXCache`. Casting `PXResult<>` to a joined DAC returns the instance that contains values from the database and has no relation with the `PXCache` instances of the corresponding DAC types.

The following diagram illustrates the database records being merged with `PXCache`.

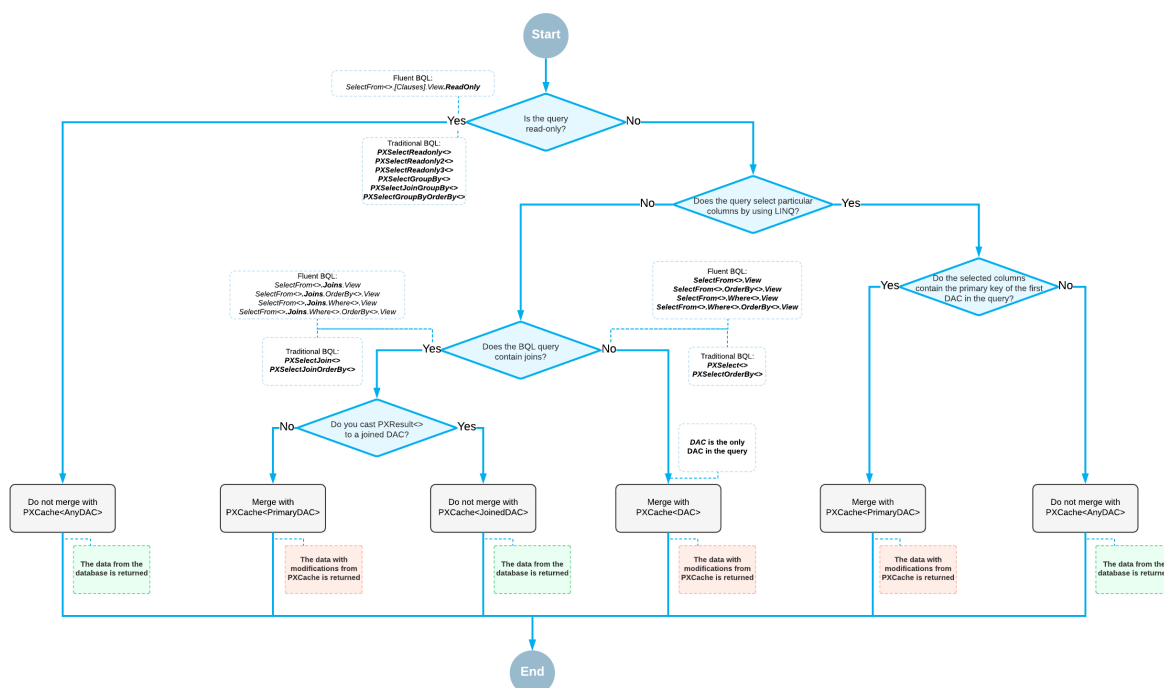


Figure: Merge with PXCache

## Comparison of Fluent BQL, Traditional BQL, and LINQ

In this topic, you can learn the main differences between the queries defined with fluent business query language (BQL), traditional BQL, and language-integrated query (LINQ).

### Comparison of Fluent BQL, Traditional BQL, and LINQ

Characteristic	Fluent BQL	Traditional BQL	LINQ
The queries can be used to define data views in graphs.	Yes	Yes	No
The queries can be defined in code.	Yes	Yes	Yes
The queries can be defined in DAC field attributes.	Yes	Yes	No
DACs are used to define database tables in the queries.	Yes	Yes	Yes
The queries can be used for dynamic query building.	Yes	Yes	Yes
The queries can be parsed and modified by the direct use of reflection—that is, by <code>Type.GetGenericArguments()</code> .	No	Yes	No
Clauses (such as <code>Join</code> , <code>Where</code> , <code>Aggregate</code> , <code>OrderBy</code> , and <code>On</code> ) can be used separately of the query.	No	Yes	No
The query language includes numbered classes (such as <code>Select2</code> and <code>Select6</code> ).	No	Yes	No

Characteristic	Fluent BQL	Traditional BQL	LINQ
Each subsequent element of the query is passed as a generic parameter of the previous one.	No	Yes	No
To create a query, a developer needs to select a suitable command overload.	No	Yes	No
IntelliSense can offer continuations that are relevant for the current query state.	Yes	No	Yes
The queries use strongly typed expressions, which makes compile-time type checks possible.	Yes	No	Yes
The queries can contain explicit brackets in conditions.	Yes	No; the <code>Where</code> clause can be used instead	Yes
You can specify particular columns of the tables to be selected.	Yes; you have to use <code>PXFieldScope</code>	Yes; you have to use <code>PXFieldScope</code>	Yes
The query is not executed until it is iterated over.	Yes	Yes	Yes

## Fluent BQL and Traditional BQL Equivalents

The fluent business query language (BQL) library defines the equivalents of traditional BQL classes listed in the following tables.

You can use the converter from BQL to fluent BQL available on GitHub (<https://github.com/Acumatica/>) to convert the previously defined BQL queries to fluent BQL.

### Data View Declarations



: All data views that contain aggregating are read-only.

Fluent BQL	Traditional BQL
<code>SelectFrom&lt;&gt;.View</code>	<code>PXSelect&lt;&gt;</code>
<code>SelectFrom&lt;&gt;.View.ReadOnly</code>	<code>PXSelectReadOnly&lt;&gt;</code>
<code>SelectFrom&lt;&gt;.OrderBy&lt;&gt;.View</code>	<code>PXSelectOrderBy&lt;,&gt;</code>
<code>SelectFrom&lt;&gt;.OrderBy&lt;&gt;.View.ReadOnly</code>	<code>PXSelectReadOnly3&lt;,&gt;</code>
<code>SelectFrom&lt;&gt;.AggregateTo&lt;&gt;.View.ReadOnly</code>	<code>PXSelectGroupBy&lt;,&gt;</code>
<code>SelectFrom&lt;&gt;.AggregateTo&lt;&gt;.OrderBy&lt;&gt;.View.ReadOnly</code>	<code>PXSelectGroupByOrderBy&lt;,,&gt;</code>
<code>SelectFrom&lt;&gt;.Where&lt;&gt;.View</code>	<code>PXSelect&lt;,&gt;</code>
<code>SelectFrom&lt;&gt;.Where&lt;&gt;.View.ReadOnly</code>	<code>PXSelectReadOnly&lt;,&gt;</code>
<code>SelectFrom&lt;&gt;.Where&lt;&gt;.OrderBy&lt;&gt;.View</code>	<code>PXSelect&lt;,,&gt;</code>
<code>SelectFrom&lt;&gt;.Where&lt;&gt;.OrderBy&lt;&gt;.View.ReadOnly</code>	<code>PXSelectReadOnly&lt;,,&gt;</code>
<code>SelectFrom&lt;&gt;.Where&lt;&gt;.AggregateTo&lt;&gt;.View.ReadOnly</code>	<code>PXSelectGroupBy&lt;,,&gt;</code>
<code>SelectFrom&lt;&gt;.Where&lt;&gt;.AggregateTo&lt;&gt;.OrderBy&lt;&gt;.View.ReadOnly</code>	<code>PXSelectGroupBy&lt;,,,&gt;</code>

Fluent BQL	Traditional BQL
SelectFrom<>. [Joins].View	PXSelectJoin<,>
SelectFrom<>. [Joins].View.ReadOnly	PXSelectReadOnly2<,>
SelectFrom<>. [Joins].OrderBy<>.View	PXSelectJoinOrderBy<,,>
SelectFrom<>. [Joins].OrderBy<>.View.ReadOnly	PXSelectReadOnly3<,,>
SelectFrom<>. [Joins].AggregateTo<>.View.ReadOnly	PXSelectJoinGroupBy<,,>
SelectFrom<>. [Joins].AggregateTo<>.OrderBy<>.View.ReadOnly	PXSelectGroupByOrderBy<,,,>
SelectFrom<>. [Joins].Where<>.View	PXSelectJoin<,,>
SelectFrom<>. [Joins].Where<>.View.ReadOnly	PXSelectReadOnly2<,,>
SelectFrom<>. [Joins].Where<>.OrderBy<>.View	PXSelectJoin<,,,>
SelectFrom<>. [Joins].Where<>.OrderBy<>.View.ReadOnly	PXSelectReadOnly2<,,,>
SelectFrom<>. [Joins].Where<>.AggregateTo<>.View.ReadOnly	PXSelectJoinGroupBy<,,,>
SelectFrom<>. [Joins].Where<>.AggregateTo<>.OrderBy<>.View.ReadOnly	PXSelectJoinGroupBy<,,,,>

### Select Commands

Fluent BQL	Traditional BQL
SelectFrom<>	Select<>
SelectFrom<>.OrderBy<>	Select3<,>
SelectFrom<>.AggregateTo<>	Select4<,>
SelectFrom<>.AggregateTo<>.OrderBy<>	Select6<,,>
SelectFrom<>.Where<>	Select<,>
SelectFrom<>.Where<>.OrderBy<>	Select<,,>
SelectFrom<>.Where<>.AggregateTo<>	Select4<,,>
SelectFrom<>.Where<>.AggregateTo<>.OrderBy<>	Select4<,,,>
SelectFrom<>. [Joins]	Select2<,>
SelectFrom<>. [Joins].OrderBy<>	Select3<,,>
SelectFrom<>. [Joins].AggregateTo<>	Select5<,,>
SelectFrom<>. [Joins].AggregateTo<>.OrderBy<>	Select6<,,,>
SelectFrom<>. [Joins].Where<>	Select2<,,>
SelectFrom<>. [Joins].Where<>.OrderBy<>	Select2<,,,>
SelectFrom<>. [Joins].Where<>.AggregateTo<>	Select5<,,,>
SelectFrom<>. [Joins].Where<>.AggregateTo<>.OrderBy<>	Select5<,,,,>

## Search Commands

Fluent BQL	Traditional BQL
SelectFrom<>.SearchFor<>	Search<>
SelectFrom<>.OrderBy<>.SearchFor<>	Search3<,>
SelectFrom<>.AggregateTo<>.SearchFor<>	Search4<,>
SelectFrom<>.AggregateTo<>.OrderBy<>.SearchFor<>	Search6<,,>
SelectFrom<>.Where<>.SearchFor<>	Search<,>
SelectFrom<>.Where<>.OrderBy<>.SearchFor<>	Search<,,>
SelectFrom<>.Where<>.AggregateTo<>.SearchFor<>	Search4<,,>
SelectFrom<>.Where<>.AggregateTo<>.OrderBy<>.SearchFor<>	Search4<,,,>
SelectFrom<>.[Joins].SearchFor<>	Search2<,>
SelectFrom<>.[Joins].OrderBy<>.SearchFor<>	Search3<,,>
SelectFrom<>.[Joins].AggregateTo<>.SearchFor<>	Search5<,,>
SelectFrom<>.[Joins].AggregateTo<>.OrderBy<>.SearchFor<>	Search6<,,,>
SelectFrom<>.[Joins].Where<>.SearchFor<>	Search2<,,>
SelectFrom<>.[Joins].Where<>.OrderBy<>.SearchFor<>	Search2<,,,>
SelectFrom<>.[Joins].Where<>.AggregateTo<>.SearchFor<>	Search5<,,,>
SelectFrom<>.[Joins].Where<>.AggregateTo<>.OrderBy<>.SearchFor<>	Search5<,,,,>

## Join Clauses

Fluent BQL	Traditional BQL
.InnerJoin<Table>.On<>	InnerJoin<Table,On>
.InnerJoin<Table>.On<>.NextJoin	InnerJoin<Table,On,NextJoin>
.InnerJoin<Table>.On<>.SingleTableOnly	InnerJoinSingleTable<Table,On>
.InnerJoin<Table>.On<>.SingleTableOnly.NextJoin	InnerJoinSingleTable<Table,On,NextJoin>
.LeftJoin<Table>.On<>	LeftJoin<Table,On>
.LeftJoin<Table>.On<>.NextJoin	LeftJoin<Table,On,NextJoin>
.LeftJoin<Table>.On<>.SingleTableOnly	LeftJoinSingleTable<Table,On>
.LeftJoin<Table>.On<>.SingleTableOnly.NextJoin	LeftJoinSingleTable<Table,On,NextJoin>
.RightJoin<Table>.On<>	RightJoin<Table,On>
.RightJoin<Table>.On<>.NextJoin	RightJoin<Table,On,NextJoin>
.RightJoin<Table>.On<>.SingleTableOnly	RightJoinSingleTable<Table,On>
.RightJoin<Table>.On<>.SingleTableOnly.NextJoin	RightJoinSingleTable<Table,On,NextJoin>
.FullJoin<Table>.On<>	FullJoin<Table,On>
.FullJoin<Table>.On<>.NextJoin	FullJoin<Table,On,NextJoin>



Fluent BQL	Traditional BQL
<code>.FullJoin&lt;Table&gt;.On&lt;&gt;.SingleTableOnly</code>	<code>FullJoinSingleTable&lt;Table, On&gt;</code>
<code>.FullJoin&lt;Table&gt;.On&lt;&gt;.SingleTableOnly.NextJoin</code>	<code>FullJoinSingleTable&lt;Table, On, NextJoin&gt;</code>
<code>.CrossJoin&lt;Table&gt;</code>	<code>CrossJoin&lt;Table&gt;</code>
<code>.CrossJoin&lt;Table&gt;.NextJoin</code>	<code>CrossJoin&lt;Table, NextJoin&gt;</code>
<code>.CrossJoin&lt;Table&gt;.SingleTableOnly</code>	<code>CrossJoinSingleTable&lt;Table&gt;</code>
<code>.CrossJoin&lt;Table&gt;.SingleTableOnly.NextJoin</code>	<code>CrossJoinSingleTable&lt;Table, NextJoin&gt;</code>

### Where Clause

Fluent BQL	Traditional BQL
<code>.Where&lt;UnaryOperator&gt;</code>	<code>Where&lt;UnaryOperator&gt;</code>
<code>.Where&lt;Operand.Comparison&gt;</code>	<code>Where&lt;Operand, Comparison&gt;</code>
<code>.Where&lt;Operand.Comparison.NextOperator&gt;</code>	<code>Where&lt;Operand, Comparison, NextOperator&gt;</code>
<code>.Where&lt;UnaryOperator.NextOperator&gt;</code>	<code>Where2&lt;UnaryOperator, NextOperator&gt;</code>

### Aggregate Clause

Fluent BQL	Traditional BQL
<code>.AggregateTo&lt;Function&gt;</code>	<code>Aggregate&lt;Function&gt;</code>
<code>.AggregateTo&lt;TFunctions&gt;.THaving</code>	<code>Aggregate&lt;TFunctions, THaving&gt;</code>
<code>.Having&lt;TCondition&gt;</code>	<code>Having&lt;TCondition&gt;</code>
<code>GroupBy&lt;Field&gt;</code>	<code>GroupBy&lt;Field&gt;</code>
<code>GroupBy&lt;Field&gt;, NextAggregate</code>	<code>GroupBy&lt;Field, NextAggregate&gt;</code>
<code>Max&lt;Field&gt;</code>	<code>Max&lt;Field&gt;</code>
<code>Max&lt;Field&gt;, NextAggregate</code>	<code>Max&lt;Field, NextAggregate&gt;</code>
<code>Min&lt;Field&gt;</code>	<code>Min&lt;Field&gt;</code>
<code>Min&lt;Field&gt;, NextAggregate</code>	<code>Min&lt;Field, NextAggregate&gt;</code>
<code>Sum&lt;Field&gt;</code>	<code>Sum&lt;Field&gt;</code>
<code>Sum&lt;Field&gt;, NextAggregate</code>	<code>Sum&lt;Field, NextAggregate&gt;</code>
<code>Avg&lt;Field&gt;</code>	<code>Avg&lt;Field&gt;</code>
<code>Avg&lt;Field&gt;, NextAggregate</code>	<code>Avg&lt;Field, NextAggregate&gt;</code>
<code>Count</code>	<code>Count</code>
<code>Count&lt;Field&gt;</code>	<code>Count&lt;Field&gt;</code>

### OrderBy Clause

Fluent BQL	Traditional BQL
<code>.OrderBy&lt;List&gt;</code>	<code>OrderBy&lt;List&gt;</code>

Fluent BQL	Traditional BQL
Field.Asc	Asc<Field>
Field.Asc, NextSort	Asc<Field,NextSort>
Field.Desc	Desc<Field>
Field.Desc, NextSort	Desc<Field,NextSort>

### Parameters

Fluent BQL	Traditional BQL
Field.FromCurrent	Current<Field>
Field.FromCurrent.NoDefault	Current2<Field>
Field.AsOptional	Optional<Field>
Field.AsOptional.NoDefault	Optional2<Field>
@P.AsBool	Required<Field>, where the property field of Field has the bool type
@P.AsByte	Required<Field>, where the property field of Field has the byte type
@P.AsShort	Required<Field>, where the property field of Field has the short type
@P.AsInt	Required<Field>, where the property field of Field has the int type
@P.AsLong	Required<Field>, where the property field of Field has the long type
@P.AsFloat	Required<Field>, where the property field of Field has the float type
@P.AsDouble	Required<Field>, where the property field of Field has the double type
@P.AsDecimal	Required<Field>, where the property field of Field has the decimal type
@P.AsGuid	Required<Field>, where the property field of Field has the Guid type
@P.AsDateTime	Required<Field>, where the property field of Field has the DateTime type
@P.AsString	Required<Field>, where the property field of Field has the string type
Argument.AsBool	Argument<bool?>
Argument.AsByte	Argument<byte?>
Argument.AsShort	Argument<short?>
Argument.AsInt	Argument<int?>
Argument.AsLong	Argument<long?>
Argument.AsFloat	Argument<float?>

Fluent BQL	Traditional BQL
<code>Argument.AsDouble</code>	<code>Argument&lt;double?&gt;</code>
<code>Argument.AsDecimal</code>	<code>Argument&lt;decimal?&gt;</code>
<code>Argument.AsGuid</code>	<code>Argument&lt;Guid?&gt;</code>
<code>Argument.AsDateTime</code>	<code>Argument&lt;DateTime?&gt;</code>
<code>Argument.AsString</code>	<code>Argument&lt;string&gt;</code>

### Logical Operators and Brackets

Fluent BQL	Traditional BQL
<code>And&lt;UnaryOperator&gt;</code>	<code>And&lt;UnaryOperator&gt;</code>
<code>And&lt;Operand.Comparison&gt;</code>	<code>And&lt;Operand, Comparison&gt;</code>
<code>And&lt;Operand.Comparison&gt;.NextOperator</code>	<code>And&lt;Operand, Comparison, NextOperator&gt;</code>
<code>And&lt;UnaryOperator.NextOperator&gt;</code>	<code>And2&lt;UnaryOperator, NextOperator&gt;</code>
<code>Or&lt;UnaryOperator&gt;</code>	<code>Or&lt;UnaryOperator&gt;</code>
<code>Or&lt;Operand.Comparison&gt;</code>	<code>Or&lt;Operand, Comparison&gt;</code>
<code>Or&lt;Operand.Comparison&gt;.NextOperator</code>	<code>Or&lt;Operand, Comparison, NextOperator&gt;</code>
<code>Or&lt;UnaryOperator.NextOperator&gt;</code>	<code>Or2&lt;UnaryOperator, NextOperator&gt;</code>
<code>Not&lt;UnaryOperator&gt;</code>	<code>Not&lt;UnaryOperator&gt;</code>
<code>Not&lt;Operand.Comparison&gt;</code>	<code>Not&lt;Operand, Comparison&gt;</code>
<code>Not&lt;Operand.Comparison.NextOperator&gt;</code>	<code>Not&lt;Operand, Comparison, NextOperator&gt;</code>
<code>Not&lt;UnaryOperator.NextOperator&gt;</code>	<code>Not2&lt;UnaryOperator, NextOperator&gt;</code>
<code>Brackets&lt;UnaryOperator&gt;</code>	<code>Where&lt;UnaryOperator&gt;</code>
<code>Brackets&lt;Operand.Comparison&gt;</code>	<code>Where&lt;Operand, Comparison&gt;</code>
<code>Brackets&lt;Operand.Comparison.NextOperator&gt;</code>	<code>Where&lt;Operand, Comparison, NextOperator&gt;</code>
<code>Brackets&lt;UnaryOperator.NextOperator&gt;</code>	<code>Where2&lt;UnaryOperator, NextOperator&gt;</code>

### Comparisons

Fluent BQL	Traditional BQL
<code>Table.field.IsEqual&lt;TOperand&gt;</code>	<code>&lt;Table.field, Equal&lt;TOperand&gt;&gt;</code>
<code>Table.field.IsNotEqual&lt;TOperand&gt;</code>	<code>&lt;Table.field, NotEqual&lt;TOperand&gt;&gt;</code>
<code>Table.field.IsGreaterEqual&lt;TOperand&gt;</code>	<code>&lt;Table.field, GreaterEqual&lt;TOperand&gt;&gt;</code>
<code>Table.field.IsGreater&lt;TOperand&gt;</code>	<code>&lt;Table.field, Greater&lt;TOperand&gt;&gt;</code>
<code>Table.field.IsLessEqual&lt;TOperand&gt;</code>	<code>&lt;Table.field, LessEqual&lt;TOperand&gt;&gt;</code>
<code>Table.field.IsLess&lt;TOperand&gt;</code>	<code>&lt;Table.field, Less&lt;TOperand&gt;&gt;</code>
<code>Table.field.IsLike&lt;TOperand&gt;</code>	<code>&lt;Table.field, Like&lt;TOperand&gt;&gt;</code>
<code>Table.field.IsNotLike&lt;TOperand&gt;</code>	<code>&lt;Table.field, NotLike&lt;TOperand&gt;&gt;</code>

Fluent BQL	Traditional BQL
Table.field.IsBetween<TOperand>	<Table.field, Between<TOperand>>
Table.field.IsNotBetween<TOperand>	<Table.field, NotBetween<TOperand>>
Table.field.IsNull	<Table.field, IsNull<TOperand>>
Table.field.IsNotNull	<Table.field, IsNotNull<TOperand>>
Table.field.IsIn<TOperand>	<Table.field, In<TOperand>>
Table.field.IsNotIn<TOperand>	<Table.field, NotIn<TOperand>>
Table.field.IsInSub<TSearch>	<Table.field, In2<TSearch>>
Table.field.IsNotInSub<TSearch>	<Table.field, NotIn2<TSearch>>
Table.field.IsIn<TConst1, ..., TConstN>	<Table.field, In3<TConst1, ..., TConstN>>
Table.field.IsNotIn<TConst1, ..., TConstN>	<Table.field, NotIn3<TConst1, ..., TConstN>>

### Case, When, Then, and Else Operators

Fluent BQL	Traditional BQL
Operand1.When<Condition1>.Else<Operand2>.When<Condition2>.[...]	Switch<Cases>.
Operand1.When<Condition1>.[...].Else<Default>	Switch<Cases,Default>
Operand.When<Condition>	Case<Condition,Operand>
Operand.When<Condition>.Else<Operand2>.When<Condition2>.[...]	Case<Condition,Operand,NextCase>

### Arithmetic Operations and Operations with Strings and Dates

Fluent BQL	Traditional BQL
Op1.Add<Op2>	Add<Op1, Op2>
Op1.Subtract<Op2>	Sub<Op1, Op2>
Op1.Multiply<Op2>	Mult<Op1, Op2>
Op1.Divide<Op2>	Div<Op1, Op2>
Op1.Concat<Op2>	Concat<Op1, Op2>
Op1.IfNullThen<Op2>	IsNull<Op1, Op2>
Op1.NullIf<Op2>	NullIf<Op1, Op2>
Date1.Diff<Date2>.Years	DateDiff<Date1, Date2, DateDiff.year>
Date1.Diff<Date2>.Quarters	DateDiff<Date1, Date2, DateDiff.quarter>
Date1.Diff<Date2>.Months	DateDiff<Date1, Date2, DateDiff.month>
Date1.Diff<Date2>.Weeks	DateDiff<Date1, Date2, DateDiff.week>
Date1.Diff<Date2>.Days	DateDiff<Date1, Date2, DateDiff.day>

Fluent BQL	Traditional BQL
<code>Date1.Diff&lt;Date2&gt;.Hours</code>	<code>DateDiff&lt;Date1, Date2, DateDiff.hour&gt;</code>
<code>Date1.Diff&lt;Date2&gt;.Minutes</code>	<code>DateDiff&lt;Date1, Date2, DateDiff.minute&gt;</code>
<code>Date1.Diff&lt;Date2&gt;.Seconds</code>	<code>DateDiff&lt;Date1, Date2, DateDiff.second&gt;</code>
<code>Date1.Diff&lt;Date2&gt;.Milliseconds</code>	<code>DateDiff&lt;Date1, Date2, DateDiff.millisecond&gt;</code>
<code>DatePart&lt;Date&gt;.Year</code>	<code>DatePart&lt;DatePart.year, Date&gt;</code>
<code>DatePart&lt;Date&gt;.Quarter</code>	<code>DatePart&lt;DatePart.quarter, Date&gt;</code>
<code>DatePart&lt;Date&gt;.Month</code>	<code>DatePart&lt;DatePart.month, Date&gt;</code>
<code>DatePart&lt;Date&gt;.Week</code>	<code>DatePart&lt;DatePart.week, Date&gt;</code>
<code>DatePart&lt;Date&gt;.WeekDay</code>	<code>DatePart&lt;DatePart.weekDay, Date&gt;</code>
<code>DatePart&lt;Date&gt;.Day</code>	<code>DatePart&lt;DatePart.day, Date&gt;</code>
<code>DatePart&lt;Date&gt;.DayOfYear</code>	<code>DatePart&lt;DatePart.dayOfYear, Date&gt;</code>
<code>DatePart&lt;Date&gt;.Hour</code>	<code>DatePart&lt;DatePart.hour, Date&gt;</code>
<code>DatePart&lt;Date&gt;.Minute</code>	<code>DatePart&lt;DatePart.minute, Date&gt;</code>
<code>DatePart&lt;Date&gt;.Second</code>	<code>DatePart&lt;DatePart.second, Date&gt;</code>

## To Execute BQL Statements

To send a request to the database, you call the `Select()` method of a `PXSelectBase`-derived class and cast the result of the query execution to a data access class (DAC) or an array of DACs, as described in this topic. The `Select()` method can accept additional parameters if a business query language (BQL) statement includes parameters.

### To Execute a BQL Statement That Defines a Data View

When a webpage requests data, you do not need to execute a data view manually; the system executes each data view automatically. If you need to manually execute a BQL statement that defines a data view, do the following:

1. Declare a data view as a member in a graph.
2. Execute the data view by calling the `Select()` method of a `PXSelectBase`-derived class.
3. Cast the result of the query execution to a DAC or an array of DACs, or iterate through DACs in the result by using the `foreach` statement. The following sample code shows the approach of iterating through DACs. For details, see [To Process the Result of the Execution of the BQL Statement](#).

```
// Business logic controller (graph) declaration
public class OrderDataEntry : PXGraph<OrderDataEntry, SalesOrder>
{
    // A data view declared as a graph member
    public SelectFrom<SalesOrder>.
        OrderBy<Asc<SalesOrder.orderNbr>>.View Orders;
    ...
    public void SomeMethod()
    {
        // An execution of the data view in code
        foreach(SalesOrder so in Orders.Select())
        {
            // The SalesOrder record selected by a data view can
```

```

        // be modified and updated through the Update() method.
        so.OrderTotal = so.LinesTotal + so.FreightAmt;
        // Update the SalesOrder data record in PXCache
        Orders.Update(so);
    }
}

```

### To Execute a BQL Statement Staticly

To execute a BQL statement statically, do the following:

1. Execute a BQL statement by using the `static Select()` method of a `PXSelectBase`-derived class. Provide a graph object as the parameter of the method.
2. Cast the result of the query execution to a DAC or an array of DACs, or iterate through DACs in the result by using the `foreach` statement. The following sample code shows the approach of iterating through DACs. For details, see [To Process the Result of the Execution of the BQL Statement](#).

```

// Business logic controller (graph) declaration
public class OrderDataEntry : PXGraph<OrderDataEntry, SalesOrder>
{
    ...
    public void SomeMethod()
    {
        // Execution through the static Select() method
        foreach(SalesOrder so in
            SelectFrom<SalesOrder>.
                OrderBy<Asc<SalesOrder.orderNbr>>>.View.Select(this))
        {
            ...
        }
    }
}

```

### To Execute a BQL Statement Dynamically

To execute a BQL statement dynamically, do the following:

1. Dynamically instantiate a data view in code. You should also provide the graph object as a parameter to the data view constructor.
2. Execute the data view by using the `Select()` method of the instance of a `PXSelectBase`-derived class.
3. Cast the result of the query execution to a DAC or an array of DACs, or iterate through DACs in the result by using the `foreach` statement. The following sample code shows the approach of iterating through DACs. For details, see [To Process the Result of the Execution of the BQL Statement](#).

```

// Business logic controller (graph) declaration
public class OrderDataEntry : PXGraph<OrderDataEntry, SalesOrder>
{
    ...
    public void SomeMethod()
    {
        // Dynamic instantiation of a data view
        PXSelectBase<SalesOrder> orders =
            new SelectFrom<SalesOrder>.
                OrderBy<Asc<SalesOrder.orderNbr>>>.View(this);

        // An execution of a dynamically created BQL statement
        foreach(SalesOrder so in orders.Select())
        {
            ...
        }
    }
}

```

```
}
```

### To Execute a BQL Statement with Parameters

Use parameters (such as `AsOptional`, `FromCurrent`, and `@P.As[Type]` in fluent BQL and `Optional`, `Current`, and `Required` in traditional BQL) to pass specific values to a BQL statement, as shown in the following example. For more details on how to construct a BQL statement with parameters, see [To Use Parameters in Fluent BQL Queries](#) and [To Use Parameters in Traditional BQL](#).

```
// Declaration of a BLC
public class ReceiptDataEntry : PXGraph<ReceiptDataEntry, Document>
{
    // When a screen associated with this BLC is first opened,
    // the Optional parameter is replaced with the default DocType value.
    public SelectFrom<Document>
        Where<Document.docType.IsEqual<Document.docType.AsOptional>> Receipts;

    // The FromCurrent parameters are replaced with the values from
    // the Current property of the PXCache<Document> object.
    public SelectFrom<DocTransaction>
        Where<DocTransaction.docNbr.IsEqual<Document.docNbr.FromCurrent>.
        And<DocTransaction.docType.IsEqual<Document.docType.FromCurrent>>>.
        OrderBy<Asc<DocTransaction.lineNbr>> ReceiptTransactions;

    public void SomeMethod()
    {
        // Select documents of the same DocType as the current document
        // has, or of the default DocType if the current document is null.
        PXResult<Document> res1 = Receipts.Select();
        foreach(Document doc in res1)
            ...

        // Select documents of the "N" DocType.
        PXResult<Document> res2 = Receipts.Select("N");
        foreach(Document doc in res2)
            ...

        // Use parameter values from the current document.
        PXResult<DocTransaction> res3 = ReceiptTransactions.Select();
        foreach(DocTransaction docTran in res3)
            ...

        // Use the @P.AsString parameter to provide values in code.
        // The result set here is the same as the res2 result set.
        PXResult<Document> res4 =
            SelectFrom<Document>.
                Where<Document.docType.IsEqual<@P.AsString>>.View
                .Select(this, "N");
        foreach(Document doc in res4)
            ...
    }
    ...
}
```

### To Execute a BQL Statement in a Data View Delegate

If the data requested from the database cannot be described by a declarative BQL statement, implement the data view delegate that is used instead of the standard `Select()` logic to retrieve data from the database; this data view delegate must satisfy the following requirements:

- The data view delegate must have the same name as the data view except for the first letter, which must be lowercase.
- The data view delegate must return `IEnumerable`, as shown in the following example.



: If the data view delegate is not defined or it returns null, the standard `Select()` logic is executed.

The following sample code defines a data view delegate.

```
// A view declaration in a graph
public SelectFrom<BalancedAPDocument>.
    LeftJoin<APInvoice>.
        On<APInvoice.docType.IsEqual<BalancedAPDocument.docType>.
            And<APInvoice.refNbr.IsEqual<BalancedAPDocument.refNbr>>>.
    LeftJoin<APPayment>.
        On<APPayment.docType.IsEqual<BalancedAPDocument.docType>.
            And<APPayment.refNbr.IsEqual<BalancedAPDocument.refNbr>>>.View
    DocumentList;

// The data view delegate
protected virtual IEnumerable documentlist()
{
    // Iterating over the result set of a complex BQL statement
    foreach (PXResult<BalancedAPDocument, APInvoice, APPayment, APAdjust> res in
        SelectFrom<BalancedAPDocument>.
            LeftJoin<APInvoice>.
                On<APInvoice.docType.IsEqual<BalancedAPDocument.docType>.
                    And<APInvoice.refNbr.IsEqual<BalancedAPDocument.refNbr>>>.
            LeftJoin<APPayment>.
                On<APPayment.docType.IsEqual<BalancedAPDocument.docType>.
                    And<APPayment.refNbr.IsEqual<BalancedAPDocument.refNbr>>>.
            LeftJoin<APAdjust>.
                On<APAdjust.adjgDocType.IsEqual<BalancedAPDocument.docType>>.
            AggregateTo<GroupBy<BalancedAPDocument.docType>,
                GroupBy<BalancedAPDocument.refNbr>,
                GroupBy<BalancedAPDocument.released>,
                GroupBy<BalancedAPDocument.prebooked>,
                GroupBy<BalancedAPDocument.openDoc>>.View.Select(this))
    {
        // Casting a result set record to DAC types
        BalancedAPDocument apdoc = (BalancedAPDocument)res;
        APAdjust adj = (APAdjust)res;
        // Checking some conditions and modifying records
        ...
    }

    return new PXResult<BalancedAPDocument, APInvoice, APPayment>(
        apdoc, res, res);
}
```

## To Process the Result of the Execution of the BQL Statement

`Select()` returns the `PXResultset<T0>` object. The type parameter (`T0`) is set to the first table selected by the business query language (BQL) statement, and `PXResultset<T0>` is a collection of `PXResult<T0>` objects. You can iterate through the result set in a `foreach` loop and obtain either data access class (DAC) instances or `PXResult<>` instances. A `PXResult<>` instance represents a whole result set record and can be cast to any of the DAC types joined in the BQL statement.

### To Get the Objects of the Primary DAC

In the `foreach` loop, cast each `PXResult<T0>` object in the collection to an object of the main DAC. The `PXResult<T0>` object is implicitly converted to the `T0` class. In the following sample code, records are selected from the `Document` table.

```
// Result set records are implicitly cast to the Document DAC.
foreach(Document doc in SelectFrom<Document>.View.Select(this))
{
    ...
}
```



```
}
```

### To Get the Objects of Joined DACs

1. In the foreach loop, cast each `PXResult<T0>` object in the collection to the needed `PXResult<T0, T1, T2, ...>` object, where `T0`, `T1`, `T2`, and other type parameters are joined DACs from the BQL statement. The `PXResult<T0, T1, T2, ...>` type must be specialized with the DACs of all joined tables.
2. Cast each `PXResult<T0, T1, T2, ...>` item to any of the listed types to get the object of this type.

The following sample code shows how to process the result set of a BQL statement joining two tables.

```
// The static Select() method is called to execute a BQL command.
PXResultset<OrderLine> result =
    SelectFrom<OrderLine>.InnerJoin<SalesOrder>.
        On<SalesOrder.orderNbr.IsEqual<OrderLine.orderNbr>>.View.Select(this);

// Iterating over the result set:
// PXResult should be specialized with the DACs of all joined tables
// to be able to cast to these DACs.
foreach(PXResult<OrderLine, SalesOrder> record in result)
{
    // Casting a result set record to the OrderLine DAC:
    OrderLine detail = (OrderLine)record;
    // Casting a result set record to the SalesOrder DAC:
    SalesOrder order = (SalesOrder)record;
    ...
}
```

## Creating Fluent BQL Queries

To query data from the database, you use the business query language (BQL), which has two dialects: fluent BQL and traditional BQL.

In this chapter, you can find information on how to create fluent BQL queries. For general information about data querying, see [Querying Data in Acumatica Framework](#). For details about building queries with traditional BQL, see [Creating Traditional BQL Queries](#).

### In This Chapter

- [Fluent Business Query Language](#)
- [Data Access Classes in Fluent BQL](#)
- [Search and Select Commands and Data Views in Fluent BQL](#)
- [Constants in Fluent BQL](#)
- [Parameters in Fluent BQL](#)
- [To Select Records by Using Fluent BQL](#)
- [To Use Parameters in Fluent BQL Queries](#)

## Fluent Business Query Language

Fluent business query language (BQL), which is described in this topic, is a dialect of BQL that is more similar to SQL than traditional BQL is. You can find all classes that can be used in fluent BQL in the `PX.Data.BQL` and `PX.Data.BQL.Fluent` namespaces.

## Fluent BQL Structure

Fluent BQL uses nesting of generic classes. That is, each section of a fluent BQL query does not depend on the other sections and can appear in only specific places of the query. The order of the sections is shown in the following code.

```
SelectFrom<>.[Joins].Where<>.AggregateTo<>.OrderBy<>
```

`SelectFrom<>` is the only mandatory part of the query. You can add to the query any number of `Join` sections and the `Where<>`, `AggregateTo<>`, and `OrderBy<>` sections of the query, depending on whether you need the corresponding clauses of the query.

The query defined with fluent BQL as described above is equivalent to the `Select` command in traditional BQL. To compose a query for different purposes (such as to define a data view or to define a `Search` command in an attribute constructor), you need to prepend additional elements to the query or append them to the query, as described in [Search and Select Commands and Data Views in Fluent BQL](#). You can find the equivalents of traditional BQL in fluent BQL in [Fluent BQL and Traditional BQL Equivalents](#).

## SelectFrom<> Section

In the `SelectFrom<>` section of the query, you use the `SelectFrom<>` class, which uses a data access class (DAC) as the type parameter. For details on DACs, see [Data Access Classes in Fluent BQL](#).

## Join Sections

Each `Join` section of the fluent BQL query consists of the following components:

- The join type (`InnerJoin<>`, `LeftJoin<>`, `RightJoin<>`, `FullJoin<>`, `CrossJoin<>`) with the joined DAC as the type parameter.
- The joining condition (`On<>`). This condition is not specified for `CrossJoin<>`.
- The single table modifier (`SingleTableOnly`). This optional part of each `Join` section forces optimization if a DAC used in the query has the `PXProjection` attribute.

The following code fragments show the `Join` sections with different types of joins.

```
.InnerJoin<TBqlTable>.On<TJoinCondition>.SingleTableOnly
.LeftJoin<TBqlTable>.On<TJoinCondition>.SingleTableOnly
.RightJoin<TBqlTable>.On<TJoinCondition>.SingleTableOnly
.FullJoin<TBqlTable>.On<TJoinCondition>.SingleTableOnly
.CrossJoin<TBqlTable>.SingleTableOnly
```

Fluent BQL queries can contain any number of `Join` sections.

## Where<> Section and On<> Subsection

Conditions in the query are defined in the `Where<>` section and the `On<>` subsections of the `Join` sections. The conditions can contain the following nested components:

- Comparisons, such as `Table.field.IsEqual<TOperand>`
- `And<>` subsections
- `Or<>` subsections
- `Brackets<>` subsections

The following code fragments show examples of an `On<>` subsection and a `Where<>` section.

```
.On<PMTask.projectID.IsEqual<PMProject.contractID>.  
  And<  
    PMTask.approverID.IsEqual<EPActivityFilter.approverID.FromCurrent>>>>  
  
.Where<PMProject.isActive.IsEqual<True>.  
  And<PMTask.taskID.IsNotNull.  
    Or<PMProject.approverID.IsEqual<  
      EPActivityFilter.approverID.FromCurrent>>>>>>
```

### AggregateTo<> and OrderBy<> Sections

The `AggregateTo<>` and `OrderBy<>` sections of a fluent BQL query accept non-empty arrays of the specific base type as the only generic parameters. To make it easier for developers to write and read of the queries, fluent BQL includes groups of aliases that embed certain array usage. These aliases are pregenerated for arrays with up to 32 elements.

The `AggregateTo<>` section can also include an optional `Having<>` subsection. In this subsection, you include conditions that can contain only logical operators, constants, parameters, and aggregated fields (that is, the fields with `.Averaged`, `.Summarized`, `.Maximized`, `.Minimized`, or `.Grouped` appended).

The following code fragments show examples of `AggregateTo<>` and `OrderBy<>` sections.

```
.AggregateTo<Sum<field1>, GroupBy<field2>, Max<field3>,  
  Min<field4>, Avg<field5>, Count<field6>>>.  
  Having<field5.Averaged.IsGreater<Zero>>>  
  
.OrderBy<field1.Asc, field2.Desc, field3.Asc>
```

## Data Access Classes in Fluent BQL

The data access classes (DACs) that are used in fluent BQL differ from the DACs that are used in traditional BQL in the declarations of the class fields. For the general information about the declaration of DACs for both traditional BQL and fluent BQL, see [Data Access Classes](#).

Each class field of a DAC (that is, each `public abstract` class of a DAC) is strongly typed, which makes it possible to perform compile-time code checks in Visual Studio. You derive class fields not from the `IBqlField` interface (as you would in traditional BQL) but from the specific fluent BQL classes that correspond to the type of the property field as shown in the following table. You assign the class field a name that starts with a lowercase letter.

Type of the Property Field	Type of the Class Field
bool	BqlBool.Field<TSelf>
byte	BqlByte.Field<TSelf>
short	BqlShort.Field<TSelf>
int	BqlInt.Field<TSelf>
long	BqlLong.Field<TSelf>
float	BqlFloat.Field<TSelf>
double	BqlDouble.Field<TSelf>
decimal	BqlDecimal.Field<TSelf>
Guid	BqlGuid.Field<TSelf>
DateTime	BqlDateTime.Field<TSelf>

Type of the Property Field	Type of the Class Field
String	BqlString.Field<TSelf>
byte[]	BqlByteArray.Field<TSelf>

The following code shows an example of the `Product` DAC declaration.

```
using System;
using PX.Data;

[Serializable]
public class Product : PX.Data.IBqlTable
{
    // The class used in BQL statements to refer to the ProductID column
    public abstract class productID : PX.Data.BQL.BqlInt.Field<productID>
    {
    }
    // The property holding the ProductID value in a record
    [PXDBIdentity(IsKey = true)]
    public virtual int? ProductID { get; set; }

    // The class used in BQL statements to refer to the AvailQty column
    public abstract class availQty : PX.Data.BQL.BqlDecimal.Field<availQty>
    {
    }
    // The property holding the AvailQty value in a record
    [PXDBDecimal(2)]
    public virtual decimal? AvailQty { get; set; }
}
```

### Simultaneous Use of DACs in Fluent BQL and Traditional BQL

The DAC fields declared in fluent BQL style can be used in traditional BQL queries without any modifications.

The class fields that are defined in the traditional BQL style (as described in [Data Access Classes](#)) can be used in fluent BQL queries if you wrap these fields in the `Use<>.As[Type]` class, where `[Type]` is one of the following: `Bool`, `Byte`, `Short`, `Int`, `Long`, `Float`, `Double`, `Decimal`, `Guid`, `DateTime`, `String`, or `ByteArray`.

The following code shows the definition of the `availQty` class field in the traditional BQL style and its use in a fluent BQL comparison.

```
public class Product : PX.Data.IBqlTable
{
    public abstract class availQty : PX.Data.IBqlField
    {
    }
    [PXDBDecimal(2)]
    public virtual decimal? AvailQty { get; set; }
}

SelectFrom<Product>.
    Where<Use<Product.availQty>.AsDecimal.IsNotEqual<Zero>>.
    View AvailableProducts;
```

Though the DAC fields in the traditional BQL style can be used in fluent BQL queries, we recommend that you use the fluent BQL style of DAC declaration for simplicity.

## Search and Select Commands and Data Views in Fluent BQL

You can use fluent business query language (BQL) to define data views and specify `Select` and `Search` commands.

## Data Views

You can use any of the following approaches to define a data view:

- Use the `PXViewOf<>` class before the fluent BQL query, as shown in the following code example.

```
PXViewOf<Product>.BasedOn<
    SelectFrom<Product>.
        Where<Product.isActive.IsEqual<True>>>.ReadOnly ActiveProducts;
```

The DACs that are specified as type parameters in `PXViewOf<>` and `SelectFrom<>` must be the same; this is checked by the compiler. You can omit `.BasedOn<>` if you want to declare a view that selects all records from one table. You append `.ReadOnly` to the view definition if you need to define a read-only data view.

- Append `.View` to the fluent BQL query, as shown in the following code example.

```
SelectFrom<Product>.
    Where<Product.isActive, Equal<True>>>.View.ReadOnly ActiveProducts;
```

You append `.ReadOnly` to the view definition if you need to define a read-only data view.

The data views defined with fluent BQL are equivalent to the corresponding traditional BQL data views. For the full list of equivalents, see [Fluent BQL and Traditional BQL Equivalents](#). Also, the fluent BQL data views have the same static methods as the traditional BQL data views have.

## Select Commands

The query defined with fluent BQL, as described in [Fluent Business Query Language](#), is equivalent to the `Select` BQL command. For the full list of equivalents, see [Fluent BQL and Traditional BQL Equivalents](#).

## Search Commands

You can use any of the following approaches to define a `Search` BQL command:

- Use the `SearchFor<>` class before the fluent BQL query, as shown in the following code example.

```
SearchFor<Product.productId>.In<
    SelectFrom<Product>.
        Where<Product.isActive.IsEqual<True>>>
```

- Append `.SearchFor<>` to the fluent BQL query, as shown in the following code example.

```
SelectFrom<Product>.
    Where<Product.isActive.IsEqual<True>>>.SearchFor<Product.productId>
```

The `Search` commands defined with fluent BQL are equivalent to the corresponding traditional BQL commands. For the full list of equivalents, see [Fluent BQL and Traditional BQL Equivalents](#).

## Dynamic Query Building

Because `SearchFor<>` and `SelectFrom<>` are derived from the `BqlCommand` class, they can be used in dynamic query building through the `WhereAnd`, `AppendJoin`, and `OrderByNew` functions. However, fluent BQL commands (which are derived from the `FbqlCommand` class) aren't decomposed by `BqlCommand.Decompose()` directly. That is, the `Decompose` function checks whether a command has a `FbqlCommand` type, retrieves the type of its underlying `BqlCommand`, and decomposes this underlying `BqlCommand` command. Therefore, the elements of the array produced by the `Decompose` function are components of the `BqlCommand`, and are not components of the passed `FbqlCommand`.

`PXViewOf<TBqlTable>` also supports all dynamic query building actions that the traditional BQL `PXView` supports.

## Constants in Fluent BQL

You can use predefined constants (such as integer `Zero`, datetime `Now`, `Today`, and `MaxDate`, string `StringEmpty`, and the Boolean values `True` and `False`) in fluent BQL queries without any changes.

If you need to use a custom constant in a fluent BQL query, you define this constant by using the class that corresponds to the C# type of the constant. The following table lists the constant classes that correspond to C# types.

C# Type	Fluent BQL Type
bool	<code>BqlBool.Constant&lt;TSelf&gt;</code>
byte	<code>BqlByte.Constant&lt;TSelf&gt;</code>
short	<code>BqlShort.Constant&lt;TSelf&gt;</code>
int	<code>BqlInt.Constant&lt;TSelf&gt;</code>
long	<code>BqlLong.Constant&lt;TSelf&gt;</code>
float	<code>BqlFloat.Constant&lt;TSelf&gt;</code>
double	<code>BqlDouble.Constant&lt;TSelf&gt;</code>
decimal	<code>BqlDecimal.Constant&lt;TSelf&gt;</code>
Guid	<code>BqlGuid.Constant&lt;TSelf&gt;</code>
DateTime	<code>BqlDateTime.Constant&lt;TSelf&gt;</code>
String	<code>BqlString.Constant&lt;TSelf&gt;</code>

The following code shows an example of the `decimal_0` constant declaration.

```
public class decimal_0 : PX.Data.BQL.BqlDecimal.Constant<decimal_0>
{
    public decimal_0()
        : base(0m)
    {
    }
}
```

### Simultaneous Use of Constants in Fluent BQL and Traditional BQL

The predefined constants and the constants defined as described in the previous section can be used in traditional BQL without any changes.

The constants defined in the traditional BQL style (that is, derived from the `Constant<Type>` class) can be used in the fluent BQL queries if you wrap these constants in the `Use<>.As<Type>` class, where `[Type]` is one of the following: `Bool`, `Byte`, `Short`, `Int`, `Long`, `Float`, `Double`, `Decimal`, `Guid`, `DateTime`, or `String`.

The following code shows the declaration of the `decimal_0` constant in traditional BQL style and its use in a fluent BQL comparison.

```
public class decimal_0 : Constant<Decimal>
{
    public decimal_0()
        : base(0m)
    {
    }
}

SelectFrom<Table>.
    Where<Table.decimalField.AsDecimal.IsEqual<Use<decimal_0>.AsDecimal>>.
```

```
View records;
```

Although the constants in the traditional BQL style can be used in fluent BQL queries, we recommend that you use the fluent BQL style of constant declaration for simplicity.

## Parameters in Fluent BQL

If you need to specify values in a fluent business query language (BQL) statement, you use fluent BQL parameters, which are replaced with the needed values in the translation to SQL. For details about how BQL statements with parameters are translated to SQL, see [Translation of a BQL Command with Parameters to an SQL Query Tree](#).

In this topic, you can find a description of the fluent BQL parameters and the difference between them.

### Use of the Current Value of the Field from PXCache

To insert into the SQL query the field value of the `Current` object from the `PXCache` object, you append `.FromCurrent` to the field name in a fluent BQL query. If the `Current` object from the `PXCache` object is null, `FromCurrent` retrieves the default value of the field. If you do not need to retrieve the default value if the `Current` object is null, you need to append `.FromCurrent.NoDefault` to the field name in a fluent BQL query. In this case, the system doesn't retrieve the default value and inserts null.



: `FromCurrent` is the equivalent of the `Current` parameter in traditional BQL.  
`FromCurrent.NoDefault` is the equivalent of the `Current2` parameter in traditional BQL.

By using the current field value from `PXCache` in the declaration of a data view, you can refer to another view to relate these data views to each other. A typical example is referencing the current master record on master-detail webpages. For details on how the current field value is used, see [To Relate Data Views to Each Another](#).

### Insertion of a Specific Value into the Query

To insert a specific value into the SQL query, you use the `@P.As[Type]` classes, where `[Type]` corresponds to the C# type of the parameter. The following table lists the fluent BQL types that correspond to C# types.

C# Type	Fluent BQL Type
bool	@P.AsBool
byte	@P.AsByte
short	@P.AsShort
int	@P.AsInt
long	@P.AsLong
float	@P.AsFloat
double	@P.AsDouble
decimal	@P.AsDecimal
Guid	@P.AsGuid
DateTime	@P.AsDateTime
String	@P.AsString



: `@P.As[Type]` is the equivalent of the `Required` parameter in traditional BQL.

By using these classes, you can pass values to the SQL query, as described in [To Pass a Field Value to the SQL Query](#) and [To Pass Multiple Field Values to the SQL Query](#).

### Insertion of an Optional Value into the Query

To insert an optional value into the query, you append `.AsOptional` to the field name in a fluent BQL query. If you specify an explicit value for this parameter during the execution of the BQL statement, `AsOptional` uses the specified value. If you don't specify an explicit value for this parameter during the execution of the BQL statement, `AsOptional` works similarly to `FromCurrent`—that is, retrieves the field value of the `Current` object from the `PXCache` object and uses the default value of the field if the `Current` object is null. You can append `.AsOptional.NoDefault` to the field name in a fluent BQL query to make the system not use the default value and insert null.



: `AsOptional` is the equivalent of the `Optional` parameter in traditional BQL. `AsOptional.NoDefault` is the equivalent of the `Optional2` parameter in traditional BQL.

By using `AsOptional`, you can pass the external presentations of the values to the SQL query, as described in [To Provide External Presentation of the Field Value to the SQL Query](#).

### Insertion of a Value from the UI Control into the Query

To insert a value from the UI control into the SQL query, you use the `Argument.As[Type]` classes, where `[Type]` corresponds to the C# type of the inserted value. The following table lists the fluent BQL types that correspond to C# types.

C# Type	Fluent BQL Type
bool	<code>Argument.AsBool</code>
byte	<code>Argument.AsByte</code>
short	<code>Argument.AsShort</code>
int	<code>Argument.AsInt</code>
long	<code>Argument.AsLong</code>
float	<code>Argument.AsFloat</code>
double	<code>Argument.AsDouble</code>
decimal	<code>Argument.AsDecimal</code>
Guid	<code>Argument.AsGuid</code>
DateTime	<code>Argument.AsDateTime</code>
String	<code>Argument.AsString</code>

By using the `Argument` classes, you can pass values to the data view delegates. For more information on how to use the `Argument` classes, see [To Pass a Value from a UI Control to a Data View](#).

## To Select Records by Using Fluent BQL

You can select records from the database by constructing a fluent business query language (BQL) statement. To construct a fluent BQL statement, you use the `SelectFrom<>` class and append the needed clauses to the statement.

This topic describes how to compose `Select` statements by using fluent BQL. For details on how to adjust these statements to define data views or to specify `Search` commands in fluent BQL, see [Search and Select Commands and Data Views in Fluent BQL](#).





: In a `SelectFrom<>` class, you configure a query to the database. The actual request to the database is performed once you cast the result of the query execution to a DAC or an array of DACs, or when you iterate through DACs in the result with the `foreach` statement. For details, see [Data Query Execution](#).

### Before You Proceed

- Make sure that the application database has the database tables from which you are going to request data, and that the application defines the data access classes (DACs) for these tables. For more information on defining DACs, see [Data Access Classes in Fluent BQL](#).
- Add references to `PX.Data.dll` and `PX.Data.BQL.Fluent.dll` in the project.
- Add the following `using` directives to your code.

```
using PX.Data.BQL.Fluent;
using PX.Data.BQL;
```

### To Compose a Fluent BQL Statement

1. Type the `SelectFrom<>` class with the needed DAC as the type parameter.

For example, suppose that you need to convert the following SQL statement to fluent BQL.

```
SELECT Product.CategoryCD, MIN(Product.BookedQty) FROM Product
INNER JOIN SupplierProduct
    ON SupplierProduct.ProductID = Product.ProductID
INNER JOIN Supplier
    ON Supplier.AccountID = SupplierProduct.AccountID
WHERE (Product.BookedQty IS NOT NULL
    AND Product.AvailQty IS NOT NULL
    AND Product.MinAvailQty IS NOT NULL
    AND (Product.Active = 1
        OR Product.Active IS NULL)
    AND (Product.BookedQty > Product.AvailQty
        OR Product.AvailQty < Product.MinAvailQty))
    OR Product.AvailQty IS NOT NULL
GROUP BY Product.CategoryCD
ORDER BY Product.UnitPrice, Product.AvailQty DESC
```

You start the corresponding fluent BQL query as follows.

```
SelectFrom<Product>
```

2. If you need to include `JOIN` clauses in the query, for each table that you want to join, do the following:
  - a. Append to the statement one of the `Join` classes—such as `InnerJoin`, `LeftJoin`, `RightJoin`, `FullJoin`, or `CrossJoin`, which are directly mapped to SQL `JOIN` clauses.
  - b. Append to the statement the `On<>` clause with the joining conditions. Adhere to the following rules when you specify the conditions:
    - Use the `And<>`, `Or<>`, and `Brackets<>` classes to logically connect the conditions and comparisons.
    - To specify the fields that should be used in the conditions, use the class fields defined in the DAC, such as `Product.productID`. (The name of the field class starts with a lowercase letter. Do not confuse it with the property field, which has the same name but starts with an uppercase letter.)
    - If you need to use constants in the fluent BQL statement, use one of the predefined BQL constants or your own constant. (For details on using constants, see [Constants in Fluent BQL](#).)

- If you need to specify the values of the parameters at run time, use the fluent BQL parameters. For information about parameters, see [Parameters in Fluent BQL](#). For information about how to use parameters, see [To Use Parameters in Fluent BQL Queries](#).

In the sample code that has been presented in this topic, you would add two `Join` classes to the statement, as follows.

```
SelectFrom<Product>.
    InnerJoin<SupplierProduct>.
        On<SupplierProduct.productID.IsEqual<Product.productID>>.
    InnerJoin<Supplier>.
        On<Supplier.accountID.IsEqual<SupplierProduct.accountID>>
```

**3.** If you need to include a `WHERE` clause in the query, append the `Where<>` clause to the statement and specify the conditions as follows:

- Use the `And<>`, `Or<>`, and `Brackets<>` classes to logically connect the conditions and comparisons.
- To specify the fields that should be used in the conditions, use the class fields defined in the DAC, such as `Product.productID`. (The name of the field class starts with a lowercase letter. Do not confuse it with the property field, which has the same name but starts with an uppercase letter.)
- If you need to use constants in the fluent BQL statement, use one of the predefined BQL constants or your own constant. (For details on using constants, see [Constants in Fluent BQL](#).)
- If you need to specify the values of the parameters at run time, use the fluent BQL parameters. For information about parameters, see [Parameters in Fluent BQL](#). For information about how to use parameters, see [To Use Parameters in Fluent BQL Queries](#).

In the sample code that has been presented in this topic, you would append the `Where<>` clause to the statement, as follows.

```
SelectFrom<Product>.
    InnerJoin<SupplierProduct>.
        On<SupplierProduct.productID.IsEqual<Product.productID>>.
    InnerJoin<Supplier>.
        On<Supplier.accountID.IsEqual<SupplierProduct.accountID>>.
    Where<
        Brackets<Product.bookedQty.IsNotNull.
            And<Product.availQty.IsNotNull>.
            And<Product.minAvailQty.IsNotNull>.
            And<Product.active.IsEqual<True>.
                Or<Product.active.IsNull>>.
            And<Product.bookedQty.IsGreater<Product.availQty>.
                Or<Product.availQty.IsLess<Product.minAvailQty>>>>.
        Or<Product.availQty.IsNotNull>>
```

**4.** If you need to group or aggregate records, append the `AggregateTo<>` clause to the statement and specify the grouping conditions and aggregation functions by using the `GroupBy` clauses and the `Min`, `Max`, `Sum`, `Avg`, and `Count` aggregation functions.

In the sample code that has been presented in this topic, you would append the `AggregateTo<>` clause to the statement as follows.

```
SelectFrom<Product>.
    InnerJoin<SupplierProduct>.
        On<SupplierProduct.productID.IsEqual<Product.productID>>.
    InnerJoin<Supplier>.
        On<Supplier.accountID.IsEqual<SupplierProduct.accountID>>.
    Where<
```

```

        Brackets<Product.bookedQty.IsNotNull.
            And<Product.availQty.IsNotNull>.
            And<Product.minAvailQty.IsNotNull>.
            And<Product.active.IsEqual<True>.
                Or<Product.active.IsNull>>.
            And<Product.bookedQty.IsGreater<Product.availQty>.
                Or<Product.availQty.IsLess<Product.minAvailQty>>>>.
        Or<Product.availQty.IsNotNull>>.
    AggregateTo<GroupBy<Product.categoryCD>,
        Min<Product.bookedQty>>

```

5. If you need to order records, append to the statement the `OrderBy<>` clause with the `Asc<>` and `Desc<>` classes as the type parameters.

In the sample code that has been presented in this topic, you would append the `OrderBy<>` clause to the statement as follows.

```

SelectFrom<Product>.
    InnerJoin<SupplierProduct>.
        On<SupplierProduct.productID.IsEqual<Product.productID>>.
    InnerJoin<Supplier>.
        On<Supplier.accountID.IsEqual<SupplierProduct.accountID>>.
    Where<
        Brackets<Product.bookedQty.IsNotNull.
            And<Product.availQty.IsNotNull>.
            And<Product.minAvailQty.IsNotNull>.
            And<Product.active.IsEqual<True>.
                Or<Product.active.IsNull>>.
            And<Product.bookedQty.IsGreater<Product.availQty>.
                Or<Product.availQty.IsLess<Product.minAvailQty>>>>.
        Or<Product.availQty.IsNotNull>>.
    AggregateTo<GroupBy<Product.categoryCD>,
        Min<Product.bookedQty>>.
    OrderBy<Product.unitPrice.Asc,
        Product.availQty.Desc>

```

## To Use Parameters in Fluent BQL Queries

If you need to specify values in a fluent business query language (BQL) statement, you use fluent BQL parameters, which are replaced with the needed values in the translation to SQL. For details about how BQL statements with parameters are translated to SQL, see [Translation of a BQL Command with Parameters to an SQL Query Tree](#).

You may need to use BQL parameters to relate data views to each other, to pass field values to the SQL query, to pass the external presentations of the values to the SQL query, or to pass values from UI controls to the SQL query.

### To Relate Data Views to Each Another

To relate data views to each another, in a data query, use the field value of the `Current` object from the `PXCache` object, as shown in the following sample code.

```

using PX.Data;
using PX.Data.BQL.Fluent;

// The view declarations in a graph
SelectFrom<Document>.View Documents;
SelectFrom<DocTransaction>.
    Where<DocTransaction.docNbr.IsEqual<Document.docNbr.FromCurrent>.
        And<DocTransaction.docType.IsEqual<Document.docType.FromCurrent>>>.View
    DocTransactions;

```

In this code, there is a many-to-one relationship between the `DocTransaction` and `Document` data access classes (DACs), and this relationship is implemented through the `DocNbr` and `DocType` key fields. The data views in the code connect the `Document` and `DocTransaction` records.



: Acumatica Framework translates the fluent BQL query of the second view in the sample code to the following SQL statement. In this SQL query, [parameter1] is the DocNbr value and [parameter2] is the DocType value retrieved from the Current property of the DocTransaction cache; [list of columns] is the list of columns of the DocTransaction table.

```
SET @P0 = [parameter1]
SET @P1 = [parameter2]

SELECT * FROM DocTransaction
WHERE DocTransaction.DocNbr = @P0
      AND DocTransaction.DocType = @P1
```

Acumatica Framework explicitly enumerates the columns of the database table in the SQL query. For details on which columns are enumerated, see [Translation of a BQL Command to SQL](#).

## To Pass a Field Value to the SQL Query

To pass a specific value to the SQL query, do the following;

1. Use the `@P.As[Type]` class of the needed type in the BQL statement, where [Type] is one of the following: Bool, Byte, Short, Int, Long, Float, Double, Decimal, Guid, DateTime, or String.
2. Specify the needed value as the `Select()` method argument. The value passed to `Select()` must be of the same type as the type of the specified field.



: The `@P.As[Type]` class must be used only in the BQL statements that are directly executed in the application code. The data views that are queried from the UI will not work if they contain this class.

The code below shows the execution of a BQL statement with a specific value passed in code.

```
using PX.Data;
using PX.Data.BQL.Fluent;
using PX.Data.BQL;

// Suppose an event handler related to the Product DAC
// is being executed.
Product product = (Product)e.Row;

// Select the Category record with the specified CategoryCD.
Category category =
    SelectFrom<Category>.
        Where<Category.categoryCD.IsEqual<@P.AsString>>.View.
        Select(this, product.CategoryCD);
```



: Acumatica Framework translates the previous fluent BQL query to the following SQL statement. In this SQL query, [parameter] is the value of the product.CategoryCD variable at the moment the `Select()` method is invoked; [list of columns] is the list of columns of the Category table.

```
SET @P0 = [parameter]

SELECT * FROM Category
WHERE Category.CategoryCD = @P0
```

Acumatica Framework explicitly enumerates the columns of the database table in the SQL query. For details on which columns are enumerated, see [Translation of a BQL Command to SQL](#).

## To Pass Multiple Field Values to the SQL Query

To pass multiple values to the SQL query, do the following:

1. Use multiple `@P.As[Type]` classes of the needed type in the fluent BQL statement, where [Type] is one of the following: Bool, Byte, Short, Int, Long, Float, Double, Decimal, Guid, DateTime, or String.

2. Specify the needed values as the `Select()` method arguments in the order in which the parameters are specified in the BQL statement. The number of `@P.As[Type]` classes must match the number of parameters passed to the `Select()` method.



: The `@P.As[Type]` classes should be used in only the BQL statements that are executed in the application code.

The following code shows an example of a fluent BQL statement with two `Required` parameters.

```
using PX.Data;
using PX.Data.BQL.Fluent;
using PX.Data.BQL;

// Suppose an event handler related to the DocTransaction DAC
// is being executed.
DocTransaction line = (DocTransaction)e.Row;

Document doc =
    SelectFrom<Document>.
        Where<Document.docNbr.IsEqual<@P.AsString>.
            And<Document.docType.IsEqual<@P.AsString>>>.View.
    Select(this, line.DocNbr, line.DocType);
```



: Acumatica Framework translates the previous fluent BQL query to the following SQL statement, where `[list of columns]` is the list of the columns of the `Document` table.

```
SET @P0 = [line.DocNbr value]
SET @P1 = [line.DocType value]

SELECT * FROM Document
WHERE Document.DocNbr = @P0
      AND Document.DocType = @P1
```

Acumatica Framework explicitly enumerates the columns of the database table in the SQL query. For details on which columns are enumerated, see [Translation of a BQL Command to SQL](#).

## To Pass a Value from a UI Control to a Data View

To pass a value from a UI control to a data view, do the following:

1. Define a data view with the `Argument.As[Type]` class, where `[Type]` specifies the data type of the expected value, as shown in the following sample BQL query.

```
using PX.Data;
using PX.Data.BQL.Fluent;

SelectFrom<TreeViewItem>.
    Where<TreeViewItem.parentID.IsEqual<Argument.AsInt>>.
    OrderBy<Asc<TreeViewItem.parentID>>.View GridDataSource;
```



: Acumatica Framework translates the previous fluent BQL query to the following SQL statement. In this SQL query, `[parameter]` will contain the value received from the UI control and passed to the `Select()` method; `[list of columns]` is the list of columns of the `TreeViewItem` table.

```
SET @P0 = [parameter]

SELECT [list of columns] FROM TreeViewItem
WHERE TreeViewItem.ParentID = @P0
ORDER BY TreeViewItem.ParentID
```

Acumatica Framework explicitly enumerates the columns of the database table in the SQL query. For details on which columns are enumerated, see [Translation of a BQL Command to SQL](#).

2. Define the data view delegate with parameters through which you can access the values passed from the UI. (You can find more information on how to use data view delegates in [To Execute a BQL Statement in a Data View Delegate.](#))



: When a BQL statement with the `Argument` parameter is executed in code, the value must be specified in the parameters of the `Select()` method.

### To Provide External Presentation of the Field Value to the SQL Query

To substitute a value in the SQL query, do the following:

1. Add the `PXSelector` attribute with a substitute key to a DAC field, as shown in the following example.

```
using PX.Data;
using PX.Data.BQL.Fluent;

[PXSelector(typeof(SearchFor<Product.productID>.In<SelectFrom<Product>>,
    new Type [] {
        typeof(Product.productCD),
        typeof(Product.productName)
    },
    SubstituteKey = typeof(Product.productCD))]
public virtual int? ProductID { get; set; }
```

In this example, `PXSelector` replaces the `ProductID` field in the user interface with the human-readable `ProductCD` field. In the UI control for this field, the user enters a `ProductCD` value. The `PXSelector` attribute implements the `FieldUpdating` event handler, which replaces the `ProductCD` value with the corresponding `ProductID` value.

2. Use the `AsOptional` class to select records by the external presentation of the field value, as shown in the following code for `OrderLine` records.

```
using PX.Data;
using PX.Data.BQL.Fluent;

// p is a Product data record.
// To select OrderLine records, pass the ProductCD value
// to the Select() method.
PXResultset<OrderLine> details =
    SelectFrom<OrderLine>.
        Where<OrderLine.productID.IsEqual<OrderLine.productID.AsOptional>>.
        View.
        Select(this, p.ProductCD);
```

3. In the `Select()` method, provide values for all `AsOptional`, `@P.As[Type]`, and `Argument.As[Type]` parameters up to the last `@P.As[Type]` or `Argument.As[Type]` parameter in the fluent BQL statement, as shown in the following sample code.

```
using PX.Data;
using PX.Data.BQL.Fluent;

// P is a Product data record.
// od is an OrderLine data record.

// At least three values (in addition to the graph reference) must
// be passed to the Select() method below.
// The second AsOptional parameter here will be replaced with the
// default UnitPrice value.
PXResultset<OrderLine> details =
    SelectFrom<OrderLine>.
        Where<OrderLine.productID.IsEqual<OrderLine.productID.AsOptional>.
        And<OrderLine.extPrice.IsLess<@P.AsDecimal>>.
        And<OrderLine.unitPrice.IsGreater<@P.AsDecimal>>.
```

```
And<OrderLine.taxRate.IsEqual<OrderLine.taxRate.AsOptional>>>.View.  
    .Select(this, p.ProductCD, od.ExtPrice, od.UnitPrice);
```



: Acumatica Framework translates the fluent BQL query in the code to the following SQL statement, where [list of columns] is the list of columns of the OrderLine table.

```
SET @P0 = [line.ProductID value or default]  
SET @P1 = [line.ExtPrice value]  
SET @P2 = [line.UnitPrice value]  
SET @P3 = [Default TaxRate value]  
  
SELECT [list of columns] FROM OrderLine  
WHERE OrderLine.ProductID = @P0  
      AND OrderLine.ExtPrice < @P1  
      AND OrderLine.UnitPrice > @P2  
      AND OrderLine.TaxRate = @P3
```

Acumatica Framework explicitly enumerates the columns of the database table in the SQL query. For details on which columns are enumerated, see [Translation of a BQL Command to SQL](#).

## Creating Traditional BQL Queries

To query data from the database, you use the business query language (BQL), which has two dialects: fluent BQL and traditional BQL.

In this chapter, you can find information on how to create traditional BQL queries. For the general information about BQL, see [Querying Data in Acumatica Framework](#). For details about building queries with traditional BQL, see [Creating Traditional BQL Queries](#).

### In This Chapter

- [Traditional Business Query Language](#)
- [Data Access Classes in Traditional BQL](#)
- [PXSelect Classes](#)
- [The Classes That Compose BQL Statements](#)
- [Parameters in Traditional BQL Statements](#)
- [Data Query Execution](#)
- [To Select Records By Using Traditional BQL](#)
- [To Filter Records](#)
- [To Order Records](#)
- [To Query Multiple Tables](#)
- [To Group and Aggregate Records](#)
- [To Use Parameters in Traditional BQL](#)
- [To Use Arithmetic Operations](#)
- [To Compose a BQL Statement from an SQL Statement](#)

## Traditional Business Query Language

When you compose a query with the traditional business query language (BQL), you work with the following classes:

- The classes that correspond to database tables (data access classes) and columns. For details on data access classes, see [Data Access Classes in Traditional BQL](#).
- The classes that define data views in a graph and select data from the database in code (PXSelect classes). For more information on these classes, see [PXSelect Classes](#).
- The classes that compose BQL statements, such as `Select`, `Search`, `Where`, `OrderBy`, `And`, and `Add`. For more information on these classes, see [The Classes That Compose BQL Statements](#).

- The classes that pass parameters to BQL statements, such as `Current`, `Required`, `Optional`, `Argument`. For details on BQL parameters, see [Parameters in Traditional BQL Statements](#).

## Data Access Classes in Traditional BQL

The data access classes (DACs) that are used in traditional BQL differ from the DACs that are used in fluent BQL in the declarations of the class fields. For details about the DAC declaration, see [Data Access Classes](#).

You derive each class field of a DAC (a `public abstract` class of a DAC) from the `IBqlField` interface and assign it a name that starts with a lowercase letter.

The following code shows an example of the `Product` data access class declared in traditional BQL style.

```
using System;
using PX.Data;

[Serializable]
public class Product : PX.Data.IBqlTable
{
    // The class used in BQL statements to refer to the ProductID column
    public abstract class productID : PX.Data.IBqlField
    {
    }
    // The property holding the ProductID value in a record
    [PXDBIdentity(IsKey = true)]
    public virtual int? ProductID { get; set; }

    // The class used in BQL statements to refer to the AvailQty column
    public abstract class availQty : PX.Data.IBqlField
    {
    }
    // The property holding the AvailQty value in a record
    [PXDBDecimal(2)]
    public virtual decimal? AvailQty { get; set; }
}
```

### Simultaneous Use of DACs in Traditional BQL and Fluent BQL

The class fields declared in traditional BQL style cannot be used in fluent BQL queries.

The class fields that are defined in the fluent BQL style (as described in [Data Access Classes in Fluent BQL](#)) can be used in traditional BQL queries without any modifications. Therefore, we recommend that you use the fluent BQL style of DAC declaration.

## PXSelect Classes

In traditional business query language (BQL), you define a data view or request database data in code by using one of the `PXSelect` classes (that is, the classes derived from `PXSelectBase`).

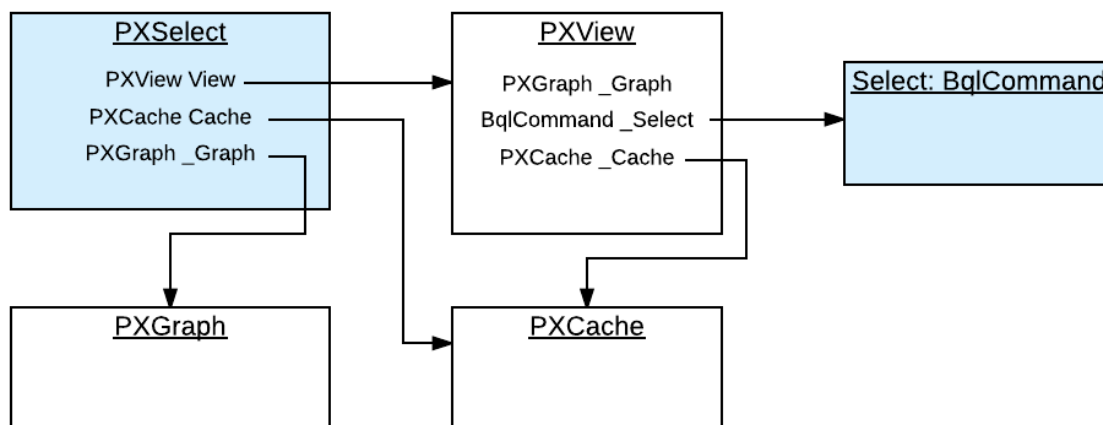
### PXSelect Classes

The instances of `PXSelect` classes are complex objects containing the following:

- A reference to the `PXView` object instantiated to process the data query
- A reference (through the `PXView` object) to the `Select` object, which is the business query language (BQL) command to be executed
- A reference to the graph
- A reference to the cache of the data access class (DAC) type that is specified in the first type parameter of `PXSelect`

That is, through the `PXSelect` classes, you can execute the BQL command and interact with the cache, as illustrated in the following diagram.





: Do not confuse the `PXSelect` classes with the `Select` classes. `PXSelect` is an aggregate of the data view, cache, and graph. You can use `PXSelect` classes to read, write, update, and delete records in the scope of a graph. `Select` classes simply represent BQL commands. You cannot read records by using a BQL command without instantiating a data view. For more information on the `Select` classes, see [The Classes That Compose BQL Statements](#).

### Types of PXSelect Classes

The first type parameter of all `PXSelect` classes is a data access class (DAC) generally bound to a database table. The resulting SQL query selects records from this table. Other type parameters (such as `Where`, `OrderBy`, `Join`, and `Aggregate`) are optional and represent clauses that can be added to the basic select statement.

Depending on the clauses that will be used in a query, you select the appropriate variant of the `PXSelect` class.

For example, if you need to use the `Where`, `OrderBy`, and `Join` clauses, you can use the `PXSelectJoin<Table, Join, Where, OrderBy>` class to create the query, as shown in the following BQL sample code.

```
PXSelectJoin<Table1,
    LeftJoin<Table2, On<Table2.field2, Equal<Table1.field1>>>,
    Where<Table1.field3, IsNotNull>,
    OrderBy<Asc<Table1.field1>>>>
```



: Acumatica Framework translates this statement to the following SQL query, where [list of columns] is the list of columns of the joined tables.

```
SELECT [list of columns] FROM Table1
LEFT JOIN Table2 ON Table2.Field2 = Table1.Field1
WHERE Table1.Field3 IS NOT NULL
ORDER BY Table1.Field1
```

Acumatica Framework explicitly enumerates the columns of the database table in the SQL query. For details on which columns are enumerated, see [Translation of a BQL Command to SQL](#).

For more information on how to use the BQL clauses, see [To Select Records By Using Traditional BQL](#).

If you need to retrieve data as it is currently stored in the database, you use one of the `PXSelect` classes that has `Readonly` in its name, such as the `PXSelectReadonly<Table>` class, or any of the `PXSelect` classes that use aggregation, such as the `PXSelectGroupBy<Table, Aggregate>` class. Otherwise, the data retrieved from the database can be merged with the data currently stored in the

cache. For more information on how the data is merged with the cache, see [Merge of the Records with PXCache](#).

### The List of PXSelect Classes

Acumatica Framework provides the following `PXSelect` classes:

- `PXSelect<Table, Where, OrderBy>`
- `PXSelect<Table, Where>`
- `PXSelect<Table>`
- `PXSelectGroupBy<Table, Aggregate>`
- `PXSelectGroupBy<Table, Where, Aggregate, OrderBy>`
- `PXSelectGroupBy<Table, Where, Aggregate>`
- `PXSelectGroupByOrderBy<Table, Aggregate, OrderBy>`
- `PXSelectGroupByOrderBy<Table, Join, Aggregate, OrderBy>`
- `PXSelectJoin<Table, Join, Where, OrderBy>`
- `PXSelectJoin<Table, Join, Where>`
- `PXSelectJoin<Table, Join>`
- `PXSelectJoinGroupBy<Table, Join, Aggregate>`
- `PXSelectJoinGroupBy<Table, Join, Where, Aggregate, OrderBy>`
- `PXSelectJoinGroupBy<Table, Join, Where, Aggregate>`
- `PXSelectJoinOrderBy<Table, Join, OrderBy>`
- `PXSelectOrderBy<Table, Join, OrderBy>`
- `PXSelectOrderBy<Table, OrderBy>`
- `PXSelectReadonly<Table, Where, OrderBy>`
- `PXSelectReadonly<Table, Where>`
- `PXSelectReadonly<Table>`
- `PXSelectReadonly2<Table, Join, Where, OrderBy>`
- `PXSelectReadonly2<Table, Join, Where>`
- `PXSelectReadonly2<Table, Join>`
- `PXSelectReadonly3<Table, Join, OrderBy>`
- `PXSelectReadonly3<Table, OrderBy>`

## The Classes That Compose BQL Statements

This topic contains an overview of the classes that you use to compose business query language (BQL) statements inside `PXSelect` and to define attributes of DACs.

### Overview of the Classes

Almost all classes that compose BQL statements are derived from the `IBqlCreator` interface, which inherits from the `IBqlVerifier` interface. These interfaces provide the following key methods:

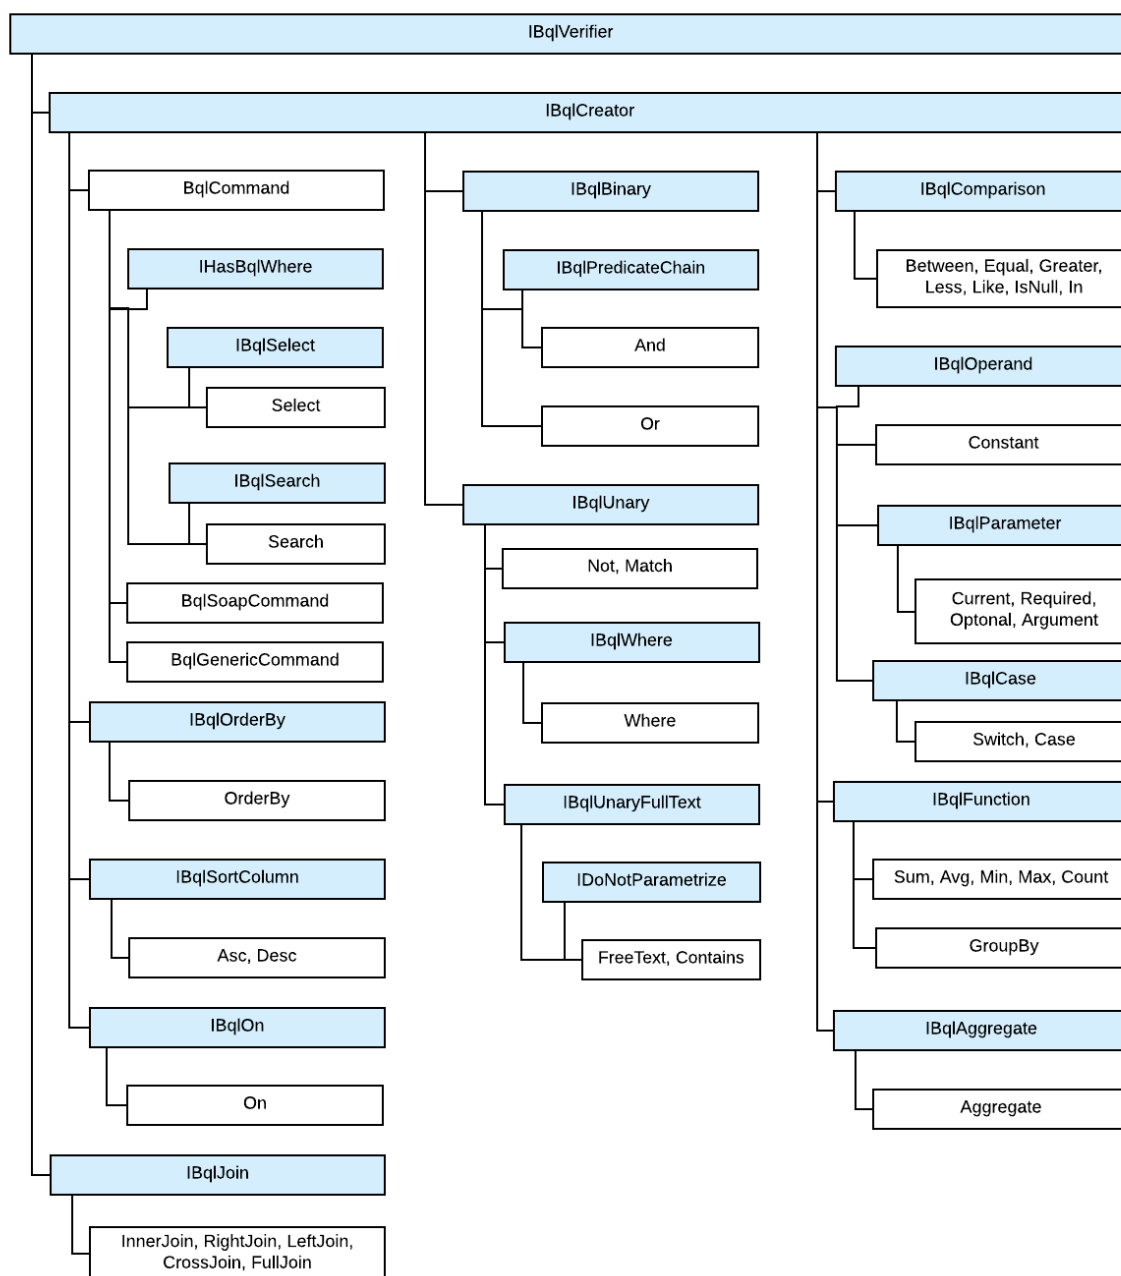
- `IBqlCreator.AppendExpression()`: Used during a BQL command preparation to translate a BQL statement into an SQL tree expression, which is then produces the SQL text to be sent to

the database maintenance server. For more information on how this method is used during BQL statement execution, see [Translation of a BQL Command to SQL](#).

- `IBqlVerifier.Verify()`: Used during the merge of the records with `PXCache` to evaluate a condition on a record retrieved from the database or calculate an expression with the record. For details on the merge, see [Merge of the Records with PXCache](#).

Depending on the purpose of each BQL class, the class also implements the methods of the interfaces derived from the `IBqlCreator` interface. For example, the aggregation functions—such as `Sum`, `Avg`, `Min`, and `Max`—implement the methods of the `IBqlFunction` interface.

The high-level overview of BQL class inheritance is illustrated in the following diagram. For descriptions of the interfaces and classes, see the API Reference.



**Figure: BQL commands**

The sections below describe the classes derived from the `BqlCommand` class.

### Select Classes

The `Select` classes, which are derived from the `BqlCommand` class, represent BQL commands and select all bound fields of the DAC and the unbound fields with specific attributes, such as `PXDBCalced`.



: More specific, the `Select` classes select all DAC fields that are decorated with the attributes that subscribe to the `PXCommandPreparing` event. For details on which fields are selected, see [Translation of a BQL Command to SQL](#).

In a BQL expression based on `Select`, the first type parameter is a DAC, as shown in the following sample BQL statement.

```
Select<Product>
```

The `Select` classes can parse themselves into SQL and provide methods for modifying the BQL command. However, you cannot directly use the `Select` class to execute the BQL query. Typically, you use `Select` in attributes in DACs, such as the `PXProjection` attribute.

### Search Classes

The `Search` classes, which are derived from the `BqlCommand` class, select one field of a DAC (while the `Select` classes select multiple fields).

In a `Search`-based statement, the first type parameter is a DAC field, as shown in the following sample BQL expression. This expression selects the `Product.unitPrice` field.

```
Search<Product.unitPrice>
```

These classes can parse themselves into SQL and provide methods for modifying the BQL command. However, you cannot directly use the `Search` class to execute the BQL query. Typically, you use `Search` in attributes in DACs, such as the `PXSelector` attribute. (`PXSelectorAttribute` requires a `Search` class and not a `Select` because the lookup control, which is configured by this attribute, displays precisely one field (usually a key field), which is what `Search` returns.)

### BqlCommand Classes

The `BqlCommand` classes represent BQL commands. The system uses the following types of `BqlCommand` classes:

- `BqlCommand`: This base class for the `Select` and `Search` classes is used by the system during the processing of data queries on the data entry forms.
- `BqlGenericCommand`: This class, which is derived from the `BqlCommand` class, is used by the system during the processing of generic inquiries.
- `BqlSoapCommand`: This class, which is derived from the `BqlCommand` class, is used by the system during the processing of reports. For details on report processing, see [Display of Reports](#).

The main purpose of `BqlCommand` classes is to convert BQL commands to SQL text. The `BqlGenericCommand` and `BqlSoapCommand` classes provide additional methods for generic inquiry and report processing.

## Parameters in Traditional BQL Statements

If you need to specify values in a business query language (BQL) statement, you use BQL parameters, which are replaced with the needed values in the translation to SQL. For details, how BQL statements with parameters are translated to SQL, see [Translation of a BQL Command with Parameters to an SQL Query Tree](#).

In this topic, you can find the description of the BQL parameters and the difference between them.

## Current and Current2

The `Current` parameter, as well as the `Current2` parameter, inserts the field value of the `Current` object from the `PXCache` object in the SQL query. If the `Current` object from the `PXCache` object is `null`, the `Current` parameter retrieves the default value of the field, while the `Current2` parameter doesn't retrieve the default value and inserts `null`.

By using the `Current` or `Current2` parameter in the declaration of a data view, you can refer to another view to relate these data views to each other. A typical example is referencing the current master record on master-detail webpages. For details on how the `Current` and `Current2` parameters are used, see [To Relate Data Views to One Another](#).

## Required

The `Required` parameter inserts a specific value into the SQL query.

By using the `Required` parameters, you can pass values to the SQL query, as described in [To Pass a Field Value to the SQL Query](#) and [To Pass Multiple Field Values to the SQL Query](#).

## Optional and Optional2

The `Optional` parameter works similarly to `Current` (as well as the `Optional2` parameter works similarly to `Current2`) if you don't specify an explicit value for this parameter during BQL statement execution. However, you can also pass an explicit value of the parameter to the SQL query.

By using the `Optional` or `Optional2` parameters, you can pass the external presentations of the values to the SQL query, as described in [To Provide External Presentation of the Field Value to the SQL Query](#).

## Argument

The `Argument` parameter passes values from UI controls to the SQL query.

By using the `Argument` parameters, you can pass values to the data view delegates. For more information on how to use the `Argument` parameter, see [To Pass a Value from a UI Control to a Data View](#).

## Traditional BQL and SQL Equivalents

The traditional business query language (BQL) library defines the following SQL function equivalents. Note that the use of the BQL equivalents may slightly differ from the use of the corresponding SQL functions. For details on each BQL class, see the API Reference.

### Correspondence Between SQL and BQL

SQL	Traditional BQL
<b>Clauses</b>	
WHERE	Where
INNER JOIN	InnerJoin
LEFT JOIN	LeftJoin
RIGHT JOIN	RightJoin
FULL JOIN	FullJoin
CROSS JOIN	CrossJoin
ON	On, On2
ORDER BY	OrderBy
ASC	Asc

SQL	Traditional BQL
DESC	Desc
GROUP BY	Aggregate, GroupBy
HAVING	Having
<b>Aggregation Functions</b>	
AVG	Avg
SUM	Sum
MIN	Min
MAX	Max
COUNT	Count
<b>Functions</b>	
ISNULL	IsNull<Operand1, Operand2>
NULLIF	NullIf
ROUND	Round
SUBSTRING	Substring
CONCAT	Add
RTRIM	RTrim
REPLACE	Replace
DATEDIFF	DateDiff
CASE	Switch, Case
<b>Arithmetic Operations</b>	
(Operand1 + Operand2)	Add<Operand1, Operand2>
(Operand1 - Operand2)	Sub<Operand1, Operand2>
(Operand1 * Operand2)	Mult<Operand1, Operand2>
(Operand1 / Operand2)	Div<Operand1, Operand2>
-Operand	Minus<Operand>
<b>Comparisons</b>	
=	Equal
<>	NotEqual
>	Greater
<	Less
<=	LessEqual
>=	GreaterEqual
LIKE	Like
NOT LIKE	NotLike
BETWEEN	Between
NOT BETWEEN	NotBetween

SQL	Traditional BQL
IS NULL	IsNull
IS NOT NULL	IsNotNull
IN	In, In2, In3
NOT IN	NotIn, NotIn2
EXISTS	Exists
<b>Logical Operators</b>	
AND	And, And2
OR	Or, Or2
NOT	Not, Not2
<b>Constants</b>	
NULL	Null
Other constants	Now, Today, Tomorrow, True, False, Zero, StringEmpty, MaxDate
<b>Full-Text Search Functions</b>	
FREETEXTTABLE	FreeText
CONTAINSTABLE	Contains

## To Select Records By Using Traditional BQL

To select records from the database, you can construct a business query language (BQL) statement. To construct a BQL statement, you use one of the generic `PXSelect` classes. You select the needed `PXSelect` class depending on the statement you need to compose, as described in the sections of this topic.



: In a `PXSelect` class, you configure a query to the database. The actual request to the database is performed once you cast the result of the query execution to a DAC or an array of DACs, or iterate through DACs in the result with the `foreach` statement. For details, see [Data Query Execution](#).

### Before You Proceed

Make sure that the application database has the database tables from which you are going to request data, and that the application defines the data access classes (DACs) for these tables. For more information on defining DACs, see [Data Access Classes in Traditional BQL](#).

### To Select All Records from a Database Table

To select all data from one database table without applying any filtering conditions or ordering, use one of the `PXSelect` classes that has DAC as the only type parameter, such as the `PXSelect<Table>` or `PXSelectReadOnly<Table>` class, as shown in the following sample BQL statement.

```
PXSelect<Product>
```

In this BQL statement, you are selecting all data records (with the values of all bound fields) from the `Product` table.



: For example, suppose that the `Product` table has two columns, `ProductID` and `UnitPrice`. In this case, Acumatica Framework translates the previous BQL statement to the following SQL query. The

framework adds ordering by the DAC key field (in ascending order) to the end of the SQL query because the BQL statement does not specify ordering.

```
SELECT Product.ProductID, Product.UnitPrice FROM Product
ORDERBY Product.ProductID
```

## To Filter Records

To filter records in the database table to be retrieved, construct a BQL statement with conditions by doing the following:

1. Use one of the `PXSelect` classes that has the `Where` type parameter, such as `PXSelect<Table, Where>`. For the full list of `PXSelect` classes, see [PXSelect Classes](#).
2. Specify the filtering conditions by using the `Where` clause, as described in [To Filter Records](#).
3. To specify the fields that should be used for filtering, use the class fields defined in the DACs, such as `Product.productID`. (The name of the field class starts with a lowercase letter. Do not confuse it with the property field, which has the same name but starts with uppercase letter.)

The following sample BQL statement selects all data records from the `Product` table that have the specified value in the `ProductID` column.

```
PXSelect<Product,
  Where<Product.productID, Equal<Required<Product.productID>>>>
```



: Acumatica Framework translates the previous BQL statement to the following SQL query. In this SQL query, `[list of columns]` is the list of columns of the `Product` table; `[parameter]` is the value passed to the `Select()` method of the `PXSelect` class, which is called when the BQL query is executed.

```
SET @P0 = [parameter];

SELECT [list of columns] FROM Product
WHERE Product.ProductID = @P0
ORDERBY Product.ProductID
```

Acumatica Framework explicitly enumerates the columns of the database table in the SQL query. For details on which columns are enumerated, see [Translation of a BQL Command to SQL](#).

## To Order Records

To order records in the database table to be retrieved, construct a BQL statement with ordering by doing the following:

1. Use one of the `PXSelect` classes that has the `OrderBy` type parameter, such as `PXSelectOrderBy<Table, OrderBy>` or `PXSelect<Table, Where, OrderBy>`. For the full list of `PXSelect` classes, see [PXSelect Classes](#).
2. Use the `OrderBy` clause to order records, as described in [To Order Records](#).
3. To specify the field that should be used for filtering, use the class field defined in the DAC, such as `Product.productID`. (The name of the field class starts with a lowercase letter. Do not confuse it with the property field, which has the same name but starts with uppercase letter.)

The following sample BQL statement selects all `Product` data records and sorts them by the `UnitPrice` field in ascending order.

```
PXSelectOrderBy<Product, OrderBy<Asc<Product.unitPrice>>>
```



: Acumatica Framework translates the previous BQL statement to the following SQL query, where `[list of columns]` is the list of columns of the `Product` table.

```
SELECT [list of columns] FROM Product
```



```
ORDER BY Product.UnitPrice
```

Acumatica Framework explicitly enumerates the columns of the database table in the SQL query. For details on which columns are enumerated, see [Translation of a BQL Command to SQL](#).

## To Query Multiple Tables

To join multiple tables, construct a BQL statement by doing the following:

1. Use one of the `PXSelect` classes that has the `Join` type parameter, such as `PXSelectJoin<Table, Join>` or `PXSelectReadOnly2<Table, Join>`.
2. In the `Join` type parameter of the `PXSelect` class, use one of the `Join` clauses—such as `InnerJoin`, `LeftJoin`, `RightJoin`, `FullJoin`, or `CrossJoin`—that are directly mapped to SQL `JOIN` clauses, as shown in the following sample BQL statement. For more information on the use of `Join` clauses, see [To Query Multiple Tables](#).

```
PXSelectJoin<SalesOrder,
    InnerJoin<OrderDetail,
        On<OrderDetail.orderNbr, Equal<SalesOrder.orderNbr>>>>
```



: Acumatica Framework translates the previous BQL statement to the following SQL query, where [list of columns] is the list of columns of the joined tables.

```
SELECT [list of columns] FROM SalesOrder
INNER JOIN OrderDetail
    ON OrderDetail.OrderNbr = SalesOrder.OrderNbr
```

Acumatica Framework explicitly enumerates the columns of the database table in the SQL query. For details on which columns are enumerated, see [Translation of a BQL Command to SQL](#).

## To Aggregate Records

To group or aggregate records, construct a BQL statement by doing the following:

1. Use one of the `PXSelect` classes with the `Aggregate` type parameter, such as `PXSelectGroupBy<Table, Aggregate>`.
2. In the `Aggregate` type parameter of the `PXSelect` class, specify the grouping conditions and aggregation functions by using the `Aggregate<Function>` class, the `GroupBy` clauses, and the `Min`, `Max`, `Sum`, `Avg`, and `Count` aggregation functions, as shown in the following sample BQL statement. For more information on the use of the grouping conditions and aggregation functions, see [To Group and Aggregate Records](#).

```
PXSelectGroupBy<Product,
    Aggregate<GroupBy<Product.categoryCD>>>
```



: Acumatica Framework translates the previous BQL statement to the following SQL query.

```
SELECT Product.CategoryCD,
    [MAX(Field) for other fields]
FROM Product
GROUP BY Product.CategoryCD
```

## To Filter Records

You construct business query language (BQL) statements with filtering conditions by using the `Where` clause in a `PXSelect` class that has the `Where` type parameter. (For more information on selecting the `PXSelect` class, see [To Select Records By Using Traditional BQL](#).) One `Where` clause can contain multiple conditions chained to one another by logical operators (`Or`, `And`, and `Not`) and nested `Where` clauses (these nested clauses are equivalent to placing conditions in brackets).

Typically, you construct a BQL statement with a condition to compare one field with another field or a constant, or to check if the field value has been specified (that is, to compare the field value with null). You can also use multiple conditions in the `Where` clause.

### To Compare a Field with Another Field

To compare one field with another field in the `Where` clause, do the following:

1. Select the comparison class that you need, such as `NotEqual`, `Greater`, or `Less`.
2. Specify the compared field in the first type parameter of the `Where` class and the comparison in the second type parameter, as shown in the following sample BQL statement.

```
PXSelect<Product, Where<Product.bookedQty, Greater<Product.availQty>>>
```



: Acumatica Framework translates the previous BQL statement to the following SQL query, where [list of columns] is the list of columns of the `Product` table.

```
SELECT [list of columns] FROM Product WHERE Product.BookedQty > Product.AvailQty
```

Acumatica Framework explicitly enumerates the columns of the database table in the SQL query. For details on which columns are enumerated, see [Translation of a BQL Command to SQL](#).

### To Compare a Field with a Constant

To compare a field with a constant in the `Where` clause, do the following:

1. Select the comparison class that you need, such as `NotEqual`, `Greater`, or `Less`.
2. Select one of the predefined constants—that is, the BQL class derived from the `Constant<Type>` class (such as Boolean values `True` and `False`, integer `Zero`, datetime `Now`, `Today`, and `MaxDate`, and string `StringEmpty`), or define your own constant as a class derived from the `Constant<Type>` class.
3. Specify the compared field in the first type parameter of the `Where` class and the comparison in the second type parameter, as shown in the following sample BQL statement.

```
PXSelect<Product, Where<Product.active, Equal<True>>>
```



: Acumatica Framework translates this BQL statement to the following SQL query, where [list of columns] is the list of columns of the `Product` table.

```
SELECT [list of columns] FROM Product WHERE Product.Active = CONVERT(BIT, 1)
```

Acumatica Framework explicitly enumerates the columns of the database table in the SQL query. For details on which columns are enumerated, see [Translation of a BQL Command to SQL](#).

### To Compare the Field Value with Null

To check whether a field value is specified, you compare the field value with null in one of the following ways:

- To check that the field is null, use the `Where<Operand, Comparison>` class, as shown in the following sample BQL statement.

```
PXSelect<Product, Where<Product.bookedQty, IsNull>>
```



: Acumatica Framework translates the previous BQL statement to the following SQL query, where [list of columns] is the list of columns of the `Product` table.

```
SELECT [list of columns] FROM Product WHERE Product.BookedQty IS NULL
```

Acumatica Framework explicitly enumerates the columns of the database table in the SQL query. For details on which columns are enumerated, see [Translation of a BQL Command to SQL](#).

- To check that the field is not null, do one of the following:
  - Use the `Where<Operator>` class and the logical operator `Not`, as shown in the following sample BQL statement.

```
PXSelect<Product, Where<Not<Product.bookedQty, IsNull>>>
```



: Acumatica Framework translates the previous BQL statement to the following SQL query, where [list of columns] is the list of columns of the `Product` table.

```
SELECT [list of columns] FROM Product WHERE NOT (Product.BookedQty IS NULL)
```

Acumatica Framework explicitly enumerates the columns of the database table in the SQL query. For details on which columns are enumerated, see [Translation of a BQL Command to SQL](#).

- Use the `Where<Operand, Comparison>` class, as shown in the following sample BQL statement..

```
PXSelect<Product, Where<Product.bookedQty, IsNotNull>>
```



: Acumatica Framework translates the previous BQL statement to the following SQL query, where [list of columns] is the list of columns of the `Product` table.

```
SELECT [list of columns] FROM Product WHERE Product.BookedQty IS NOT NULL
```

Acumatica Framework explicitly enumerates the columns of the database table in the SQL query. For details on which columns are enumerated, see [Translation of a BQL Command to SQL](#).



**Important:** The predefined constant `Null` cannot be used in the `Where` clause with `Equal` to select records with null fields. The `Null` constant is used in `Switch` conditions.

## To Use Multiple Conditions in One Where Clause

To specify multiple comparisons in one `Where` clause, do one of the following:

- To specify multiple comparisons that are connected with the same logical operator, use the `Where<Operand, Comparison, NextOperator>` class and specify its type parameters as follows:
  - In the first type parameter, specify the first compared field.
  - In the second type parameter, specify the first comparison, such as `NotEqual`, `Greater`, or `Less`.
  - In the third type parameter, specify the logical operator, such as `And`, `And2`, `Or`, or `Or2`. You can chain any number of comparisons to one another by using binary operators with three type parameters, as shown in the following sample BQL statement.

```
PXSelect<Product,
    Where<Product.bookedQty, Greater<Product.availQty>,
        Or<Product.availQty, Less<Product.minAvailQty>,
        Or<Product.availQty, IsNull>>>>
```



: Acumatica Framework translates the previous BQL statement to the following SQL query, where [list of columns] is the list of columns of the `Product` table.

```
SELECT [list of columns] FROM Product
WHERE Product.BookedQty > Product.AvailQty
```

```
OR Product.AvailQty < Product.MinAvailQty
OR Product.AvailQty IsNull
```

Acumatica Framework explicitly enumerates the columns of the database table in the SQL query. For details on which columns are enumerated, see [Translation of a BQL Command to SQL](#).

- To write more complex conditional expressions with logical operators of different types, use nested `Where` or `Where2` clauses. For more information on writing complex BQL statements, see [To Compose a BQL Statement from an SQL Statement](#).

## To Order Records

You construct business query language (BQL) statements that include ordering of records by using the `OrderBy` clause in one of the `PXSelect` classes that has the `OrderBy` type parameter. (For more information on selecting the `PXSelect` class, see [To Select Records By Using Traditional BQL](#).)

By default, if the BQL statement does not specify ordering, Acumatica Framework adds ordering by the data access class (DAC) key fields (in the order of field declaration) in ascending order to the end of the SQL query. You can order the records by one column or multiple columns, or by a condition.

### To Order Records by One Column

To order records by one column in ascending or descending order, use the `OrderBy` class and the `Asc<Field>` or `Desc<Field>` class, as shown in the following sample BQL statement.

```
PXSelectOrderBy<Product, OrderBy<Asc<Product.unitPrice>>>
```

In this statement, all `Product` data records are selected and are sorted by the `UnitPrice` field in ascending order.



: Acumatica Framework translates this BQL statement to the following SQL query, where [list of columns] is the list of columns of the `Product` table.

```
SELECT [list of columns] FROM Product
ORDER BY Product.UnitPrice
```

Acumatica Framework explicitly enumerates the columns of the database table in the SQL query. For details on which columns are enumerated, see [Translation of a BQL Command to SQL](#).

### To Order Records by Multiple Columns

To order records by multiple columns, use the `OrderBy` class and the `Asc<Field, NextField>` or `Desc<Field, NextField>` class, as shown in the following sample BQL statement.

```
PXSelectOrderBy<Product,
    OrderBy<Asc<Product.unitPrice, Desc<Product.availQty>>>>
```



: Acumatica Framework translates the previous BQL statement to the following SQL query, where [list of columns] is the list of columns of the `Product` table.

```
SELECT [list of columns] FROM Product
ORDER BY Product.UnitPrice, Product.AvailQty DESC
```

Acumatica Framework explicitly enumerates the columns of the database table in the SQL query. For details on which columns are enumerated, see [Translation of a BQL Command to SQL](#).

## To Order Records by a Condition

To sort data records according to a condition, put the `Switch` clause inside `Asc` or `Desc` in `OrderBy`, as shown in the following sample BQL statement.

```
PXSelectOrderBy<Product,
  OrderBy<Asc<
    Switch<Case<Where<Product.availQty, Greater<Product.bookedQty>>, True>,
      False>>>>
```

In this statement, the records with `AvailQty` values less or equal to `BookedQty` values are ordered first.



: Acumatica Framework translates the previous BQL statement to the following SQL query, where [list of columns] is the list of columns of the `Product` table.

```
SELECT [list of columns] FROM Product
ORDER BY
  ( CASE
    WHEN Product.AvailQty > Product.BookedQty THEN 1
    ELSE 0
  END )
```

Acumatica Framework explicitly enumerates the columns of the database table in the SQL query. For details on which columns are enumerated, see [Translation of a BQL Command to SQL](#).

## To Query Multiple Tables

You construct business query language (BQL) statements that join multiple tables by using one of the `Join` clauses in one of the `PXSelect` classes that has the `Join` type parameter. (For more information on selecting the `PXSelect` class, see [To Select Records By Using Traditional BQL](#).) In BQL statements, you can join multiple database tables by using the following clauses directly mapped to SQL `JOIN` clauses:

- `InnerJoin` returns all records where there is at least one match in both tables.
- `LeftJoin` returns all records from the left table, and the matched records from the right table. Where there are no matched records from the right table, null values are inserted.
- `RightJoin` returns all records from the right table, and the matched records from the left table. Where there are no matched records from the left table, null values are inserted.
- `FullJoin` returns all records when there is a match in one of the tables.
- `CrossJoin` returns the entire Cartesian product of the two tables.

### To Join Two Tables (Inner Join, Left Join, Right Join, or Full Join)

To join two tables, use one of the `Join` clauses with two type parameters (such as `InnerJoin<Table, On>`) and the `On<Operand, Comparison>` or `On<Operator>` class to specify a conditional expression for joining, as shown in the following sample BQL statement.

```
PXSelectJoin<SalesOrder,
  InnerJoin<OrderDetail,
    On<OrderDetail.orderNbr, Equal<SalesOrder.orderNbr>>>>
```



: Acumatica Framework translates the previous BQL statement to the following SQL query, where [list of columns] is the list of columns of the joined tables.

```
SELECT [list of columns] FROM SalesOrder
INNER JOIN OrderDetail
  ON OrderDetail.OrderNbr = SalesOrder.OrderNbr
```

Acumatica Framework explicitly enumerates the columns of the database table in the SQL query. For details on which columns are enumerated, see [Translation of a BQL Command to SQL](#).

## To Cross-Join Two Tables

To cross-join two tables, use the `CrossJoin<Table>` class, as shown in the following sample BQL statement.

```
PXSelectJoin<Product, CrossJoin<Supplier>>
```



: Acumatica Framework translates the previous BQL statement to the following SQL query, where [list of columns] is the list of columns of the joined tables.

```
SELECT [list of columns] FROM Product CROSS JOIN Supplier
```

Acumatica Framework explicitly enumerates the columns of the database table in the SQL query. For details on which columns are enumerated, see [Translation of a BQL Command to SQL](#).

## To Join Multiple Tables

To specify multiple join clauses, use the following instructions:

- Use a `Join` clause with three type parameters (such as `InnerJoin<Table, On, NextJoin>`). Each subsequent join clause is specified as the last type parameter of the previous join clause, as shown in the following sample BQL statement.

```
PXSelectJoin<SalesOrder,
    InnerJoin<OrderDetail,
        On<OrderDetail.orderNbr, Equal<SalesOrder.orderNbr>>,
    LeftJoin<Employee,
        On<Employee.employeeID, Equal<SalesOrder.employeeID>>>>>
```



: Acumatica Framework translates this BQL statement to the following SQL query, where [list of columns] is the list of columns of the joined tables.

```
SELECT [list of columns] FROM SalesOrder
INNER JOIN OrderDetail
    ON OrderDetail.OrderNbr = SalesOrder.OrderNbr
LEFT JOIN Employee
    ON Employee.EmployeeID = SalesOrder.EmployeeID
```

Acumatica Framework explicitly enumerates the columns of the database table in the SQL query. For details on which columns are enumerated, see [Translation of a BQL Command to SQL](#).

- Use the `On` conditions to specify conditional expressions for joining, as shown in the following sample BQL statement. In subsequent join clauses, the `On` conditions can refer to fields from any joined table, and can contain any number of conditions chained by logical operators as in filtering conditions.

```
PXSelectJoin<SalesOrder,
    InnerJoin<OrderDetail,
        On<OrderDetail.orderNbr, Equal<SalesOrder.orderNbr>>,
    LeftJoin<Employee,
        On<Employee.employeeID, Equal<SalesOrder.employeeID>>,
    RightJoin<Product,
        On<Product.productID, Equal<OrderDetail.productID>,
        And<Product.unitPrice, Equal<OrderDetail.unitPrice>>>>>>>
```



: Acumatica Framework translates the previous BQL statement to the following SQL query, where [list of columns] is the list of columns of the joined tables.

```
SELECT [list of columns] FROM SalesOrder
INNER JOIN OrderDetail
    ON OrderDetail.OrderNbr = SalesOrder.OrderNbr
LEFT JOIN Employee
    ON Employee.EmployeeID = SalesOrder.EmployeeID
RIGHT JOIN Product
    ON (Product.ProductID = OrderDetail.ProductID AND
```

```
Product.UnitPrice = OrderDetail.UnitPrice)
```

Acumatica Framework explicitly enumerates the columns of the database table in the SQL query. For details on which columns are enumerated, see [Translation of a BQL Command to SQL](#).

## To Group and Aggregate Records

You construct business query language (BQL) statements that group and aggregate records by using the `Aggregate` clause in one of the `PXSelect` classes that has the `Aggregate` type parameter. (For more information on selecting the `PXSelect` class, see [To Select Records By Using Traditional BQL](#).)

### To Group and Aggregate Records

To group and aggregate records, follow the instructions below:

1. Specify all grouping conditions (the `GroupBy` clause) and aggregation functions (such as `Min`, `Max`, `Sum`, `Avg`, and `Count`) in the `Aggregate` clause, as shown in the following sample BQL statement. Fields specified in `GroupBy` clauses are selected as is; an aggregation function is applied to all other fields. The default `Max` function is used if no function is specified for a field. If a data field has the `PXDBScalar` attribute, `NULL` is inserted for that field.

```
PXSelectGroupBy<Product,
    Aggregate<GroupBy<Product.categoryCD>>>
```



: Acumatica Framework translates the previous BQL statement to the following SQL query.

```
SELECT Product.CategoryCD,
       [MAX(Field) for other fields]
FROM Product
GROUP BY Product.CategoryCD
```

2. If necessary, insert another `GroupBy` clause or aggregation function as the second type parameter of the previous `GroupBy` clause or aggregation function, as shown in the following sample BQL statement.

```
PXSelectGroupBy<Product,
    Aggregate<GroupBy<Product.categoryCD,
        Sum<Product.availQty,
        Sum<Product.bookedQty,
        GroupBy<Product.stockUnit,
        Min<Product.unitPrice>>>>>>
```



: Acumatica Framework translates the previous BQL statement to the following SQL query.

```
SELECT Product.CategoryCD, Product.StockUnit,
       SUM(Product.AvailQty), SUM(Product.AvailQty), MIN(Product.UnitPrice),
       [MAX(Field) for other fields]
FROM Product
GROUP BY Product.CategoryCD, Product.StockUnit
```

## To Use Parameters in Traditional BQL

You may need to use BQL parameters if you need to relate data views to each other, to pass field values to the SQL query, to pass the external presentations of the values to the SQL query, or to pass values from UI controls to the SQL query. For more information on the BQL parameters, see [Parameters in Traditional BQL Statements](#).

### To Relate Data Views to One Another

To relate data views to one another, use the `Current` parameter, as shown in the following sample code.

```
// The view declarations in a graph
PXSelect<Document> Documents;
```

```
PXSelect<DocTransaction,
    Where<DocTransaction.docNbr, Equal<Current<Document.docNbr>>,
        And<DocTransaction.docType, Equal<Current<Document.docType>>>>>
    DocTransactions;
```

In this code, there is a many-to-one relationship between the `DocTransaction` and `Document` data access classes (DACs), and this relationship is implemented through the `DocNbr` and `DocType` key fields. The views in the code connect the `Document` and `DocTransaction` records.



: Acumatica Framework translates the BQL query of the second view in the sample BQL code to the following SQL statement. In this SQL query, `[parameter1]` is the `DocNbr` value and `[parameter2]` is the `DocType` value taken from the `Current` property of the `DocTransaction` cache; `[list of columns]` is the list of columns of the `DocTransaction` table.

```
SET @P0 = [parameter1]
SET @P1 = [parameter2]

SELECT * FROM DocTransaction
WHERE DocTransaction.DocNbr = @P0
AND DocTransaction.DocType = @P1
```

Acumatica Framework explicitly enumerates the columns of the database table in the SQL query. For details on which columns are enumerated, see [Translation of a BQL Command to SQL](#).

### To Pass a Field Value to the SQL Query

To pass a specific value to the SQL query, use the `Required` parameter in the BQL statement and specify the needed value as the `Select()` method argument. The value passed to `Select()` must be of the same type as the type of the specified field.



: The `Required` parameter should be used only in the BQL statements that are directly executed in the application code. The data views that are queried from the UI will not work if they contain `Required` parameters.

The code below shows the execution of a BQL statement with the `Required` parameter.

```
// Suppose an event handler related to the Product DAC
// is being executed
Product product = (Product)e.Row;

// Select the Category record with the specified CategoryCD
Category category =
    PXSelect<Category,
        Where<Category.categoryCD, Equal<Required<Category.categoryCD>>>>
        .Select(this, product.CategoryCD);
```



: Acumatica Framework translates the previous BQL query to the following SQL statement. In this SQL query, `[parameter]` is the value of the `product.CategoryCD` variable at the moment the `Select()` method is invoked; `[list of columns]` is the list of columns of the `Category` table.

```
SET @P0 = [parameter]

SELECT * FROM Category
WHERE Category.CategoryCD = @P0
```

Acumatica Framework explicitly enumerates the columns of the database table in the SQL query. For details on which columns are enumerated, see [Translation of a BQL Command to SQL](#).

### To Pass Multiple Field Values to the SQL Query

To pass multiple values to the SQL query, use multiple `Required` parameters in the BQL statement and specify the needed values as the `Select()` method arguments in the order in which the parameters are specified in the BQL statement. The number of `Required` parameters must match the number of parameters passed to the `Select()` function.





: The Required parameters should be used in only the BQL statements that are executed in the application code.

The following code shows an example of a BQL statement with two Required parameters.

```
// Suppose an event handler related to the DocTransaction DAC
// is being executed
DocTransaction line = (DocTransaction)e.Row;
...
Document doc =
    PXSelect<Document,
        Where<Document.docNbr, Equal<Required<DocTransaction.docNbr>,
            And<Document.docType, Equal<Required<DocTransaction.docType>>>>>
    .Select(this, line.DocNbr, line.DocType);
```



: Acumatica Framework translates the previous BQL query to the following SQL statement, where [list of columns] is the list of columns of the Document table.

```
SET @P0 = [line.DocNbr value]
SET @P1 = [line.DocType value]

SELECT * FROM Document
WHERE Document.DocNbr = @P0
      AND Document.DocType = @P1
```

Acumatica Framework explicitly enumerates the columns of the database table in the SQL query. For details on which columns are enumerated, see [Translation of a BQL Command to SQL](#).

### To Provide External Presentation of the Field Value to the SQL Query

To substitute the value in the SQL query, do the following:

1. Add the `PXSelector` attribute with a substitute key to a DAC field, as shown in the following example.

```
[PXSelector(typeof(Search<Product.productID>),
    new Type [] {
        typeof(Product.productCD),
        typeof(Product.productName)
    },
    SubstituteKey = typeof(Product.productCD))]
public virtual int? ProductID { get; set; }
```

In this example, `PXSelector` replaces the `ProductID` field in the user interface with the human-readable `ProductCD` field. In the UI control for this field, the user enters a `ProductCD` value. The `PXSelector` attribute implements the `FieldUpdating` event handler, which replaces the `ProductCD` value with the corresponding `ProductID` value.

2. Use the `Optional` parameter to select records by the external presentation of the field value, as shown in the following code for `OrderDetail` records.

```
// Product data record obtained
Product p = ...
// Selecting OrderDetail records: ProductCD value is passed
// to the Select() method.
PXSelect<OrderDetail,
    Where<OrderDetail.ProductID, Equal<Optional<OrderDetail.ProductID>>>>
    .Select(this, p.ProductCD);
```

3. In the `Select()` method, provide values for all `Optional`, `Required`, and `Argument` parameters up to the last `Required` or `Argument` parameter in the BQL statement, as shown in the following sample code.

```
// Related OrderDetail and Product records obtained
```

```

OrderDetail od = ...
Product p = ...

// At least three values (in addition to graph reference) must
// be passed to the Select() method below.
// The second Optional parameter here will be replaced with the
// default UnitPrice value.
PXResultset<OrderDetail> details =
    PXSelect<OrderDetail,
        Where<OrderDetail.productID, Equal<Optional<OrderDetail.productID>>,
            And<OrderDetail.extPrice, Less<Required<OrderDetail.extPrice>>,
            And<OrderDetail.unitPrice, Greater<Required<OrderDetail.unitPrice>>,
            And<OrderDetail.taxRate, Equal<Optional<OrderDetail.taxRate>>>>>>
        .Select(this, p.ProductCD, od.ExtPrice, od.UnitPrice);

```



: Acumatica Framework translates the BQL query in the code to the following SQL statement, where [list of columns] is the list of columns of the OrderDetail table.

```

SET @P0 = [line.ProductID value or default]
SET @P1 = [line.ExtPrice value]
SET @P2 = [line.UnitPrice value]
SET @P3 = [Default TaxRate value]

SELECT [list of columns] FROM OrderDetail
WHERE OrderDetail.ProductID = @P0
      AND OrderDetail.ExtPrice < @P1
      AND OrderDetail.UnitPrice > @P2
      AND OrderDetail.TaxRate = @P3

```

Acumatica Framework explicitly enumerates the columns of the database table in the SQL query. For details on which columns are enumerated, see [Translation of a BQL Command to SQL](#).

## To Pass a Value from a UI Control to a Data View

To pass a value from a UI control to a data view, do the following:

1. Define the `PXSelect` data view with the `Argument` parameter whose type parameter specifies the data type of the expected value, as shown in the following sample BQL query.

```

PXSelect<TreeViewItem,
    Where<TreeViewItem.parentID, Equal<Argument<int?>>>,
    OrderBy<Asc<TreeViewItem.parentID>>> GridDataSource;

```



: Acumatica Framework translates the previous BQL query to the following SQL statement. In this SQL query, [parameter] will contain the value received from the UI control and passed to the `Select()` method; [list of columns] is the list of columns of the TreeViewItem table.

```

SET @P0 = [parameter]

SELECT [list of columns] FROM TreeViewItem
WHERE TreeViewItem.ParentID = @P0
ORDER BY TreeViewItem.ParentID

```

Acumatica Framework explicitly enumerates the columns of the database table in the SQL query. For details on which columns are enumerated, see [Translation of a BQL Command to SQL](#).

2. Define the data view delegate with the parameters through which you can access the values passed from the UI. (You can find more information on how to use data view delegates in [#unique\\_124/unique\\_124\\_Connect\\_42\\_optional\\_method.](#))



: When a BQL statement with the `Argument` parameter is executed in code, the value must be specified in the parameters of the `Select()` method.

## To Use Arithmetic Operations

Arithmetic operations—such as `Add<Operand1, Operand2>`, `Sub<Operand1, Operand2>`, `Mult<Operand1, Operand2>`, `Div<Operand1, Operand2>`, and `Minus<Operand>`—are used primarily in attributes to calculate the value of a field from other fields. Arithmetic operations can also be used as operands in `Where` and `OrderBy` clauses in business query language (BQL) statements.

### To Use Arithmetic Operations in Attributes

To calculate an expression in an attribute, do the following:

1. Compose the expression by using arithmetic operations. For example, you can calculate product reorder discrepancy by using the following BQL expression, where the `decimal_0` constant represents the 0 decimal value. `IsNull` returns the first argument if it is not null or the second argument otherwise.

```
Minus<
    Sub<Sub<IsNull<Product.availQty, decimal_0>,
        IsNull<Product.bookedQty, decimal_0>>,
        Product.minAvailQty>>
```



: Acumatica Framework translates the previous BQL statement to the following SQL query.

```
-( (ISNULL(Product.AvailQty, .0) - ISNULL(Product.BookedQty, .0))
  - Product.MinAvailQty)
```

2. Use the calculated expression in an attribute (such as `PXDBCalced`) to define a calculated field that is not bound to a database column, as shown in the following sample code.

```
// Data field definition in a DAC
[PXDecimal(2)]
[PXDBCalced(typeof(Minus<
    Sub<Sub<IsNull<Product.availQty, decimal_0>,
        IsNull<Product.bookedQty, decimal_0>>,
        Product.minAvailQty>>),
    typeof(Decimal))]]
public virtual decimal? Discrepancy { get; set; }
```

### To Use Arithmetic Operations in BQL Statements

To use arithmetic operations in a conditional expression in a BQL statement, do the following:

1. Compose the expression by using arithmetic operations. For example, you can calculate product reorder discrepancy by using the following BQL expression, where the `decimal_0` constant represents the 0 decimal value. `IsNull` returns the first argument if it is not null or the second argument otherwise.

```
Minus<
    Sub<Sub<IsNull<Product.availQty, decimal_0>,
        IsNull<Product.bookedQty, decimal_0>>,
        Product.minAvailQty>>
```



: Acumatica Framework translates the previous BQL statement to the following SQL query.

```
-( (ISNULL(Product.AvailQty, .0) - ISNULL(Product.BookedQty, .0))
  - Product.MinAvailQty)
```

2. Use the calculated expression in a BQL statement, as shown in the following example.

```
PXSelect<Product,
    Where<Minus<
        Sub<Sub<IsNull<Product.availQty, decimal_0>,
            IsNull<Product.bookedQty, decimal_0>>,
```

```
Product.minAvailQty>>,
NotEqual<decimal_0>>>
```



: Acumatica Framework translates the previous BQL statement to the following SQL query, where [list of columns] is the list of columns of the Product table.

```
SELECT [list of columns] FROM Product
WHERE -((ISNULL(Product.AvailQty, .0) - ISNULL(Product.BookedQty, .0))
- Product.MinAvailQty) <> .0
```

Acumatica Framework explicitly enumerates the columns of the database table in the SQL query. For details on which columns are enumerated, see [Translation of a BQL Command to SQL](#).

## To Compose a BQL Statement from an SQL Statement

If you are familiar with the construction of SQL statements, you may want to first construct an SQL statement and then translate it to business query language (BQL). You can follow the instructions described in this topic to translate SQL statements to BQL statements.

### To Translate an SQL Statement to BQL

To translate an SQL statement to BQL, do the following:

1. Construct an SQL statement that selects the data you need.

For example, suppose that you need to convert to BQL the following SQL statement. In this SQL query, we use the \* sign to indicate that all columns of the Product table should be selected.

```
SELECT * FROM Product
INNER JOIN SupplierProduct
    ON SupplierProduct.ProductID = Product.ProductID
INNER JOIN Supplier
    ON Supplier.AccountID = SupplierProduct.AccountID
WHERE (Product.BookedQty IS NOT NULL
    AND Product.AvailQty IS NOT NULL
    AND Product.MinAvailQty IS NOT NULL
    AND (Product.Active = 1
        OR Product.Active IS NULL)
    AND (Product.BookedQty > Product.AvailQty
        OR Product.AvailQty < Product.MinAvailQty))
    OR Product.AvailQty IS NOT NULL
ORDER BY Product.UnitPrice, Product.AvailQty DESC
```

2. Replace the names of columns with the names of class fields that correspond to the columns in data access classes (DACs). That is, change the uppercase letter in the name of each column to the lowercase, as shown in the following sample code. In this sample code, the changes are shown in bold type.

```
SELECT * FROM Product
INNER JOIN SupplierProduct
    ON SupplierProduct.productID = Product.productID
INNER JOIN Supplier
    ON Supplier.accountID = SupplierProduct.accountID
WHERE (Product.bookedQty IS NOT NULL
    AND Product.availQty IS NOT NULL
    AND Product.minAvailQty IS NOT NULL
    AND (Product.active = 1
        OR Product.active IS NULL)
    AND (Product.bookedQty > Product.availQty
        OR Product.availQty < Product.minAvailQty))
    OR Product.availQty IS NOT NULL
ORDER BY Product.unitPrice, Product.availQty DESC
```

3. If your SQL statement contains constants, replace it with either one of the predefined constants or your own constant. (For details on using constants, see [To Compare a Field with a Constant](#) and the description of the `Constant<ConstType>` class.) If you need to change the values of the constants at runtime, replace the constants with parameters, as described in [To Use Parameters in Traditional BQL](#).
4. Find the JOIN, WHERE, GROUP BY, and ORDER BY clauses that you have in the SQL statement. Depending on the included clauses, select one of the `PXSelect` classes, and replace `SELECT * FROM` with this class in your SQL statement. For details on selection of the `PXSelect` class, see [To Select Records By Using Traditional BQL](#). For the list of all `PXSelect` classes, see [PXSelect Classes](#).

In the sample code that has been presented in this topic, you would use the `PXSelectJoin<Table, Join, Where, OrderBy>` class, and you would change the sample code as follows. (The changes are shown in bold type.)

```
PXSelectJoin<Product,
INNER JOIN SupplierProduct
    ON SupplierProduct.productID = Product.productID
INNER JOIN Supplier
    ON Supplier.accountID = SupplierProduct.accountID,
WHERE (Product.bookedQty IS NOT NULL
      AND Product.availQty IS NOT NULL
      AND Product.minAvailQty IS NOT NULL
      AND(Product.active = 1
          OR Product.active IS NULL)
      AND(Product.bookedQty > Product.availQty
          OR Product.availQty < Product.minAvailQty))
      OR Product.availQty IS NOT NULL,
ORDER BY Product.unitPrice, Product.availQty DESC>
```

5. If your SQL statement includes JOIN clauses, do the following:
  - a. Replace the last JOIN clause with the corresponding BQL `Join` clause. You would change the sample code of this topic as follows. (The changes are shown in bold type.)

```
PXSelectJoin<Product,
INNER JOIN SupplierProduct
    ON SupplierProduct.productID = Product.productID
InnerJoin<Supplier,
    ON Supplier.accountID = SupplierProduct.accountID>,
WHERE (Product.bookedQty IS NOT NULL
      AND Product.availQty IS NOT NULL
      AND Product.minAvailQty IS NOT NULL
      AND(Product.active = 1
          OR Product.active IS NULL)
      AND(Product.bookedQty > Product.availQty
          OR Product.availQty < Product.minAvailQty))
      OR Product.availQty IS NOT NULL,
ORDER BY Product.unitPrice, Product.availQty DESC>
```

- b. Chain other JOIN clauses to one another, as described in [To Query Multiple Tables](#). You would change the sample code of this topic as follows. (The changes are shown in bold type.)

```
PXSelectJoin<Product,
InnerJoin<SupplierProduct,
    ON SupplierProduct.productID = Product.productID,
InnerJoin<Supplier,
    ON Supplier.accountID = SupplierProduct.accountID>>,
WHERE (Product.bookedQty IS NOT NULL
      AND Product.availQty IS NOT NULL
      AND Product.minAvailQty IS NOT NULL
      AND(Product.active = 1
          OR Product.active IS NULL)
```

```

        AND(Product.bookedQty > Product.availQty
            OR Product.availQty < Product.minAvailQty))
    OR Product.availQty IS NOT NULL,
ORDER BY Product.unitPrice, Product.availQty DESC>

```

c. Replace each ON clause, as follows:

- For a single condition or groups that start with a simple condition, replace the ON clause with **On**.
- For groups that start with a group of conditions, replace the ON clause with **On2**.

With these replacements, the sample code used in this topic would be changed to the following code. (The changes are shown in bold type.)

```

PXSelectJoin<Product,
InnerJoin<SupplierProduct,
    On<SupplierProduct.productID = Product.productID>,
InnerJoin<Supplier,
    On<Supplier.accountID = SupplierProduct.accountID>>>,
WHERE (Product.bookedQty IS NOT NULL
    AND Product.availQty IS NOT NULL
    AND Product.minAvailQty IS NOT NULL
    AND(Product.active = 1
        OR Product.active IS NULL)
    AND(Product.bookedQty > Product.availQty
        OR Product.availQty < Product.minAvailQty))
    OR Product.availQty IS NOT NULL,
ORDER BY Product.unitPrice, Product.availQty DESC>

```

6. If your SQL statement includes a WHERE clause, replace the WHERE clause and each pair of parentheses that encloses each group of conditions in the WHERE clause with a **Where**, **Where2**, **Not**, or **Not2** clause, as follows:

- **Where** is used for groups that start with a simple condition.
- **Not** is used for groups that start with a simple condition but are preceded with the logical NOT.
- **Where2** is used for groups that start with a group of conditions.
- **Not2** is used for groups that start with a group of conditions but preceded with the logical NOT.

With these replacements, the sample code used in this topic would be changed to the following code. (The changes are shown in bold type.)

```

PXSelectJoin<Product,
InnerJoin<SupplierProduct,
    On<SupplierProduct.productID = Product.productID>,
InnerJoin<Supplier,
    On<Supplier.accountID = SupplierProduct.accountID>>>,
Where2<Where<Product.bookedQty, IS NOT NULL
    AND Product.availQty IS NOT NULL
    AND Product.minAvailQty IS NOT NULL
    AND Where<Product.active = 1,
        OR Product.active IS NULL>
    AND Where<Product.bookedQty > Product.availQty,
        OR Product.availQty < Product.minAvailQty>>,
    OR Product.availQty IS NOT NULL>,
OrderBy<Asc<Product.unitPrice, Desc<Product.availQty>>>>

```

7. In each BQL **Where** or **On** clause, replace the logical operators (either AND or OR) to **And**, **Or**, **And2**, or **Or2**, as follows:

- a. Replace the last AND or OR in each BQL Where or On clause with the And or Or operator, respectively, as shown in the following code. (The changes are shown in bold type.)

```
PXSelectJoin<Product,
  InnerJoin<SupplierProduct,
    On<SupplierProduct.productID = Product.productID>,
  InnerJoin<Supplier,
    On<Supplier.accountID = SupplierProduct.accountID>>>,
Where2<Where<Product.bookedQty, IS NOT NULL
  AND Product.availQty IS NOT NULL
  AND Product.minAvailQty IS NOT NULL
  AND Where<Product.active = 1,
    Or<Product.active IS NULL>>
  And<Where<Product.bookedQty > Product.availQty,
    Or<Product.availQty < Product.minAvailQty>>>>,
Or<Product.availQty IS NOT NULL>>,
OrderBy<Asc<Product.unitPrice, Desc<Product.availQty>>>>
```

- b. In each BQL Where or On clause, if the AND or OR is located before a simple condition, replace it with And or Or, respectively. If the condition is preceded by NOT, wrap it in Not. With these replacements, the sample code used in this topic would be changed to the following code. (The changes are shown in bold type.)

```
PXSelectJoin<Product,
  InnerJoin<SupplierProduct,
    On<SupplierProduct.productID = Product.productID>,
  InnerJoin<Supplier,
    On<Supplier.accountID = SupplierProduct.accountID>>>,
Where2<Where<Product.bookedQty, IS NOT NULL,
  And<Product.availQty IS NOT NULL,
  And<Product.minAvailQty IS NOT NULL,
  AND Where<Product.active = 1,
    Or<Product.active IS NULL>>
  And<Where<Product.bookedQty > Product.availQty,
    Or<Product.availQty < Product.minAvailQty>>>>,
Or<Product.availQty IS NOT NULL>>,
OrderBy<Asc<Product.unitPrice, Desc<Product.availQty>>>>
```

- c. In each BQL Where or On clause, if the AND or OR is located before a group of conditions, replace it with And2<Operator, NextOperator> or Or2<Operator, NextOperator>, respectively. The first parameter in a logical operator is Where (or Where2). If the condition is preceded by NOT, place Not before a group in a Where clause. The following sample code implements these changes (shown in bold type).

```
PXSelectJoin<Product,
  InnerJoin<SupplierProduct,
    On<SupplierProduct.productID = Product.productID>,
  InnerJoin<Supplier,
    On<Supplier.accountID = SupplierProduct.accountID>>>,
Where2<Where<Product.bookedQty, IS NOT NULL,
  And<Product.availQty IS NOT NULL,
  And<Product.minAvailQty IS NOT NULL,
  And2<Where<Product.active = 1,
    Or<Product.active IS NULL>>,
  And<Where<Product.bookedQty > Product.availQty,
    Or<Product.availQty < Product.minAvailQty>>>>,
Or<Product.availQty IS NOT NULL>>,
OrderBy<Asc<Product.unitPrice, Desc<Product.availQty>>>>
```

8. In each Where or On clause, replace the groups that use arithmetic operations with the corresponding BQL operators, as described in [To Use Arithmetic Operations](#).

9. In each `Where` or `On` clause, replace each comparison with the corresponding comparison operator, such as `Equal`, `Greater`, or `IsNull`. For more information on constructing comparisons, see [To Filter Records](#).

The following sample code includes these changes (shown in bold type).

```
PXSelectJoin<Product,
InnerJoin<SupplierProduct,
    On<SupplierProduct.productID, Equal<Product.productID>>Equal<SupplierProduct.accountID>>>>,
Where2<Where<Product.bookedQty, IsNullIsNullIsNullEqual<True>IsNullGreater<Product.availQty>Less<Product.minAvailQty>>>>>>>>,
    Or<Product.availQty, IsNull

```

10. Align logical operators of the same level so that they have the same indentation and so that each simple condition is placed on a separate line. Do not add line breaks before nested `Where` clauses.
11. If your SQL statement includes the `GROUP BY` clause, do the following:
- Replace the `GROUP BY` clause with the `Aggregate` clause.
  - Chain the `GroupBy` clause and aggregation functions (such as `Min`, `Max`, `Sum`, `Avg`, and `Count`) to one another as described in [To Group and Aggregate Records](#).
12. If your SQL statement includes the `ORDER BY` clause, do the following:
- Replace the `ORDER BY` clause with the `OrderBy` clause. The following sample code shows this change (with changes shown in bold type).

```
PXSelectJoin<Product,
InnerJoin<SupplierProduct,
    On<SupplierProduct.productID, Equal<Product.productID>>,
InnerJoin<Supplier,
    On<Supplier.accountID, Equal<SupplierProduct.accountID>>>>,
Where2<Where<Product.bookedQty, IsNotNull,
    And<Product.availQty, IsNotNull,
    And<Product.minAvailQty, IsNotNull,
    And2<Where<Product.active, Equal<True>,
        Or<Product.active, IsNull>>,
    And<Where<Product.bookedQty, Greater<Product.availQty>,
        Or<Product.availQty, Less<Product.minAvailQty>>>>>>>>,
    Or<Product.availQty, IsNotNull>>,
OrderBy<Product.unitPrice, Product.availQty DESC>>
```

- Chain the `Asc` and `Desc` operators to one another, as described in [To Order Records](#). The following sample code shows this change (with changes shown in bold type).

```
PXSelectJoin<Product,
InnerJoin<SupplierProduct,
    On<SupplierProduct.productID, Equal<Product.productID>>,
InnerJoin<Supplier,
    On<Supplier.accountID, Equal<SupplierProduct.accountID>>>>,
Where2<Where<Product.bookedQty, IsNotNull,
    And<Product.availQty, IsNotNull,
    And<Product.minAvailQty, IsNotNull,
    And2<Where<Product.active, Equal<True>,
        Or<Product.active, IsNull>>,
    And<Where<Product.bookedQty, Greater<Product.availQty>,
        Or<Product.availQty, Less<Product.minAvailQty>>>>>>>>
```



```
Or<Product.availQty, Less<Product.minAvailQty>>>>>>>>,
Or<Product.availQty, IsNotNull>>,
OrderBy<Asc<Product.unitPrice, Desc<Product.availQty>>>>
```

- 13.** Check that the final statement is correct by doing the following:
- Check that all lines except the last line of the BQL statement end with a comma.
  - Ensure that the number of closing angle brackets equals the number of opening angle brackets.

## Creating LINQ Queries

To query data from the database, you can use language-integrated query (LINQ), which is a part of .NET Framework. In the code of Acumatica Framework-based applications, you can use both the standard query operators (provided by LINQ libraries) and the Acumatica Framework-specific operators that are designed to query database data.

In this chapter, you can find information on how to create LINQ queries. For the general information about data querying, see *Querying Data in Acumatica Framework*.

## In This Chapter

- *LINQ in Acumatica Framework*
- *Deferred LINQ Query Execution*
- *Fallback to the LINQ to Objects Mode*
- *To Select Records by Using LINQ*
- *To Append LINQ Expressions to BQL Statements*

## LINQ in Acumatica Framework

You can use language-integrated query (LINQ) provided by the `System.Linq` library when you need to select records from the database in the code of Acumatica Framework-based applications or if you want to apply additional filtering to the data of a BQL query. However, you still have to use business query language (BQL) to define the data views in graphs and to specify the data queries in attributes of data fields. For details about BQL, see [Creating Fluent BQL Queries](#) and [Creating Traditional BQL Queries](#). For more information about the differences between LINQ and BQL, see [Comparison of Fluent BQL, Traditional BQL, and LINQ](#).

## Data Access Classes in LINQ

In LINQ expressions, to access data from the database tables, you use data access classes (DACs). For details on DACs, see [Data Access Classes](#).

You use property fields of DACs when you need to specify table columns in LINQ expressions. (The name of the property field starts with an uppercase letter. Do not confuse it with the class field, which has the same name but starts with lowercase letter.)

## Query Syntax

To configure a LINQ query, you can use the following variants of syntax:

- Query expressions, which use standard query operators from the `System.Linq` namespace (such as `where` or `orderby`) or Acumatica Framework-specific operators from the `PX.Data.SQLTree` namespace (such as `SQL.BinaryLen`, which is shown in the following example of this syntax).

```
ProductMaint graph = PXGraph.CreateInstance<ProductMaint>();
var goods = from p in graph.Select<Product>()
            where
```

```

        p.ProductCD.Length == 5 &&
        p.GroupMask.Length == 4 &&
        (p.WorkGroupID & 0b10) != 0
    select new
    {
        p.ProductID,
        p.ProductCD,
        p.ProductName,
        Len = p.ProductName.Length,
        BLen = SQL.BinaryLen( p.ProductName) + 1,
        p.GroupMask,
        p.WorkGroupID
    };

```

- Explicit (method-based) syntax. The arguments of the methods used in this syntax are lambda expressions. In these expressions, you can use the standard C# operators and Acumatica Framework-specific operators from the `PX.Data.SQLTree` namespace (such as `SQL.BinaryLen`, which is shown in the following code). The code below is equivalent to the query expression shown above.

```

ProductMaint graph = PXGraph.CreateInstance<ProductMaint>();
var goods = graph.Select<Product>()
    .Where( p =>
        p.ProductCD.Length == 5 &&
        p.GroupMask.Length == 4 &&
        (p.WorkGroupID & 0b10) != 0)
    .Select( p => new
    {
        p.ProductID,
        p.ProductCD,
        p.ProductName,
        Len = p.ProductName.Length,
        BLen = SQL.BinaryLen(p.ProductName) + 1,
        p.GroupMask, p.WorkGroupID
    });

```

For details about composing LINQ queries, see [To Select Records by Using LINQ](#). In the code examples of this guide, we use explicit syntax.

### Simultaneous Use of LINQ and BQL

The `Select` method of all `PXSelect` classes of Acumatica Framework return `PXResultset<T0>`, which implements the `IQueryable<PXResult<T0>>` interface. That is, you can work with the query expression defined with BQL by using LINQ. The following code shows an example of additional filtering of data of the BQL query.

```

//BQL statement
var Products = new PXSelect<Product,
    Where<Product.productCD, Like<string_D>>>(graph);
//Use of LINQ for the result of the BQL query
var goods = Products.Select()
    .Where(p => p.GetItem<Product>().StockUnit == "item");
//Execution of the query
foreach (var good in goods) {
    var prod = good.GetItem<Product>();
}

```

However, you cannot work with the query defined with LINQ by using BQL.

For details about how to use LINQ and BQL simultaneously, see [To Append LINQ Expressions to BQL Statements](#).

## Deferred LINQ Query Execution

Queries defined with LINQ in Acumatica Framework implement the `IQueryable` interface—that is, for these queries, the system generates expression trees and executes the queries only when they are iterated over. For details about query execution, see [Data Query Execution](#).

### Execution of the LINQ Query in Code

To execute the query, you do one of the following:

- Call the `ToList` or `ToArray` method for the query, as shown in the following code.

```
//query is a LINQ expression.
var data = query.ToList();
```

- Iterate the query by using the `foreach` statement, as shown in the following code.

```
//query has the IQueryable<Product> type.
//Product is a DAC.
foreach (Product record in query)
{
    ...
}
```

### Explicit Merge of Records with PXCache

The system merges with `PXCache` the records the system has retrieved from the database by using the LINQ queries, as described in [Merge of the Records with PXCache](#).

You may need to explicitly merge the records retrieved from the database with a particular `PXCache` object. To merge the records with `PXCache` explicitly, you use the `Merge` method, as shown in the following code example.

```
ProductMaint graph = PXGraph.CreateInstance<ProductMaint>();
var query = new PXSelectReadOnly<Product>(graph).Select()
    .Select(r => r.GetItem<Product>())
    .Where(p => SQL.Like(p.ProductName, "%d%") && p.StockUnit != null)
    .OrderBy(p => p.ProductID);
var result = query.Merge(p => p).ToArray();
```

You also can specify that the query should not be merged with `PXCache` by using the `ReadOnly` method, as shown in the following code example.

```
ProductMaint graph = PXGraph.CreateInstance<ProductMaint>();
var query = new PXSelect<Product>(graph).Select()
    .Select(r => r.GetItem<Product>())
    .Where(p => SQL.Like(p.ProductName, "%d%") && p.StockUnit != null);
var result = query.ReadOnly().ToArray();
```

## Fallback to the LINQ to Objects Mode

With LINQ, you may not be able to filter records by using custom C# functions. For example, suppose that your C# function filters records by a regular expression, which cannot be converted to standard SQL functions. If the system cannot convert a custom C# function in a LINQ statement to an SQL query tree, the system falls back to LINQ to Objects mode—that is, the system executes the data query in memory, which can lead to degradation of the application's performance.

The following code shows an example of a situation when the system falls back to LINQ to Objects mode. In this example, the system selects from the database all records from the `CRCCase` table, and

then, in memory, orders the retrieved records by the `Date` column and selects the records that satisfy the condition specified by using the `MyHelpers.IsHighPriority` function.

```
// MyHelpers.IsHighPriority is a custom function.
var results = graph
    .Select<CRCase>()
    .OrderByDescending(c => c.Date)
    .Where(c => MyHelpers.IsHighPriority(c));

foreach (CRCase case in results)
{
    ...
}
```



: If the system falls back to LINQ to Objects, only the results of the base `PXSelectBase` query are merged with `PXCache` as described in [Merge of the Records with PXCache](#). The `Merge` and `ReadOnly` methods do not affect the merge of records with `PXCache` for the queries that caused fallback.

The LINQ fallback is supported in Acumatica Framework for compatibility with previous versions. The system writes to the trace log about all situations in which the system falls back to LINQ to Objects mode. Therefore, we strongly recommend that you investigate the trace log for such issues and fix the issues in one of the following ways:

- Remove the custom C# functions that cause fallback so that the full query is executed in the database.
- Append the `AsEnumerable()` method to the part of the query that can be converted to SQL, and add after it the conditions that include custom C# functions. In this case, the system does not waste resources trying to build the SQL query tree for the whole query. Instead, the system builds the SQL query tree for the part of the query that has `AsEnumerable()` appended and performs the corresponding request to the database, while the custom C# conditions of the query are processed in memory. For example, the code example above can be modified as follows.

```
// MyHelpers.IsHighPriority is a custom function.
var results = graph
    .Select<CRCase>()
    .OrderByDescending(c => c.Date).AsEnumerable()
    .Where(c => MyHelpers.IsHighPriority(c));

foreach (CRCase case in results)
{
    ...
}
```



: The system executes the following SQL query for the code above, where `[list of columns]` is the list of columns of the `CRCase` table.

```
SELECT [list of columns] FROM CRCase
ORDER BY CRCase.Date DESC
```

## To Select Records by Using LINQ

To select records from the database by using language-integrated query (LINQ), you use the standard query operators (provided by LINQ libraries), as described in this topic. In the LINQ queries, you use the property fields of DACs to specify the columns of database tables. (The name of the property field starts with an uppercase letter. Do not confuse it with the class field, which has the same name but starts with lowercase letter.)



: After you have composed a LINQ expression, to execute the query defined by this expression, you have to call the `ToList` or `ToArray` method for the query, or iterate the query by using the `foreach` statement. For example, the following code executes the query defined by a LINQ expression.

```
//query is a LINQ expression
```

```
var data = query.ToList();
```

For details about the execution of LINQ expressions, see [Deferred LINQ Query Execution](#).

### Before You Proceed

- Add the `using` directives shown below to your code.

```
using PX.Data;
using PX.Data.SQLTree;
using System.Linq;
```

- Make sure that the application database has the database tables from which you are going to request data, and that the application defines the data access classes (DACs) for these tables. For more information on defining DACs, see [Data Access Classes](#).

### To Filter Records

To filter records in the database table to be retrieved, construct the LINQ expression by using the `Where` LINQ method and the needed conditions. In the conditions, use the property field defined in the DAC, such as `Product.ProductID`.

The following LINQ expression uses the C# logical operators (`||`, `&&`, and `!`) to define multiple conditions.

```
ProductMaint graph = PXGraph.CreateInstance<ProductMaint>();
IQueryable<Product> query = graph.Select<Product>().Where(prod =>
    prod.BookedQty > prod.AvailQty
    || prod.AvailQty < prod.MinAvailQty
    || prod.AvailQty == null);
```



: This LINQ expression is equivalent to the following SQL query. In this SQL query, [list of columns] is the list of columns of the `Product` table.

```
SELECT [list of columns] FROM Product
WHERE Product.BookedQty > Product.AvailQty
    OR Product.AvailQty < Product.MinAvailQty
    OR Product.AvailQty IsNull
```

### To Order Records

To order records to be retrieved from the database table, construct the LINQ expression by using the `OrderBy` or `OrderByDescending` LINQ method and the needed property fields of the DAC, such as `Product.ProductID`.

The following sample LINQ expression selects all `Product` data records and sorts them by the `UnitPrice` field in ascending order.

```
ProductMaint graph = PXGraph.CreateInstance<ProductMaint>();
IQueryable<Product> query = graph.Select<Product>().OrderBy(prod => prod.UnitPrice)
    .ThenByDescending(prod => prod.AvailQty);
```



: This LINQ expression is equivalent to the following SQL query, where [list of columns] is the list of columns of the `Product` table.

```
SELECT [list of columns] FROM Product
ORDER BY Product.UnitPrice, Product.AvailQty DESC
```

### To Query Multiple Tables

To join multiple tables, construct the LINQ expression by using the `Join`, `LeftJoin`, `GroupJoin`, and `FullJoin` LINQ methods and the needed property fields of DACs, such as `SalesOrder.OrderNbr`.

The following sample LINQ expression performs an inner join of the `SalesOrder` and `OrderDetail` DACs by the `OrderNbr` field.

```
SalesOrderEntry graph = PXGraph.CreateInstance<SalesOrderEntry>();
var query = graph.Select<SalesOrder>()
    .Join(graph.Select<OrderLine>(),
        ord => ord.OrderNbr, ordDet => ordDet.OrderNbr,
        (ord, ordDet) => new { SalesOrder = ord, OrderDetail = ordDet });
```



: This LINQ expression is equivalent to the following SQL query, where [list of columns] is the list of columns of the `Product` table.

```
SELECT [list of columns] FROM SalesOrder
INNER JOIN OrderDetail
    ON OrderDetail.OrderNbr = SalesOrder.OrderNbr
```

## To Group or Aggregate Records

To group or aggregate records, do the following:

1. Construct the LINQ expression by using the `GroupBy` LINQ method and the needed property fields of the DAC, such as `Product.CategoryCD`. (The name of the property field starts with an uppercase letter. Do not confuse it with the class field, which has the same name but starts with lowercase letter.)
2. Append to the expression the `Select` LINQ method.



: You have to use the `Select` method after `GroupBy` to eliminate the number of requests that the system performs to the database.

The following sample LINQ expression groups the records of the `Product` table by the `CategoryCD` field.

```
ProductMaint graph = PXGraph.CreateInstance<ProductMaint>();
var query = graph.Select<Product>().GroupBy(prod => prod.CategoryCD).
    Select(group => new { CategoryCD = group.Key });
```



: This LINQ expression is equivalent to the following SQL query.

```
SELECT Product.CategoryCD
FROM Product
GROUP BY Product.CategoryCD
```

## To Select Particular Columns of Records

To select particular columns, specify the corresponding property fields of DACs in the `Select` clause.

The following example selects the values of the `ProductCD` and `AvailQty` fields for all records of the `Product` table.

```
ProductMaint graph = PXGraph.CreateInstance<ProductMaint>();
var results = graph.Select<Product>()
    .Select(p => new { ProductCD = p.ProductCD, AvailQty = p.AvailQty });
```



: The system executes the following SQL query for the code above.

```
SELECT Product.ProductCD, Product.AvailQty FROM Product
```

## To Append LINQ Expressions to BQL Statements

By using LINQ, you can work with a query expression that is defined with BQL. In this topic, you can find out how to apply additional filtering to a BQL statement and join a table to a BQL statement by using LINQ.

### Before You Proceed

- Add the `using` directives shown below to your code.

```
using PX.Data;
using PX.Data.SQLTree;
using System.Linq;
```

- Make sure that the application database has the database tables from which you are going to request data, and that the application defines the data access classes (DACs) for these tables. For more information on defining DACs, see [Data Access Classes](#).

### To Append LINQ Expressions to BQL Statements

To append LINQ expressions to BQL statements, do the following:

1. Configure a BQL query derived from `PXSelectBase` either in fluent BQL or in traditional BQL.
2. Call the `Select()` method of `PXSelectBase`, and append the LINQ query to the result. Because the result of the `Select()` method call is a `PXResultset<>` object, you need to cast it to a DAC type by using the `PXResult.GetItem<DacType>()` method or direct casting.

The following example appends LINQ joining and filtering to a BQL query.

```
using PX.Data;
using PX.Data.SQLTree;
using System.Linq;
using PX.Data.BQL.Fluent;

ProductMaint graph = PXGraph.CreateInstance<ProductMaint>();

//Configure a BQL query
var products = new SelectFrom<Product>().
    Where<Product.productCD.IsLike<string_D>>.View(graph);

//Append joining and filtering by using LINQ
var goods = products.Select().Join(graph.Select<SupplierProduct>(),
    p => p.GetItem<Product>().ProductID,
    sp => sp.ProductID,
    (p, sp) => new { p = p.GetItem<Product>(), sp }
).Where(sp => sp.p.UnitPrice >= 0 && sp.sp.LastPurchaseDate == null);

//Execute the query
foreach (var item in goods)
{
    ...
}
```



**:** The system executes the following SQL query for the code above. In this SQL query, [list of columns] is the list of columns of the `Product` and `SupplierProduct` tables.

```
SELECT [list of columns] FROM Product
INNER JOIN SupplierProduct
    ON Product.ProductID = SupplierProduct.ProductID
WHERE Product.ProductCD LIKE 'D' AND Product.UnitPrice>=0
    AND SupplierProduct.LastPurchaseDate IS NULL
```

## Defining Relationships Between DACs

In this chapter, you can learn how to define relationships between data access classes (DACs) by using the special classes for primary and foreign keys.

### In This Chapter

- *Establishment of a Relationship Between Two DACs*
- *To Define a Primary Key*
- *To Define a Foreign Key*

## Establishment of a Relationship Between Two DACs

To define a relationship between two tables, you need to define the primary keys of the parent and child tables. Also, in the child table, you need to define the foreign key that refers to the primary key of the parent table.

In the code of an Acumatica Framework-based application, you can define this relationship as follows:

- To define a primary key of a table, for the set of key fields of the data access class (DAC) that corresponds to the table, you set the `IsKey` property of the data type attribute to `true`.
- To define a foreign key of a table, in the DAC that corresponds to the table, you mark the field that contains the foreign key with one of the following attributes: `PXForeignReference`, `PXSelector`, or `PXParent`.

To select a record from the database by its primary or foreign key, you can use a `Select` statement in business query language (BQL) or use the methods of the attributes mentioned above.

Another way to define a relationship between two tables is to use the classes and methods that are specially designed for the definition of primary and foreign keys. This approach, which is described in this topic, provides the following advantages:

- These classes provide static information that a compiler can use to identify errors in the code.
- You can use runtime information about primary keys to select records by their keys.
- These classes and methods have no other meanings and use cases; conversely, the `PXForeignReference`, `PXSelector`, and `PXParent` attributes can be used for other purposes.
- These classes and methods are optimized for the selection of records from the database; therefore, using them improves database access performance on record selection.

### Definition of a Primary Key

You define a primary key of a DAC by using the `PrimaryKeyOf<Table>.By<keyFields>` class. With this class, you can define simple keys (with one key field) and compound keys (with up to five key fields). In the primary key definition, you have to define the `public Find` method, which calls the `protected FindBy` method. A definition of a compound key is shown in the following example.

```
using PX.Data.ReferentialIntegrity.Attributes;

public partial class SOLine : PX.Data.IBqlTable
{
    public class PK : PrimaryKeyOf<SOLine>.By<orderType, orderNbr, lineNbr>
    {
        public static SOLine Find(
            PXGraph graph, string orderType, string orderNbr, int lineNbr)
            => FindBy(graph, orderType, orderNbr, lineNbr);
    }

    public abstract class orderType : PX.Data.IBqlField { }
```



```

    public abstract class orderNbr : PX.Data.IBqlField { }
    public abstract class lineNbr : PX.Data.IBqlField { }
}

```

### Definition of a Foreign Key

You can define a foreign key based on the primary key of the referenced table, as shown in the following code.

```

//Definition of the primary key
public partial class SOOrder : PX.Data.IBqlTable
{
    public class PK : PrimaryKeyOf<SOOrder>.By<orderType, orderNbr>
    {
        public static SOOrder Find(
            PXGraph graph, string orderType, string orderNbr) =>
            FindBy(graph, orderType, orderNbr);
    }

    public abstract class orderType : PX.Data.IBqlField { }
    public abstract class orderNbr : PX.Data.IBqlField { }
}

//Definition of the foreign key based on the primary key
public partial class SOLine : PX.Data.IBqlTable
{
    public class SOOrderFK : SOOrder.PK.ForeignKeyOf<SOLine>
        .By<orderType, orderNbr> { }

    public abstract class orderType : PX.Data.IBqlField { }
    public abstract class orderNbr : PX.Data.IBqlField { }
}

```

### Selection of a Record by Key Fields

If a primary key is defined for a DAC, you can select a record by using the values of the key fields of the record, as shown in the following example.

```

SOLine line = SOLine.PK.Find(
    this, split.OrderType, split.OrderNbr, split.LineNbr.Value);

```



**Tip** : The Find method encapsulates a `PXSelectReadOnly<Table, Where<...>>.SelectWindowed(graph, 0, 1, keys)` call. Therefore, the code above can replace the following code written using BQL.

```

SOLine line = PXSelectReadOnly<SOLine,
    Where<SOLine.orderType, Equal<Required<SOLine.orderType>>,
        And<SOLine.orderNbr, Equal<Required<SOLine.orderNbr>>,
        And<SOLine.lineNbr, Equal<Required<SOLine.lineNbr>>>>>
>.Select(this, split.OrderType, split.OrderNbr, split.LineNbr);

```

You can also select a record by using a record of the same type with the key fields specified, as shown in the following example.

```

InventoryItem actualItem = InventoryItem.PK.Find(this, notActualItem);

```

If a foreign key is defined for a DAC, you can select the parent and child records, as shown in the following code.

```

//Selection of the parent record
SOOrder order = SOLine.SOOrderFK.FindParent(this, soLine);
//Selection of the child records
IEnumerable<SOLine> lines = SOLine.SOOrderFK.SelectChildren(this, soOrder);

```

## Use of Primary and Foreign Keys in Attributes

You can use static foreign keys, defined as described in [Definition of a Foreign Key](#), for the configuration of the `PXForeignReference` and `PXParent` attributes, as shown in the following example.

```
public partial class SOLine : PX.Data.IBqlTable
{
    public class SOOrderFK : SOOrder.PK.ForeignKeyOf<SOLine>
        .By<orderType, orderNbr> { }
    public class InventoryFK : InventoryItem.PK.ForeignKeyOf<SOLine>
        .By<inventoryID> { }

    public abstract class orderType : PX.Data.IBqlField { }

    [...]
    [PXParent(typeof(SOOrderFK))]
    public virtual String OrderNbr { get; set; }
    public abstract class orderNbr : PX.Data.IBqlField { }

    [...]
    [PXForeignReference(typeof(InventoryFK))]
    public virtual Int32? InventoryID { get; set; }
    public abstract class inventoryID : PX.Data.IBqlField { }
}
```

Because all primary keys, which are defined as described in this topic, implement the `IPrimaryKey` interface, you can use primary and foreign keys in the scope of custom attributes, as shown in the following example.

```
public class SomeAttribute : PXEventSubscriberAttribute, ...
{
    private readonly IPrimaryKey _pk;
    public SomeAttribute(Type pkType)
    {
        _pk = (IPrimaryKey)Activator.CreateInstance(pkType);
    }

    public void SomeHandler(PXCache cache, PXSomeEventArgs e)
    {
        IBqlTable row = _pk.Find(cache.Graph, e.NewValue);
        ...
        _pk.StoreCached(cache.graph, row);
        ...
        row = _pk.Find(cache.Graph, row);
    }
}
```

## To Define a Primary Key

You can define the primary key of a data access class (DAC) by using the `PrimaryKeyOf<Table>.By<keyFields>` class. With this class, you can define simple keys (with one key field) and compound keys (with up to five key fields) and select records by using these keys, as described in this topic.

### To Define a Simple Primary Key and Select a Record by This Key

1. In the DAC, declare a `PrimaryKeyOf<Table>.By<keyFields>` descendant with the public `Find` method, which calls the protected `FindBy` method, as shown in the following code.

```
using PX.Data.ReferentialIntegrity.Attributes;

public partial class InventoryItem : PX.Data.IBqlTable
{
    public class PK : PrimaryKeyOf<InventoryItem>.By<inventoryID>
    {

```

```

        public static InventoryItem Find(PXGraph graph, int inventoryID)
            => FindBy(graph, inventoryID);
    }

    public abstract class inventoryID : PX.Data.IBqlField { }
}

```

2. Use the primary key to select a record, as shown in the following code.

```
InventoryItem item = InventoryItem.PK.Find(this, soLine.InventoryID.Value);
```

### To Define a Compound Primary Key and Select a Record by This Key

1. In the data access class (DAC), declare a `PrimaryKeyOf<Table>.By<keyFields>` descendant with the `public Find` method, which has the needed number (up to five) of key fields. The `Find` method must call the protected `FindBy` method with the same number of key fields, as shown in the following code.

```

using PX.Data.ReferentialIntegrity.Attributes;

public partial class SOLine : PX.Data.IBqlTable
{
    public class PK : PrimaryKeyOf<SOLine>
        .By<orderType, orderNbr, lineNbr>
    {
        public static SOLine Find(
            PXGraph graph, string orderType, string orderNbr, int lineNbr)
            => FindBy(graph, orderType, orderNbr, lineNbr);
    }

    public abstract class orderType : PX.Data.IBqlField { }
    public abstract class orderNbr : PX.Data.IBqlField { }
    public abstract class lineNbr : PX.Data.IBqlField { }
}

```

2. Use the compound primary key to select a record, as shown in the following code.

```
SOLine line = SOLine.PK.Find(
    this, split.OrderType, split.OrderNbr, split.LineNbr.Value);
```

## To Define a Foreign Key

You can define a foreign key based on the primary key of the parent data access class (DAC) and select records by this key, as described in this topic.

### To Define a Foreign Key and Use It to Select Records

1. In the DAC of the parent table, define the primary key, as described in [To Define a Primary Key](#). The following code shows an example of the definition of the primary key that is used in the following code examples in this instruction.

```

public partial class SOOrder : PX.Data.IBqlTable
{
    public class PK : PrimaryKeyOf<SOOrder>.By<orderType, orderNbr>
    {
        public static SOOrder Find(
            PXGraph graph, string orderType, string orderNbr)
            => FindBy(graph, orderType, orderNbr);
    }

    public abstract class orderType : PX.Data.IBqlField { }
    public abstract class orderNbr : PX.Data.IBqlField { }
}

```

2. In the DAC of the child table, define the foreign key based on the primary key of the parent table, as shown below.

```
public partial class SOLine : PX.Data.IBqlTable
{
    public class SOOrderFK : SOOrder.PK.ForeignKeyOf<SOLine>
        .By<orderType, orderNbr> { }

    public abstract class orderType : PX.Data.IBqlField { }
    public abstract class orderNbr : PX.Data.IBqlField { }
}
```

3. Use the foreign key to select the parent record or child records, as shown in the following code.

```
//Select the parent record
SOOrder order = SOLine.SOOrderFK.FindParent(this, soLine);
//Select the child records
IEnumerable<SOLine> lines = SOLine.SOOrderFK.SelectChildren(this, soOrder);
```

## Working with Data in Cache and Session

---

In this chapter, you can learn how to store graph data in the session and use slots to cache data objects.

### In This Chapter

- [Session](#)
- [Session Sharing Between Application Servers](#)
- [Storing of Graph Data in the Session](#)
- [Use of Slots to Cache Data Objects](#)

### Session

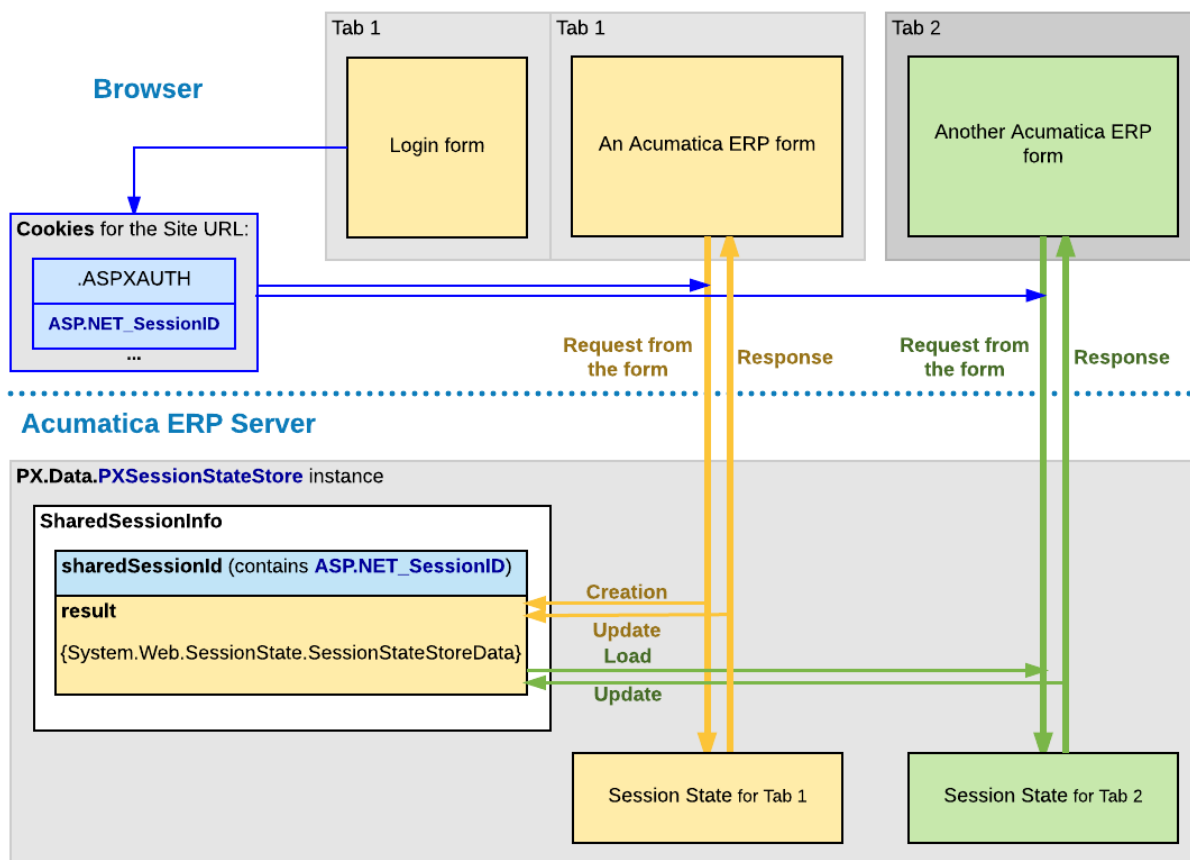
The Acumatica ERP server creates a separate session for each browser tab or window that opens an Acumatica ERP form.

The server creates the first session for a user after the user authorization when the starting form is loading. Then the server does the following:

- Saves the user authorization data (.ASPXAUTH) and the session ID (ASP.NET\_SessionID) in the browser cookies for the website URL
- Creates the shared session data to be used for the Acumatica ERP forms opened in new browser tabs and windows
- Saves the shared session data in the storage that is specified in the website configuration

When the user opens a form of Acumatica ERP in a new browser tab, the server creates a new session that is based on the previous session data. To access the shared data, the server uses the session ID from the cookies, which are added to the request by the browser.

The following diagram shows how the server of Acumatica ERP manages the shared session data that is used for multiple sessions of a single user.



**Figure: Use of shared data for multiple sessions of a user**

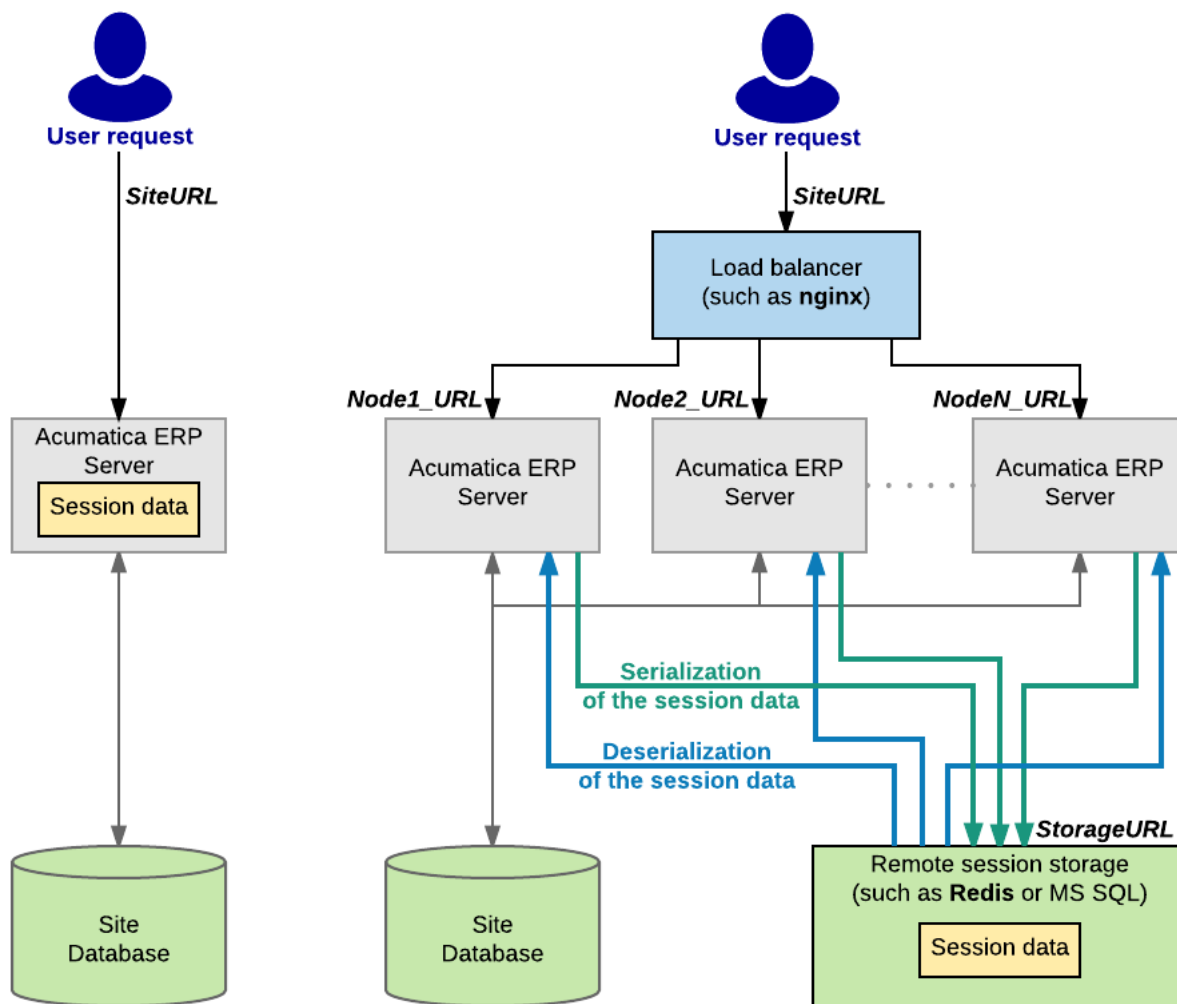
If the session data has been changed during the processing of a request, the server updates the data in the shared session data store. For example, if the user clicks **Copy** on a form toolbar to copy the form data, the data is stored in the shared session, so that it is accessible for the **Paste** action in another session of the same user.

To distinguish different sessions that have the same `ASP.NET_SessionID`, the server adds to each new session a unique identifier that consists of the *W* character and a number value wrapped in parentheses. In the browser, you can see such an identifier in the site URL, as with the bolded part in the following example: `http://localhost/MySite/(W(3))/Main.aspx?ScreenId=AR301000`.

## Session Sharing Between Application Servers

To achieve horizontal scalability and fault tolerance, an application written with Acumatica Framework can be configured to run in a cluster of application servers behind a load balancer. With this configuration, it is not possible to predict the application server that will receive the next request from the client. In this model, session specific data must be shared between the application servers.

The following diagram shows difference in storing session data on a stand-alone server and in a cluster. On a stand-alone Acumatica ERP server, session data is stored in the server memory. In a cluster of application servers, session data must be serialized and stored in a high-performance remote server, such as Redis or Microsoft SQL.



**Figure: Storing session data on a stand-alone server and in a cluster**

The cost of serialization and the amount of data that need to be shared between application servers is often the main challenge to scaling complex business applications horizontally.

Acumatica Framework implements the following techniques to address issues related to session-state management without sacrificing performance, fault tolerance, or scalability:

- Objects on the application server are created on each request and disposed after the request execution. The application state is preserved in the session through the serialization mechanism.
- Data serialized into the session is minimized to store only modified data (inserted, deleted, held and modified records). (Serialization and retrieval times are directly proportional to the size of the serialized data.)
- The rest of the data is extracted from the database on demand and built around the session data. (A custom algorithm that extracts only the data required for the current request execution from the database is implemented.)
- A custom serialization mechanism is implemented to serialize only relevant data and reduce the amount of service information. (The standard serialization mechanism implemented in the Microsoft .NET platform is generic and cannot be optimized when used for a specific task.)
- Hash tables, constraints, relations, and indexes concerned with the execution of business logic are created strictly on demand. This technique allows the user to avoid execution of these operations on each request if not needed. (Creation of indexes, constraints, hash tables, and relations consumes a significant amount of CPU and runtime memory.)

## Storing of Graph Data in the Session

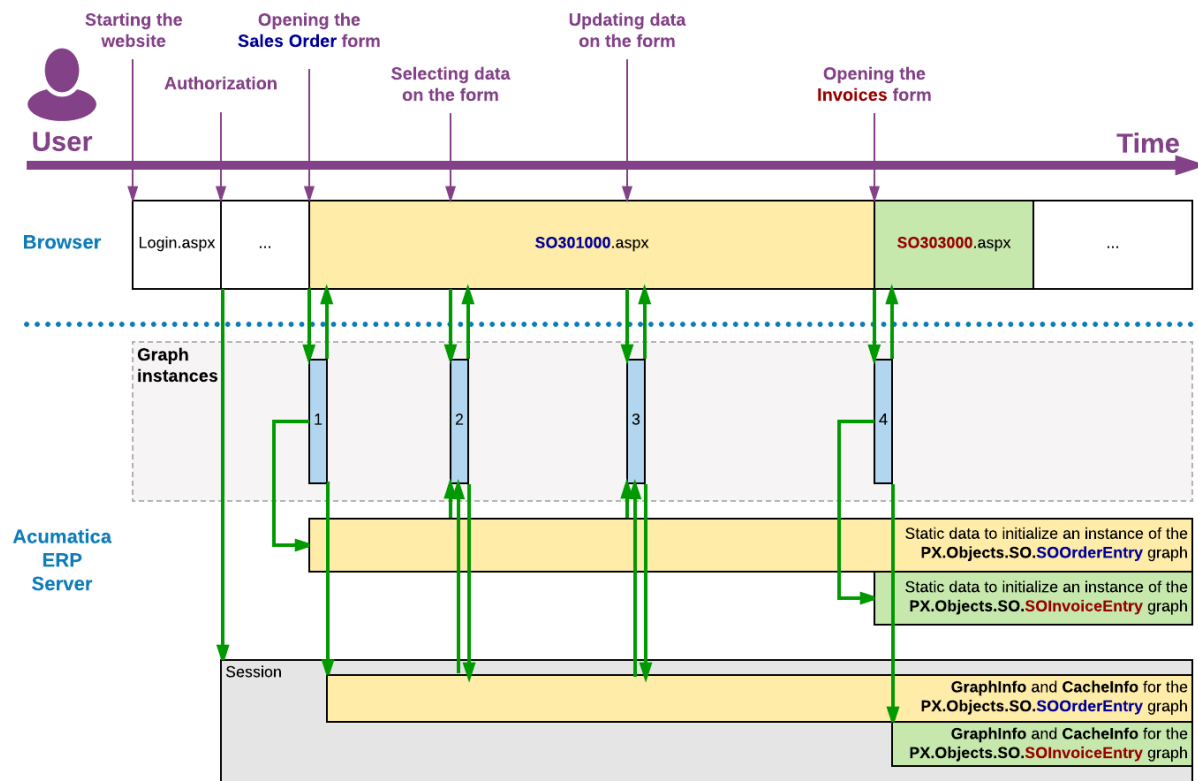
Acumatica ERP keeps all modified data records in the cache. Therefore, you do not change a data record in the database directly when you modify its value in the user interface.



The system commits the changes to the database in the following cases:

- The user clicks **Save**.
- A request is sent through the Web Service API.
- The `Actions.PressSave` method is invoked on the graph instance.
- The `PXAutoSaveAttribute` attribute is defined for a data access class. As a result, the `PXCache<>.Unload` method automatically stores in the database all the changes of the appropriate data records at the end of each round trip.

A graph instance exists on the server only while a user request is being processed, and it is destroyed right after this processing. The following diagram shows that a graph instance is created to process a user request on the Acumatica ERP server and destroyed once processing is completed.



**Figure: Processing of user requests**

In the diagram above, the blue rectangles labeled 1, 2, 3, and 4 indicate the lifetime of graph instances. After a graph instance completes the processing of a request, the system stores the graph state in the session. It also stores the inserted, deleted, held, and modified records of the cache that are required to restore the state and data of the graph for the processing of the subsequent user request on the same Acumatica ERP form.



On a stand-alone Acumatica ERP server, session data is stored in the server memory. In a cluster of Acumatica ERP servers, session data must be serialized and stored in external high-performance session state storage. (For more information on storing session data in a cluster, see [Session Sharing Between Application Servers](#).)

For a user request on an Acumatica ERP form, the following operations are executed in the system:

1. The application server creates a graph instance that is specified in the `TypeName` property of the `PXDataSource` control of the form. (For more information about the initialization of graph views, caches, actions, and event handlers, see [Initialization of an Event Handler Collection](#).)

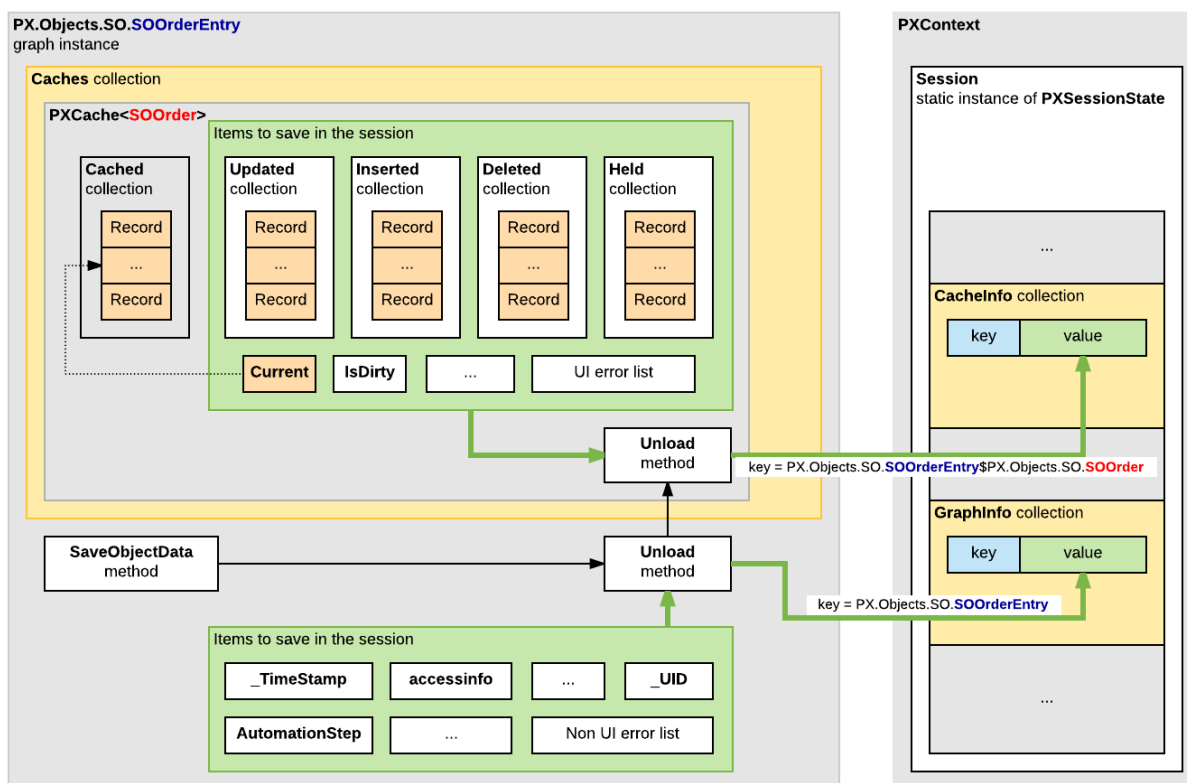
2. If the user session contains graph data that has been stored during a previous request, the system loads the graph state and the cache data from the session.
3. The graph instance processes the requested data on the data view that is specified in the ASPX code in the `DataMember` property of the control container for the data to be processed. To process the data, the system invokes the `ExecuteSelect`, `ExecuteInsert`, `ExecuteUpdate`, or `ExecuteDelete` method of the graph, based on the request type. The invoked method implements the logic of the appropriate scenario to add the request data to the cache and to execute the event handlers defined for the data fields and records in the cache. (See [Data Manipulation Scenarios](#) for details.) The cache then merges the data retrieved from the database with the data restored from the session, and the application accesses the data as if the entire data set had been preserved from the time of the previous request.
4. The graph instance returns the request results to the `PXDataSource` control of the form.
5. The system stores in the session the graph state and the modified data of the cache.



: Because the graph instance is no longer being used by the application server, the .NET Framework garbage collector then clears the memory allocated for the graph instance.

While a graph is instantiated, all the cached data of the graph is saved in the appropriate `PXCache` objects that are created in the graph instance based on the data access class (DAC) declarations. To preserve the modified entity data between user requests, the cache controller saves the `Updated`, `Inserted`, `Deleted`, and `Held` collections of each `PXCache` object in the session.

The following diagram shows how the graph state and cache data are stored in the `Session` object.



**Figure: The graph data stored in the session**

In the diagram above, notice the following:

- The items of the graph instance are stored in the `GraphInfo` collection of the `Session` object as a key-value pair, where the `key` is equal to the full name of the graph.
- The items of a `PXCache` object are stored in the `CacheInfo` collection of the `Session` object as a key-value pair, where the `key` consists of the following parts separated by the `$` symbol:
  1. The full name of the graph



## 2. The full name of the DAC



: When you instantiate a graph from code, the system will not load data from the session, because you may want to perform redirection or other processing. You can direct the system to load this data by using the `PXPreserveScope` class, as the following code snippet shows.

```
using (new PXPreserveScope())
{
    GraphName graph = PXGraph.CreateInstance<GraphName>();
    graph.Load();
    ...
}
```

## Use of Slots to Cache Data Objects

If you have to cache a data object from your code, you can use the slots provided by the `PXContext` and `PXDatabase` classes. By using these slots, you can cache any type of data object without restrictions.

A slot provided by the `PXContext` class exists in the memory of the application server only during the current HTTP request. Therefore, you can use these slots for quick data exchange between different server modules while the server processes a single request.

A slot of the `PXDatabase` class is stored in the server memory until you clear the slot. Therefore, you can use such a slot to cache a data object for a long time—for example, to read the cached data during a future HTTP request.

If a `PXDatabase` slot is used to cache the data that is obtained from the database tables, you can use a special API to automatically update the data in the slot when any of these tables has been changed.

For detailed information on using slots, see the sections of this topic.

### Caching Data in `PXContext` Slots

If you need to keep a data object during a single HTTP request, we recommend that you cache the object in a slot provided by the `PXContext` class.

You can use the following public static methods of the class to save a data object in a slot.

Method	Description
<code>public static ObjectType SetSlot&lt;ObjectType&gt;(ObjectType value)</code>	Stores the specified data object under the key that is created on the base of the object type.
<code>public static ObjectType SetSlot&lt;ObjectType&gt;(string key, ObjectType value)</code>	Stores the specified data object under the key that is defined by the first parameter.

The following example shows how you can save the `MyData` object in the slot of the current HTTP context under the key that is the same as the object type.

```
PXContext.SetSlot<MyDataType>(MyData);
```

To get a data object that is cached in the current HTTP context, you can use the following methods of the `PXContext` class.

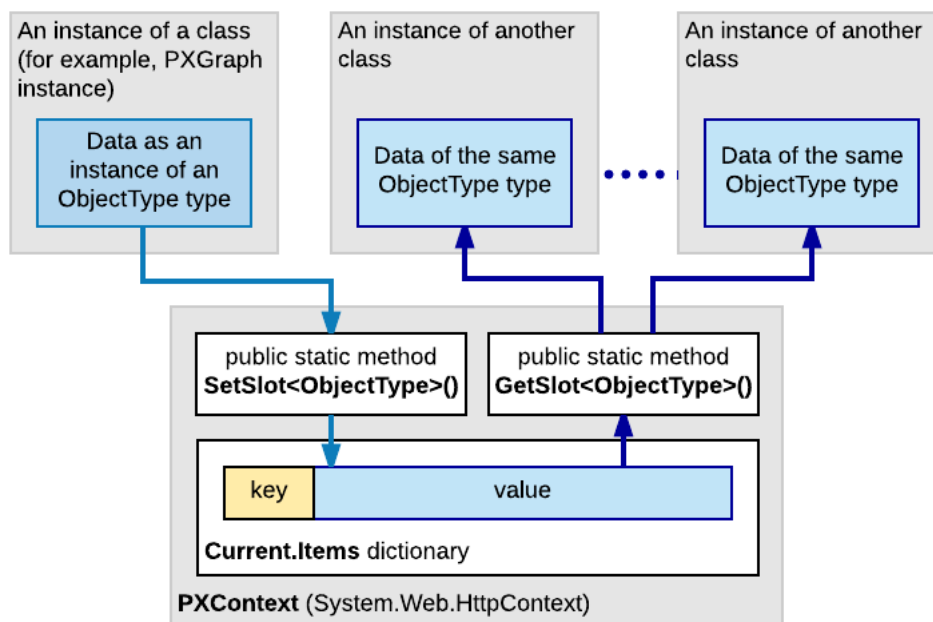
Method	Description
<code>public static ObjectType GetSlot&lt;ObjectType&gt;()</code>	Returns the data object that is cached under the key that is created on the base of the object type.

Method	Description
<pre>public static ObjectType GetSlot&lt;ObjectType&gt; (string key)</pre>	Returns the data object that is cached under the specified key.

The following example shows how you can get from the slot of the current HTTP context the *MyData* object that is cached under the *MyData22* key.

```
var MyData = PXContext.GetSlot<MyDataType>("MyData22");
```

The following diagram illustrates how you can use a data object cached by using a slot provided by the *PXContext* class.



**Figure: Caching data in a slot of the *PXContext* class**


You do not need to delete the data saved in the *PXContext* class slots, because the system deletes these slots from the server memory along with the data of the current HTTP context created for the current request.

### Caching Data in *PXDatabase* Slots

If you need to keep a data object in the server memory for a long time, we recommend that you cache the object in a slot provided by the *PXDatabase* class.

You can use the following public static methods of the class to cache a data object in a slot and to get the cached object from the slot.

Method	Description
<pre>public static ObjectType GetSlot&lt;ObjectType&gt; (string key, params Type[] tables)</pre>	If the <i>PXDatabase</i> slots contain a valid data object of the specified type saved under the key defined by the first parameter, returns this data object. Otherwise, the method creates a new object of the specified type, saves this empty object in the slot under the key defined by the first parameter, and returns the data object that is used by the calling code to save the needed data. The list of the table types specified in the <i>params</i> parameter is used to invalidate the slot if any table of the list has been changed in the database.

Method	Description
	 : If this method is used to cache a data object of an <code>ObjectType</code> class inherited from the <code>IPrefetchable&lt;&gt;</code> interface, the <code>GetSlot&lt;&gt;</code> method invokes the <code>Prefetch</code> method of the object without a parameter.
<pre>public static ObjectType GetSlot&lt;ObjectType, Parameter&gt; (string key, Parameter parameter, params Type[] tables)</pre>	<p>Is used for caching a data object of an <code>ObjectType</code> class inherited from the <code>IPrefetchable&lt;&gt;</code> interface to provide automatic update of the object in the slot. If the <code>PXDatabase</code> slots contain a valid data object of the specified type saved under the key defined by the first parameter, the method returns this data object. Otherwise, the method does the following:</p> <ol style="list-style-type: none"> <li>1. Creates a new object of the specified type</li> <li>2. To create or update data in the object, invokes the <code>Prefetch</code> method with the <code>parameter</code> specified in the second parameter</li> <li>3. Saves this object in the slot under the key defined by the first parameter</li> <li>4. Returns the data object to the calling method</li> </ol> <p>The list of the table types specified in the <code>params</code> parameter is used to invalidate the slot in the case if any table of the list has been changed in the database. The use of this method is described below in the <a href="#">Automatically Updating Data in a PXDatabase Slot</a> section.</p>

The following example shows how you can use the `GetSlot<ObjectType> (string key, params Type[] tables)` method to cache data under the *MyData* key in the slot of the `PXDatabase` class.

```
...
Dictionary<string, string[]> dict =
    PXDatabase.GetSlot<Dictionary<string, string[]>>(
        "MyData", typeof(Table1), typeof(Table2), typeof(Table3));
lock ((System.Collections.ICollection)dict).SyncRoot
{
    ...
    List<string> myList = new List<string>();
    ...
    string key = "myListKey";
    dict[key] = myList.ToArray();
}
...
```

After the data object has been cached, you can access the object by using the following instruction.

```
Dictionary<string, string[]> dict =
    PXDatabase.GetSlot<Dictionary<string, string[]>>(
        "MyData", typeof(Table1), typeof(Table2), typeof(Table3));
```

You can clear a slot provided by the `PXDatabase` class by means of the following public static methods of the class.

Method	Description
<pre>public static void ResetSlot&lt;ObjectType&gt; (string key, params Type[] tables)</pre>	Sets to <i>null</i> the value of the slot that has the specified key.

Method	Description
<code>public static void ResetSlots()</code>	Sets to <i>null</i> the value of each slot that is provided by the <code>PXDatabase</code> class.

The following example shows how you can clear the slot created in the example above.

```
PXDatabase.ResetSlot<MyDataType>(
    "MyData", typeof(Table1), typeof(Table2), typeof(Table3));
```

### Automatically Updating Data in a `PXDatabase` Slot

If a data object that is to be cached depends on data in the database, we recommend that you inherit the object class from the `IPrefetchable<>` interface and develop in this class the `Prefetch` method, to provide automatic updating of data in the object. Then the `GetSlot<ObjectType, Parameter>(string key, Parameter parameter, params Type[] tables)` method of the `PXDatabase` class will use the `Prefetch` method to update the data in the slot, if required. (See the description of the method in [Caching Data in `PXDatabase` Slots](#).)

For example, suppose that you need to develop a data provider that selects data from multiple tables of the database and caches the data in `PXDatabase` slots. To do this, you can develop the provider class based on the following code.

```
public abstract class MyProvider : ProviderBase
{
    // Here you can add abstract definitions for all the methods of
    // the PXDatabaseMyProvider class
}

public class PXDatabaseMyProvider : MyProvider
{
    private class MyDataObject : IPrefetchable<PXDatabaseMyProvider>
    {
        public MyDataType MyData = new MyData();

        public void Prefetch(PXDatabaseMyProvider provider)
        {
            // Here you can implement the code to generate data of the MyData object.
        }
    }

    private MyDataObject MyDataObj
    {
        get
        {
            return PXDatabase.GetSlot<MyDataObject, PXDatabaseMyProvider>(
                "MYDATA_SLOT_KEY",
                this, typeof(Table1), typeof(Table2), typeof(Table3)
                /* ,... Add here the types of all tables, any change in which
                should make the slot invalid. */ );
        }
    }

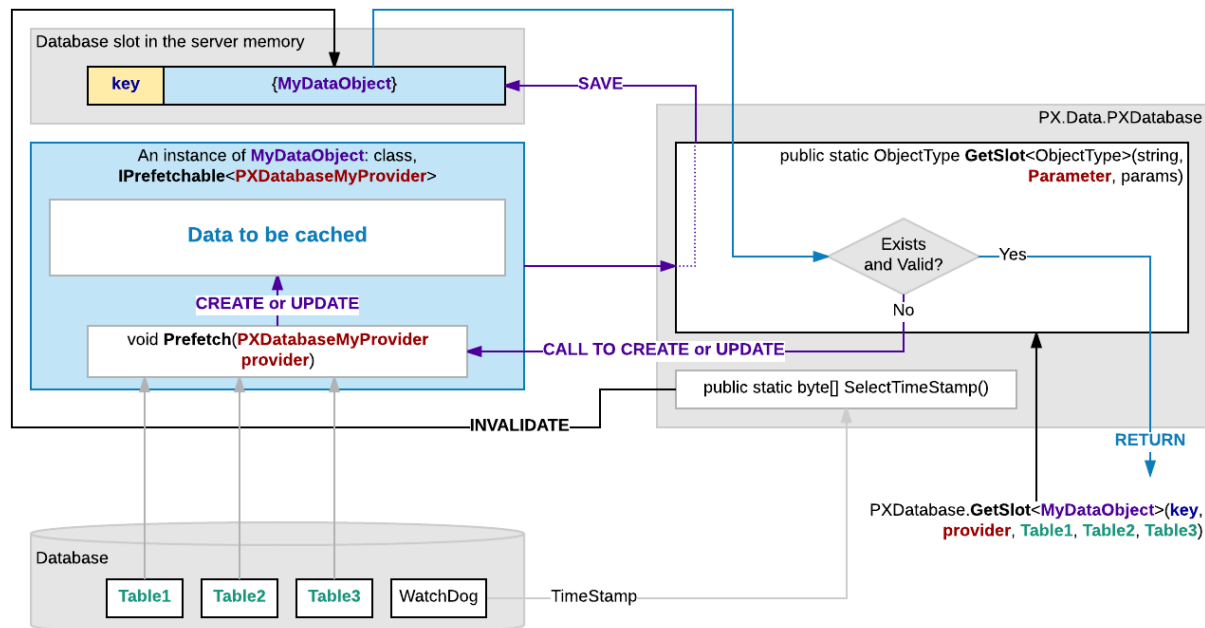
    // Here you need to add the code for all the methods that are defined
    // in the MyProvider abstract class.
    // These methods can be used to manage the MyData object.
    ...
}
```

The code above contains declarations of the following classes:

- The `MyProvider` abstract class, which derives from the `System.Configuration.Provider.ProviderBase` public abstract class and is used to define implementation of the `PXDatabaseMyProvider` class.
- The `PXDatabaseMyProvider` class, which contains the following:
  - The `MyDataObject` private class, which derives from the `IPrefetchable<PXDatabaseMyProvider>` interface and contains the following members:
    - The `MyData` data object to be cached
    - The `Prefetch` method, which creates or updates the data object
  - An implementation of the methods that are declared in the `MyProvider` abstract class and used to manage to the `MyData` object. To access the data object stored in the database slot, in these methods, you can use the `MyDataObject` property of the `PXDatabaseMyProvider` class, as the following instruction shows.

```
MyDataObject data = MyDataObj;
```

For the code above, the following diagram shows how the data object is cached and automatically updated in the `PXDatabase` slot.



**Figure: Automatic update of the cached data**



: If you have discovered that a `PXDatabase` slot returns legacy data, you can invoke the `SelectTimeStamp()` public static method of the `PXDatabase` class to invalidate all the `PXDatabase` slots that contain data obtained from the database tables that have been changed. Then the `GetSlot` method invokes the `Prefetch` method and updates the data in the slot.

# Implementing Business Logic

---

The topics in this part of the guide explain how to implement business logic of an application based on Acumatica Framework.

## In This Part

- [Working with Events](#)
- [Working with Attributes](#)
- [Working with Attachments](#)
- [Configuring the UI from the Back End](#)
- [Executing Code Asynchronously](#)
- [Localizing Applications](#)
- [Reusing Business Logic](#)

## Working with Events

---

The Acumatica Framework provides its own event model in which events related to the manipulation of data records and data fields are raised in a particular order within certain scenarios. An *event handler* is a method invoked by the Acumatica Framework once the corresponding event is raised.

By implementing event handlers, application developers can add business logic for the manipulation of data within business logic controllers (BLCs). This business logic includes the validation and calculation of field values, the management of data records (inserting, updating, or deleting), the checks for duplicate records, and the implementation of the presentation logic of the user interface.

## In This Chapter

- [Event Handlers](#)
- [Execution of Event Handlers](#)
- [Data Manipulation Scenarios](#)
- [Sequence of Events: Insertion of a Data Record](#)
- [Sequence of Events: Update of a Data Record](#)
- [Sequence of Events: Deletion of a Data Record](#)
- [Sequence of Events: Display of a Data Record](#)
- [Sequence of Events: Saving of Changes to the Database](#)
- [List of Events](#)

## Event Handlers

The Acumatica Framework raises events in the context of a graph. An event handler can be implemented in a graph, as well as in an attribute of a data field.

## Graph and Attribute Event Handlers

*Graph event handlers* are defined as methods in a business logic controller (BLC) class for a particular data access class (DAC) or a particular DAC field. See the topics in the [PX.Data Delegates](#) section in the API Reference for each event for an example of a graph event handler declaration.

*Attribute event handlers* are defined as methods in attribute classes. The logic of the event handlers is attached to all DAC objects or data fields annotated with these attributes. The attribute in which an attribute event handler is implemented must be derived from the `PXEventSubscriberAttribute` class and must implement the interface of the `IPX<EventName>Subscriber` form (where `<EventName>` is replaced with the name of the needed event), as shown in the following example.

```
// The attribute implements handlers for the FieldVerifying
// and RowPersisting events
public class MyAttribute : PXEventSubscriberAttribute,
                        IPXFieldVerifyingSubscriber,
                        IPXRowPersistingSubscriber
{
    public virtual void FieldVerifying(PXCache sender,
                                      PXFieldVerifyingEventArgs e)
    {
        ...
    }

    public virtual void RowPersisting(PXCache sender,
                                      PXRowPersistingEventArgs e)
    {
        ...
    }
}
```

## Execution of Event Handlers

In this topic, you can find information about how event handlers are executed and how to add and remove event handlers at run time.

### Execution of Event Handlers

All event handlers for a particular event share the same `PXCache` instance that has raised this event. The system creates a `PXCache` instance to control the modified data records of a particular data access class (DAC) type. The `PXCache` instance is always available as the first argument in an event handler; the second argument provides the specific data that corresponds to the event.

Once an event is raised, the order in which the associated event handlers are executed may differ. For some events, the chain of graph event handlers is executed before the attribute event handlers are; the attribute event handlers are executed only if the `Cancel` property of the event arguments doesn't equal `true` after the execution of the graph event handlers.

For other events, the attribute event handlers are executed first, and the graph event handlers are executed next. The topics in the [PX.Data Delegates](#) section in the API Reference for each event shows the order in which the system invokes handlers for a particular event.

### Dynamic Addition of Event Handlers

A business logic controller (BLC) includes the collections of graph event handlers for all events except `CacheAttached`. Each of these collections holds event handlers for a particular event and has the same name as the event. (You can find more information about how these collections are initialized in the [Initialization of an Event Handler Collection](#) section in this topic.) By using the methods of these collections, you can add and remove graph event handlers in code at run time.

A method added as an event handler must have the signature of a graph event handler, but doesn't need to follow the naming convention for graph event handlers. If you want to add a method as an

event handler, you invoke the `AddHandler<>()` method on the corresponding collection. For example, if the event is related to a row, it is invoked as follows.

```
RowEventName.AddHandler<DACName>(MethodName);
```

The event is invoked as follows if it is related to a field.

```
FieldEventName.AddHandler<DACName.fieldName>(MethodName);
```

When the `AddHandler<>()` method is invoked, event handlers are added to the collection as follows:

- Event handlers are added to the beginning of the collection for any event whose name ends with *ing* except the `RowSelecting` event.
- Event handlers are added to the end of the collection for any event whose name ends with *ed* and for the `RowSelecting` event.

To remove a handler, you should invoke the `RemoveHandler<>()` method.

### Initialization of an Event Handler Collection

On each round trip, the `PXGraph()` constructor does the following while it initializes a graph instance:

1. Creates the `Cashes`, `Views`, and `Actions` collections and other required collections. All of these collections are initially empty.
2. If the graph instance is being created on the Acumatica ERP server for the first time:
  - a. Obtains the metadata of this graph from the appropriate assembly (which is `PX.Objects` for most graphs in the application).
  - b. By using the metadata, emits the `InitializeDelegate` method, which is designed to subscribe event handlers and to initialize graph views, caches, and actions. To process the metadata for the fields declared in the graph, the constructor invokes the `PXGraph.ProcessFields` static method. To process the metadata of the methods that are defined in the graph, the constructor invokes the `PXGraph.ProcessMethods` static method.



:

The `ProcessMethods` method processes the metadata of the methods that are declared in the graph and all extensions of the graph. According to the naming convention for event handlers, the `_` symbol is a separator, so this method tries to split the name of each processed method into segments. If the name of the processed method has fewer than two segments or more than three segments, the processed method is skipped.

If the name of the processed method adheres to the naming convention for record event handlers, the processed method is added to the `_EventsRow` collection of the `PXCache<DACName>` object that is instantiated in the graph instance based on the DAC declaration. For example, the `SOOrder_RowSelected` event handler is added to the `_EventsRow` collection of the `PXCache<SOOrder>` cache object as an element with the `RowSelected` key.

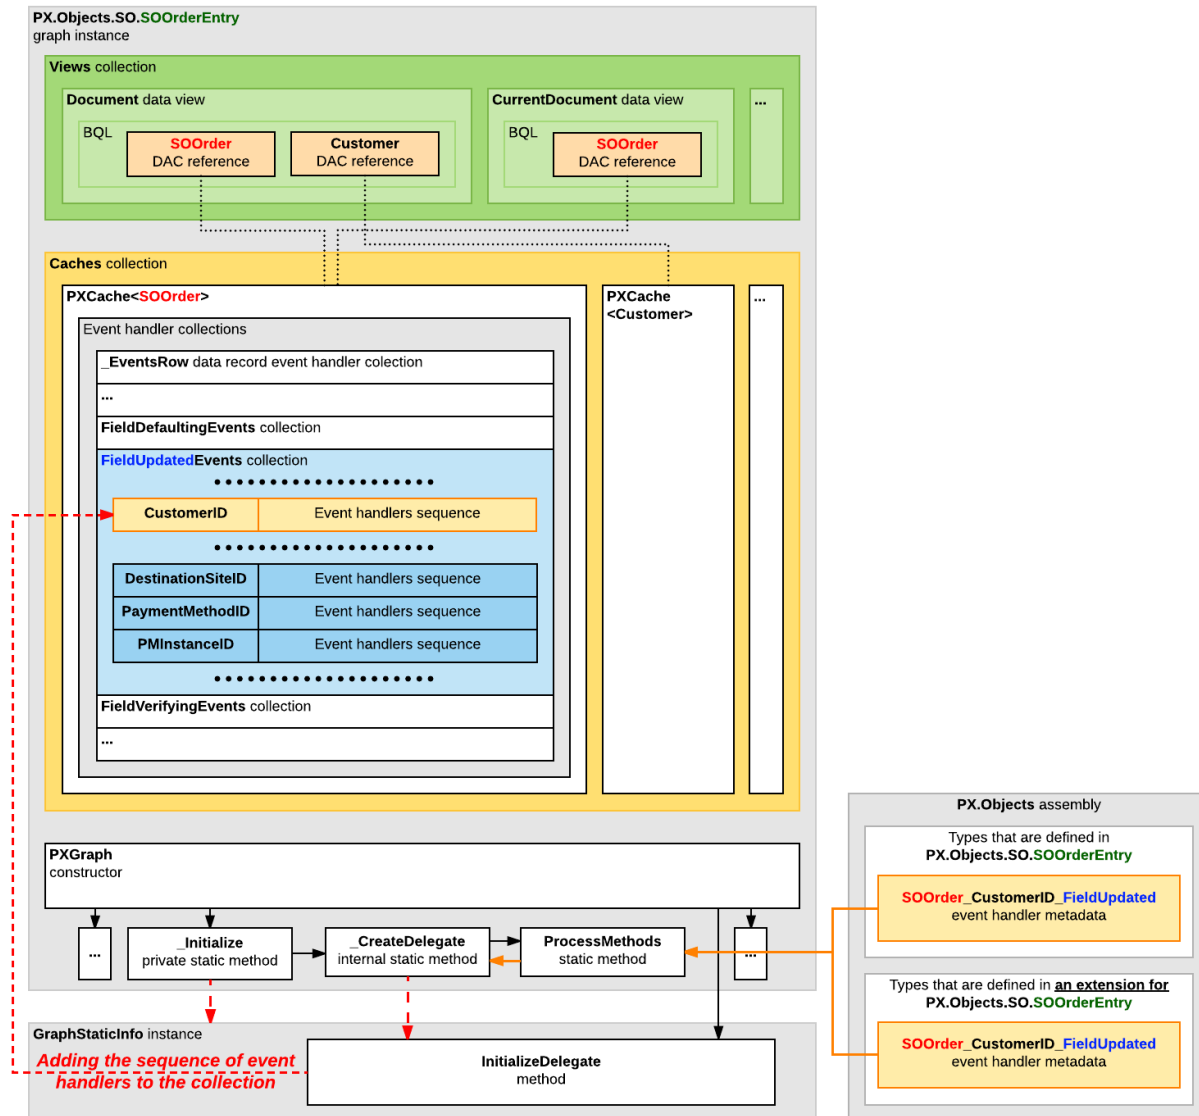
If the name of the processed method adheres to the naming convention for field event handlers, the processed method is added to the `EventNameEvents` collection of the `PXCache<DACName>` object. For example, the `SOOrder_CustomerID_FieldUpdated` event handler is added to the `FieldUpdatedEvents` collection of the `PXCache<SOOrder>` cache object as an element with the `CustomerID` key.

- c. Saves the graph metadata and the `InitializeDelegate` emitted method in the Acumatica ERP server memory as the `GraphStaticInfo` static object shared for the entire application instance.



- From the `GraphStaticInfo` static object, invokes the `InitializeDelegate` method, which initializes graph views, caches, and actions; the method also adds event handler delegates to the appropriate event handler collections of the relevant `PXCache` objects.

The following diagram shows how an instance of the `PX.Objects.SO.SOOrderEntry` graph uses the `PX.Objects` assembly metadata to add the `SOOrder_CustomerID_FieldUpdated()` event handler (described in the graph and graph extensions) to the `FieldUpdatedEvents` collection of the `PXCache<SOOrder>` cache object.



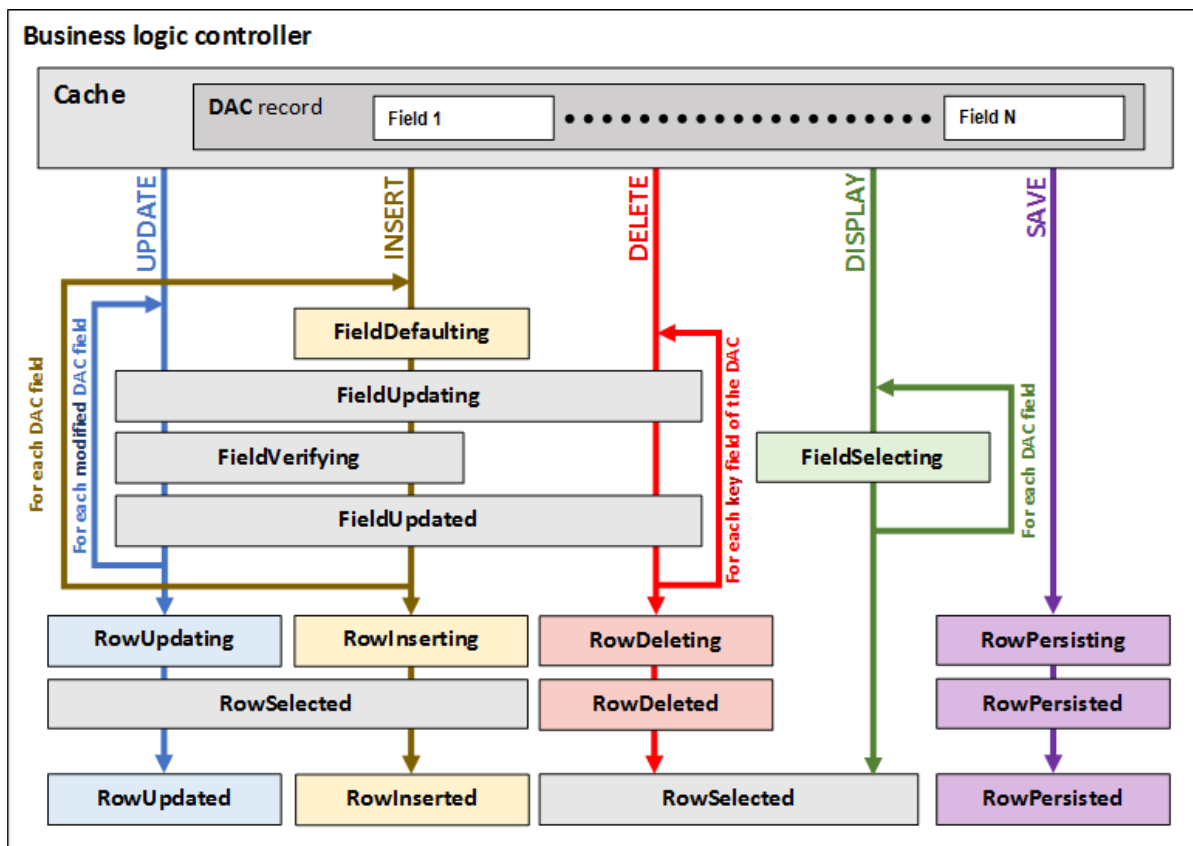
**Figure: Addition of an event handler to the appropriate collection**

In the collection, the `CustomerID` field name is used as a key, and the delegate of the event handler sequence for the field processing is used as a value.

## Data Manipulation Scenarios

Most events are raised within common scenarios related to the manipulation of data records. The scenarios are invoked by Acumatica Framework when users perform certain actions in the user interface, when the corresponding requests are made to the web services API, and when special methods are executed within the business logic controller (BLC).

The following diagram shows how different types of event handlers are invoked.



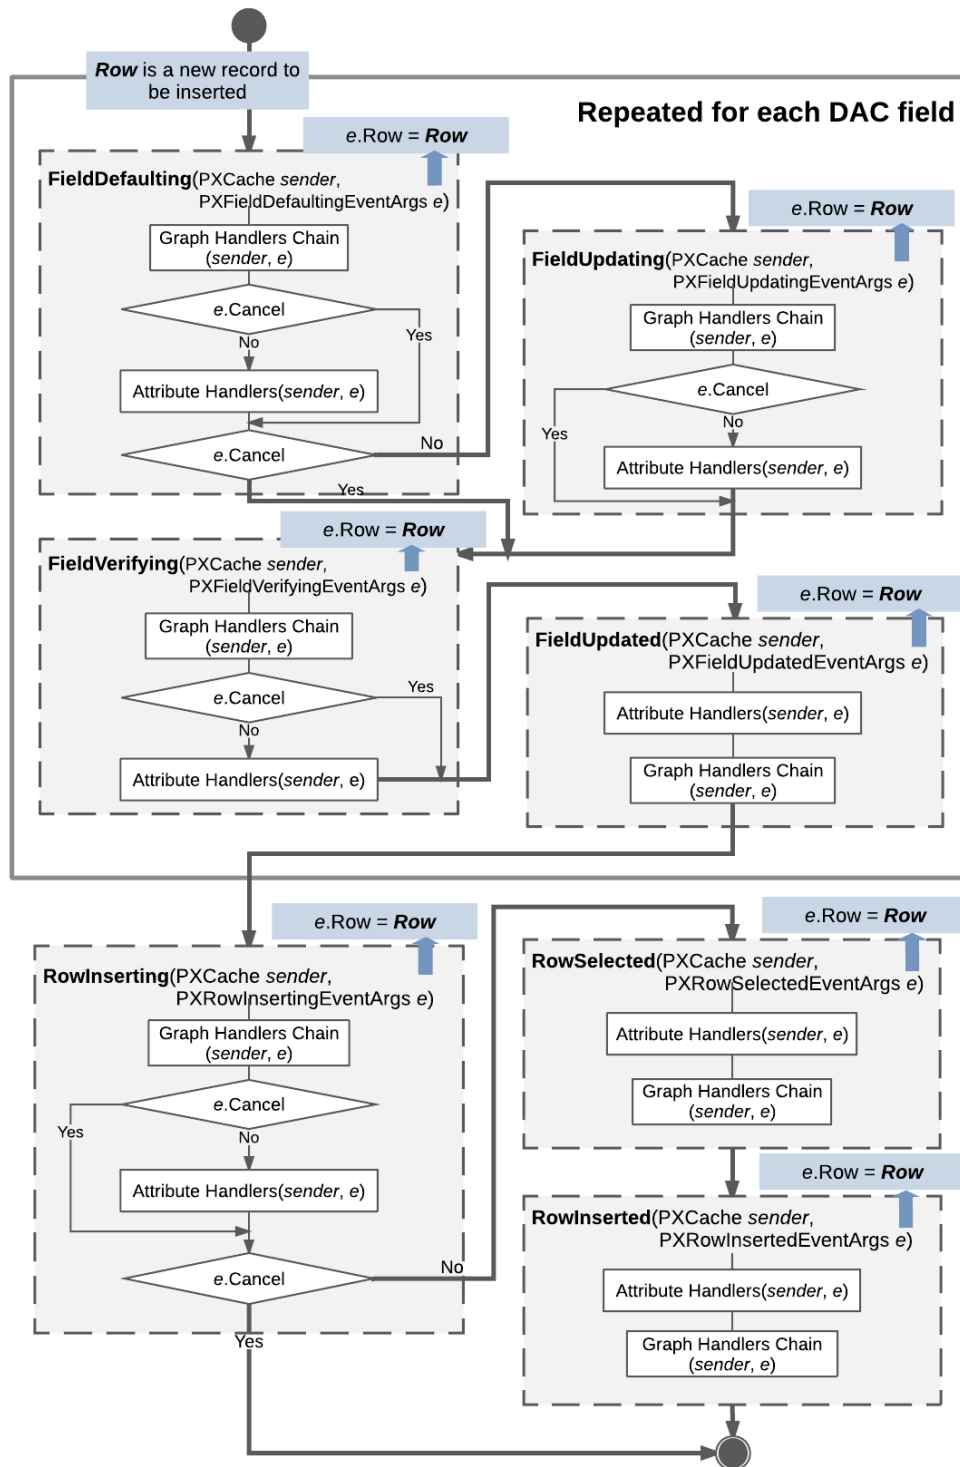
**Figure: Use of event handlers while the basic data operations are processed**

For details on how Acumatica Framework processes the basic data operations, see the following topics:

- [Sequence of Events: Insertion of a Data Record](#)
- [Sequence of Events: Update of a Data Record](#)
- [Sequence of Events: Deletion of a Data Record](#)
- [Sequence of Events: Display of a Data Record](#)
- [Sequence of Events: Saving of Changes to the Database](#)

## Sequence of Events: Insertion of a Data Record

The figure below illustrates the sequence of events raised during the insertion of a data record.



**Figure: Insertion of a data record**

The system inserts a data record—as an instance of a data access class (DAC)—when a user creates a new data record in the user interface, a request to insert a record is sent to the web services API, or the `Insert()` method of a data view is called in code. The data record is actually inserted into the `PXCache` object that corresponds to the DAC of the data record. The inserted data record has the `Inserted` status and is available through the `Inserted` and `Dirty` collections of the `PXCache` object.

When a data record is inserted, data field events are raised for each data field in the following order:

1. `FieldDefaulting`
2. `FieldUpdating` if the `e.Cancel` property equals `true`
3. `FieldVerifying`
4. `FieldUpdated`

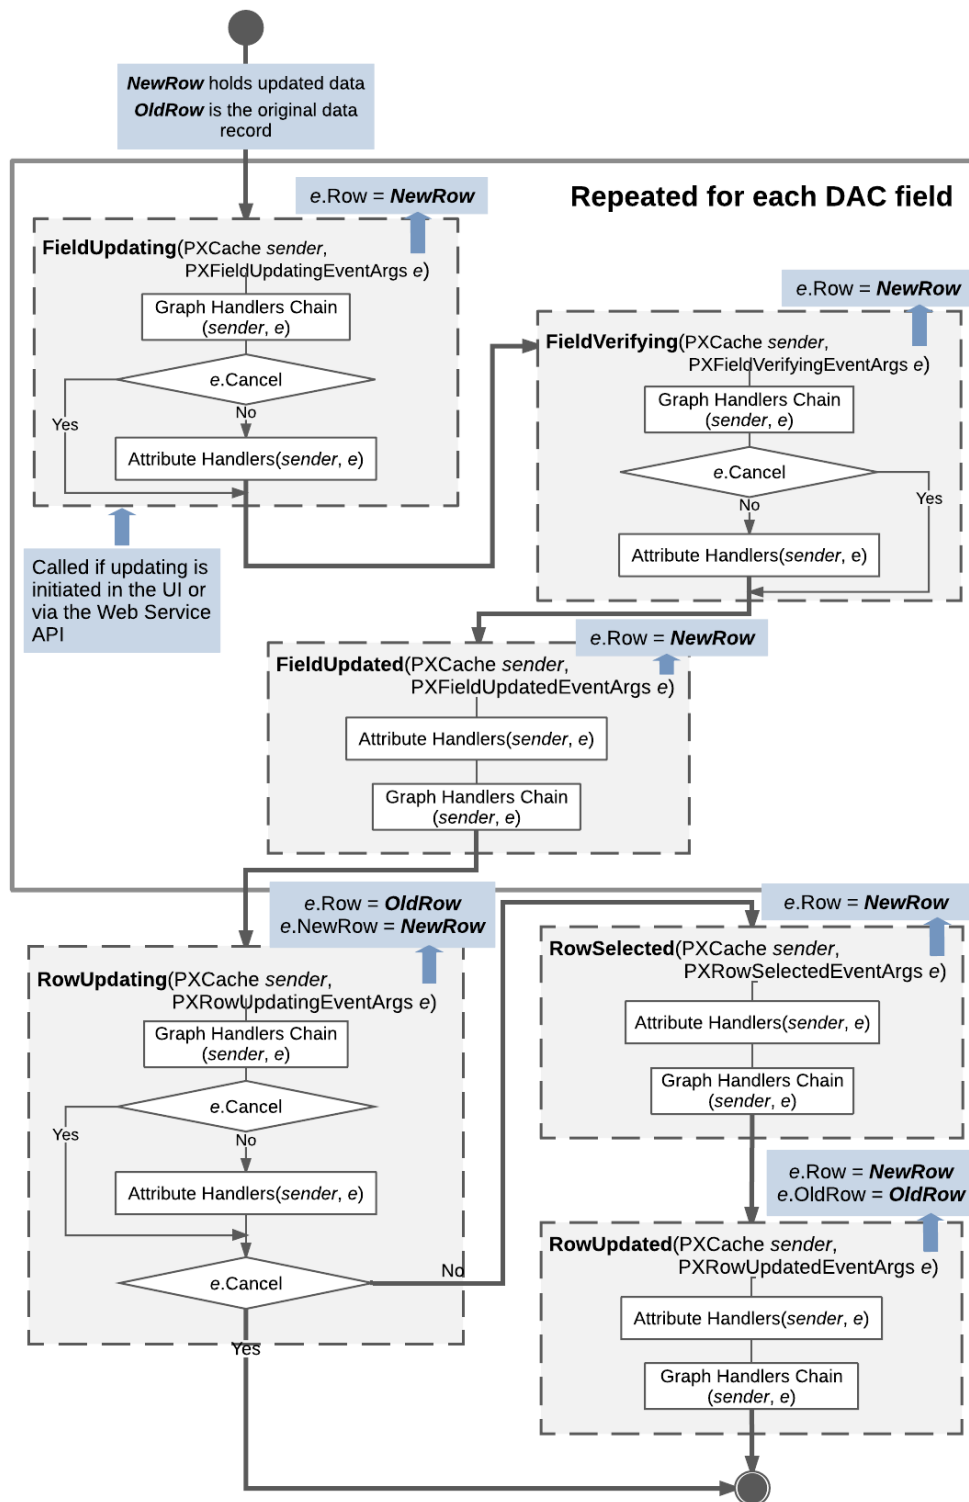
Next, the following data record events are raised:

1. `RowInserting` (If the `e.Cancel` property is `true`, no further events are raised.)
2. `RowSelected`
3. `RowInserted`

The instance of the inserted data record is available in the `e.Row` property of event arguments.

## Sequence of Events: Update of a Data Record

The figure below illustrates the sequence of events raised during the update of a data record.



**Figure: Update of a data record**

A data record is updated when a user modifies the data record on the user interface, the request is sent through the Web Service API, or the `Update()` method is invoked on the data view. Updated data records, which the system gives the `Updated` status, are later available through the `Updated` and `Dirty` collections of the appropriate `PXCache` object.

The `RowUpdating` event is fired before the update happens, while the `RowUpdated` event is fired after the update. The developer can handle these events and has access to the updated data record and the previous version of the data record that is kept in the `PXCache` object. The actual update happens between these two events when the data record is copied to the `PXCache` object.

When a data record is updated, the following data field events are raised for each updated data field:

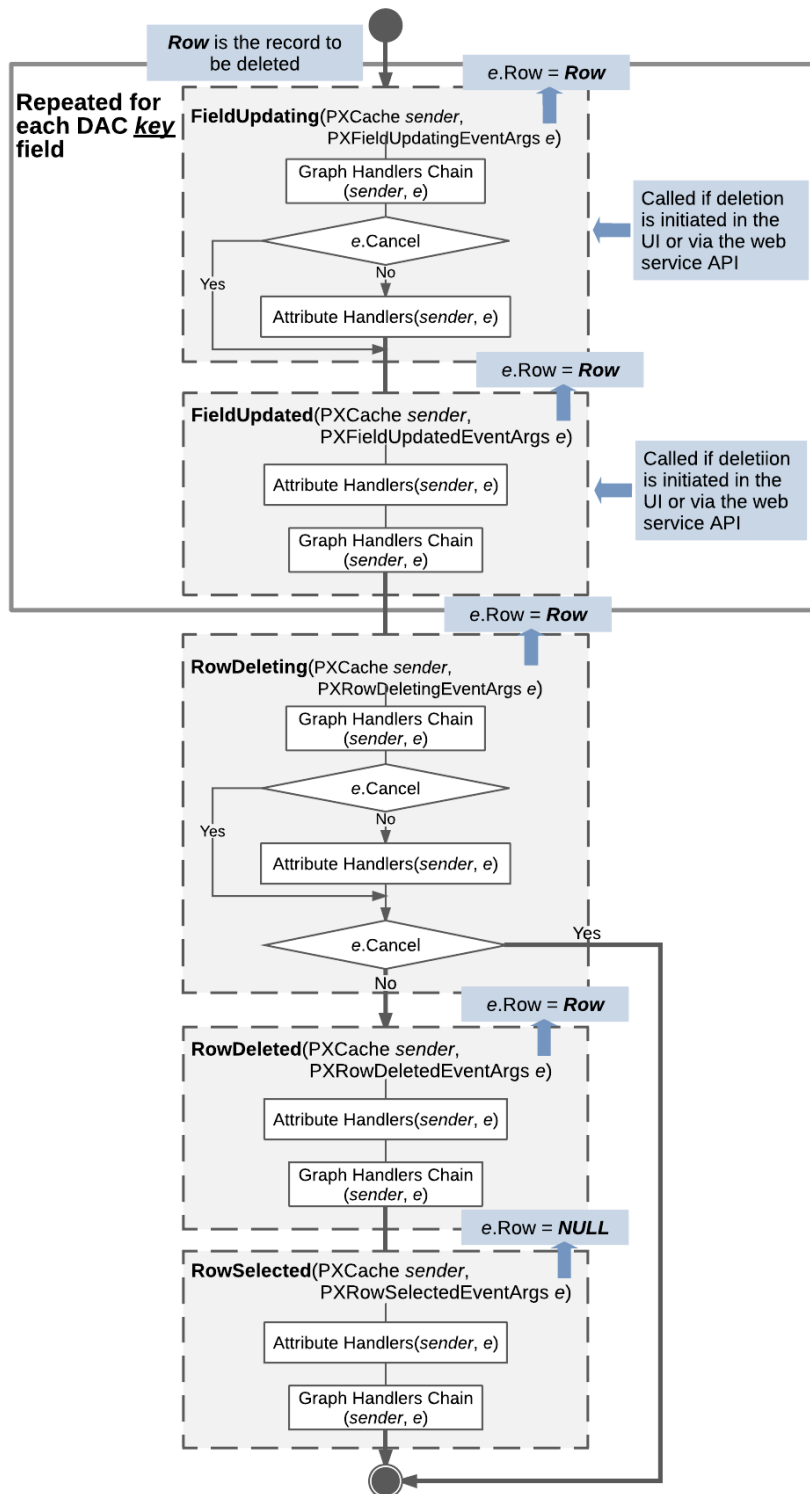
1. `FieldUpdating`
2. `FieldVerifying`
3. `FieldUpdated`

Next, data record events are raised as follows:

1. `RowUpdating` is raised. At this moment, in the `e` variable, which represents event data, `e.Row` holds the data record version from the cache, while `e.NewRow` holds the updated data record. You can still stop the update by throwing a `PXException` instance.
2. If `e.Cancel` doesn't equal `true`:
  - a. `RowSelected` is raised. Only the updated data record can be accessed through `e.Row`.
  - b. `RowUpdated` is raised. `e.Row` now holds the updated instance, while `e.OldRow` holds a copy of the old data record with the previous values.

## Sequence of Events: Deletion of a Data Record

The figure below illustrates the sequence of events raised during the deletion of a data record.



**Figure: Deletion of a data record**

A data record is deleted when a user deletes the record on the user interface, the deletion request is sent through the Web Service API, or the `Delete()` method of a data view is invoked in code. As a result of the deletion, the data record gets the `Deleted` status if it already exists in the database, or the `InsertedDeleted` status if the record has just been inserted into the `PXCache` object and deletion

from the database is not required. The data record is later available through the `Deleted` and `Dirty` collections of the `PXCache` object.

If the deletion has been initiated by a user on the UI or through the web services API, the following field events are raised for each key data field before any other events are raised:

1. `FieldUpdating`
2. `FieldUpdated`

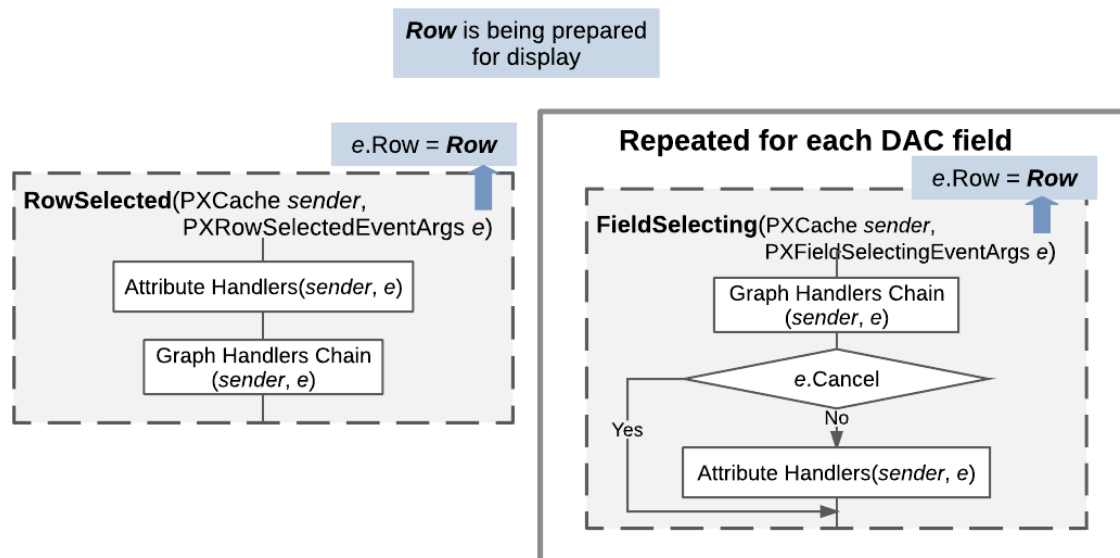
Next, regardless of how the deletion was initiated, data record events are raised as follows:

1. `RowDeleting` is raised. At this point, the developer can still stop the deletion by throwing a `PXException` instance. In the `e` variable representing event data, `e.Row` holds the data record being deleted.
2. If `e.Cancel` doesn't equal `true`:
  - a. `RowDeleted` is raised, and `e.Row` still holds the data record.
  - b. `RowSelected` is raised, and `e.Row` equals `NULL`.

## Sequence of Events: Display of a Data Record

Each time a data record is displayed in the user interface or retrieved through the Web Service API, the `RowSelected` and `FieldSelecting` events are raised for each data field. For both events, the `e.Row` property of event arguments holds the data record that is being displayed or retrieved.

The diagram below illustrates this process in more detail.




**Figure: Display of a data record**

## Sequence of Events: Saving of Changes to the Database

The following figure illustrates the sequence of events that are raised when a data record is saved.





**Figure: Save (commit) of a data record to the database**

While a user is inserting, updating, or deleting a data record, no changes are committed to the database. The system stores the modified data records in the session, and you can access them through the appropriate `PXCache` object. The system commits the changes to the database when the user clicks **Save** in the user interface, the save request is sent through the Web Service API, or `Actions.PressSave()` is invoked on the business logic controller (BLC) instance. In both cases, the `Persist()` method of the graph is invoked. The `Actions.PressSave()` method additionally checks that the `Save` action exists in the graph and is enabled. The `Save` action then invokes the `Persist()` method.

When changes are saved to the database, events are raised as follows:

1. `RowPersisting` is raised. At this moment, a database transaction has already been opened. If any of the handlers sets `e.Cancel` to `true`, the process will be canceled for the currently processed data record without an error being reported to the user. To cancel the process of committing changes and indicate the error to the user, you should throw the `PXException` exception.
2. If `e.Cancel` doesn't equal `true`:
  - a. `RowPersisted` is raised. The commit operation for the current data record (available through `e.Row` in the handler) is completed, but the transaction is still open: `e.TranStatus` equals `Open`.
  - b. `RowPersisted` is raised one more time, either with `e.TranStatus` equal to `Completed` (if all changes have been saved successfully) or with `e.TranStatus` equal to `Aborted` if an error has occurred and all changes have been canceled.

## List of Events

In this topic, you can find the list of all events by category. You can get more information about any of these events by navigating to the applicable topic in the API Reference.

### Data Field Events

- [\*PXFieldDefaulting\*](#)
- [\*PXFieldVerifying\*](#)
- [\*PXFieldUpdating\*](#)
- [\*PXFieldUpdated\*](#)
- [\*PXFieldSelecting\*](#)

### Data Record Events

- [\*PXRowSelected\*](#)
- [\*PXRowInserting\*](#)
- [\*PXRowInserted\*](#)
- [\*PXRowUpdating\*](#)
- [\*PXRowUpdated\*](#)
- [\*PXRowDeleting\*](#)
- [\*PXRowDeleted\*](#)

### Database-Related Events

- [\*PXCommandPreparing\*](#)
- [\*PXRowSelecting\*](#)
- [\*PXRowPersisting\*](#)
- [\*PXRowPersisted\*](#)

### Exception-Handling Event

- [\*PXExceptionHandling\*](#)

## Code Reuse Through Attributes

```

{
    DocTransaction old = e.OldRow as DocTransaction;
    DocTransaction trn = e.Row as DocTransaction;
    if ((trn != null) && (trn.TranQty != old.TranQty ||
        trn.UnitPrice != old.UnitPrice))
    {
        Document doc = Receipts.Current;
        if (doc != null)
        {
            doc.TotalAmt -= old.TranQty * old.UnitPrice;
            doc.TotalAmt += trn.TranQty * trn.UnitPrice;
            Receipts.Update(doc);
        }
    }
}

```

This logic can be used in multiple forms of the application, and therefore can be moved into an `Attribute` class. The attribute is used to annotate a data field in the data access class. Then it can be reused anywhere in the code, as in the example below.

```

public class DocTransaction : PX.Data.IBqlTable
{
    ...
    #region TotalAmt
    public abstract class totalAmt : PX.Data.IBqlField
    {
    }
    [PXDBDecimal(2)]
    [PXDefault(TypeCode.Decimal, "0.00")]
    [PXUIField(DisplayName = "Line Total", Enabled = false)]
    [DeltaMultiply(typeof(DocTransaction.tranQty), typeof(DocTransaction.unitPrice),
        typeof(Document.totalAmt))]
    public virtual decimal? ExtPrice { get; set; }
    #endregion
    ...
}

```

In this example, the logic of updating the receipt total on an update of the transaction is implemented inside the `DeltaMultiply` attribute. This logic is triggered after each update, delete, or insert operation on the `DocTransaction` data access class instance and updates totals on the receipt level in the appropriate `Document` data access class instance.

Acumatica Framework provides a wide range of predefined attributes that can be used for defining data types, database mapping, referential integrity, data format validation, and default values for the field. The following code shows an example of how you can implement the logic from the above example by using the predefined `PXFormula` attribute, which is used for implementing calculations of data fields.

```

public class DocTransaction : PX.Data.IBqlTable
{
    ...
    #region TotalAmt
    public abstract class totalAmt : PX.Data.IBqlField
    {
    }
    [PXDBDecimal(2)]
    [PXDefault(TypeCode.Decimal, "0.00")]
    [PXUIField(DisplayName = "Line Total", Enabled = false)]
    [PXFormula(typeof(Mult<DocTransaction.tranQty, DocTransaction.unitPrice>),
        typeof(SumCalc<Document.totalAmt>))]
    public virtual decimal? ExtPrice { get; set; }
    #endregion
    ...
}

```

Because the data access classes are shared within an application, formatting, custom logic, and any constraints implemented in attributes are reused in each business logic controller that utilizes each data

access class. By reusing code through attributes, you can move shared application functionality into attributes and avoid code duplication, while still enforcing application integrity.

## Mandatory Attributes

In this topic, you can learn about the mandatory attributes of data access class (DAC) fields and actions.

### Mandatory Attributes of DAC Fields

For each field defined in a DAC, you must specify the following attributes:

- A data type attribute, which is either a bound field data type attribute that binds the field to a database column of a particular data type, or an unbound field data type attribute that indicates that the field is unbound. For lists of these attributes, see [Bound Field Data Types](#) and [Unbound Field Data Types](#).
- The `PXUIField` attribute, which is mandatory for all fields that are displayed in the user interface. For details on the `PXUIField` attribute, see [UI Field Configuration](#).

The example below demonstrates a declaration of a DAC field bound to a database column and displayed in the user interface.

```
// The data access class for the POReceiptFilter database table
[Serializable]
public partial class POReceiptFilter : IBqlTable
{
    ...
    // The type declaration of a DAC field
    public abstract class receiptType : PX.Data.IBqlField
    {
    }
    // The value declaration of a DAC field
    // Put attributes before this declaration
    [PXDBString(2, IsFixed = true)]
    [PXUIField(DisplayName = "Type", Enabled = false)]
    public virtual String ReceiptType { get; set; }
    ...
}
```

### Mandatory Attributes of Actions

A declaration of a method that implements an action in a business logic controller must be preceded with the `PXButton` attribute or one of its successors and the `PXUIField` attribute. For details on the `PXUIField` attribute, see [UI Field Configuration](#).

The example below demonstrates a declaration of an action handler.

```
public PXAction<SalesOrder> ViewDocument;

[PXUIField(DisplayName = "View Document",
           MapEnableRights = PXCachedRights.Select,
           MapViewRights = PXCachedRights.Select)]
[PXButton]
public virtual IEnumerable viewDocument(PXAdapter adapter)
{
    ...
}
```

## Use of Attributes

To apply the attribute business logic to an entity, you should place the attribute on the entity declaration. At run time, you can call the static methods of a particular attribute to adjust the attribute's behavior.

### Attributes on the Entity Declaration

An attribute may be placed on a declaration of a class or a class member, with or without parameters. The parameters that are possible for an attribute depend on the constructor parameters and the properties defined in the attribute. The parameters of the selected constructor are placed first without names, and the named property settings follow them, as shown in the following example.

```
[PXDefault(false, PersistingCheck = PXPersistingCheck.Nothing)]
public virtual Boolean? Released { get; set; }
```

Here the `PXDefault` attribute is created with the constructor that takes a Boolean-type parameter (set to `false`). Additionally, the `PersistingCheck` property is specified.

### Setting of Attribute Properties at Run Time

You should call static methods defined in the attribute class to change the properties at run time. The static methods can affect a single attribute instance or multiple attribute instances related to a specific data record or all data records in a particular cache object. The following example shows an invocation of a static method.

```
PXUIFieldAttribute.SetVisible<APInvoice.curyID>(cache, doc, true);
```

When calling such a method, you typically specify the cache object, a data record related to this cache object, and the data access class (DAC) field. The method affects the attribute instance created for this field and the specified data record. If you pass `null` as the data record, the method affects attribute instances related to all data records in the specified cache object.

## Bound Field Data Types

The following attributes bind a data access class (DAC) field to a database column of a specific type.

Attribute	C# Data Type	Database Data Type	Comment
<a href="#">PXDBBool</a>	bool?	bit	Boolean value
<a href="#">PXDBByte</a>	byte?	tinyint	One-byte integer value
<a href="#">PXDBDate</a>	DateTime?	datetime or smalldatetime	Date and time
<a href="#">PXDBTime</a>	DateTime?	smalldatetime	Time without date
<a href="#">PXDBDateAndTime</a>	DateTime?	datetime or smalldatetime	Date and time values represented by separate input controls in the user interface
<a href="#">PXDBDecimal</a>	decimal?	decimal	Sixteen-byte floating point numeric value with a specific precision
<a href="#">PXDBDecimalString</a>	decimal?	decimal	A decimal value with a value selected by a user from the list of predefined values
<a href="#">PXDBDouble</a>	double?	float	Eight-byte floating point value
<a href="#">PXDBFloat</a>	float?	real	Four-byte floating point value

Attribute	C# Data Type	Database Data Type	Comment
<i>PXDBGuid</i>	Guid?	uniqueidentifier	Sixteen-byte unique value
<i>PXDBIdentity</i>	int?	int	Four-byte auto-incremented integer value
<i>PXDBLongIdentity</i>	int64?	bigint	Eight-byte auto-incremented integer value
<i>PXDBShort</i>	short?	smallint	Two-byte integer value
<i>PXDBInt</i>	int?	int	Four-byte integer value
<i>PXDBLong</i>	int64?	bigint	Eight-byte integer value
<i>PXDBString</i>	string	char, varchar, nchar, or nvarchar	Common string
<i>PXDBEmail</i>	string	nvarchar	Email address
<i>PXDBLocalizedString</i>	string	char, varchar, nchar, or nvarchar	Localized string
<i>PXDBCryptString</i>	string	char, varchar, nchar, or nvarchar	Encrypted string
<i>PXDBText</i>	string	nvarchar or varchar	Text
<i>PXDBTimeSpan</i>	int?	int	Date and time value represented by minutes passed from 01/01/1900
<i>PXDBTimeSpanLong</i>	int?	int	Duration in time as a number of minutes
<i>PXDBTimestamp</i>	byte[]	timestamp	Eight-byte unique, automatically generated binary numbers within a database
<i>PXDBBinary</i>	byte[]		Arbitrary array of bytes
<i>PXDBVariant</i>	byte[]	variant	Variant data type

Acumatica Framework also includes other attributes that are used in special cases to bind a DAC field to database columns.

## Unbound Field Data Types

The following table contains unbound field attributes. You use unbound attributes when you define a custom fields of your own that are not bound to any database fields.

Attribute	C# Data Type	Comment
<i>PXBool</i>	bool?	Boolean value
<i>PXByte</i>	byte?	One-byte integer value
<i>PXDate</i>	DateTime?	Date and time

Attribute	C# Data Type	Comment
<i>PXDateTime</i>	DateTime?	Date and time values represented by separate input controls in the user interface
<i>PXDecimal</i>	Decimal?	Sixteen-byte floating point numeric value with a specific precision
<i>PXDouble</i>	double?	Eight-byte floating point value
<i>PXFloat</i>	float?	Four-byte floating point value
<i>PXGuid</i>	Guid?	Sixteen-byte unique value
<i>PXShort</i>	short?	Two-byte integer value
<i>PXInt</i>	int?	Four-byte integer value
<i>PXLong</i>	int64?	Eight-byte integer value
<i>PXString</i>	string	String of characters
<i>PXTimeSpan</i>	int?	Date and time value represented by minutes passed from 01/01/1900
<i>PXTimeSpanLong</i>	int?	Duration in time as a number of minutes
<i>PXVariant</i>	byte[]	Random array of bytes

## UI Field Configuration

By using the `PXUIField` attribute, you can configure the layout of input controls and buttons. The attribute is mandatory for all data access class (DAC) fields displayed in the user interface.

### Setting of the `PXUIField` Attribute

You can add the `PXUIField` attribute in the following ways:

- To a DAC field declaration to configure the field input control, as shown in the following example

```
[PXDBDate()]
[PXUIField(DisplayName = "Pay Date")]
public virtual DateTime? PayDate { get; set; }
```

- To the declaration of a method that implements an action to configure the action button, as shown in the following sample code

```
[PXUIField(DisplayName = "View Document",
    MapEnableRights = PXCacheRights.Select,
    MapViewRights = PXCacheRights.Select)]
[PXButton]
public virtual IEnumerable viewDocument(PXAdapter adapter)
{
    ...
}
```

The attribute's properties determine the control layout in the user interface. You can specify the display name, specify whether the control is visible and available, set the error marker, and specify the access rights to view and use the control.



## Setting of the Properties of the PXUIField Attribute at Run Time

You can use the static methods (such as `SetEnabled` and `SetRequired`) of the `PXUIFieldAttribute` class to set the properties of a control at run time. The `PXUIFieldAttribute` static methods can be called in the graph constructor or the `RowSelected` event handlers.



: The `RowSelected` event handler is raised when the user interface controls are prepared to be displayed. This happens each time the webpage sends a request to the server.

If you want to modify the `Visible`, `Enabled`, and `Required` properties for all detail rows in a grid, you use the `RowSelected` event handler of the primary view DAC. If you want to set the `Enabled` property of a field in particular row in a grid, you use the `RowSelected` event handler of the DAC that includes this field.

If the grid column layout is configured at run time, you set the `data` parameter of the corresponding method to `null`. This indicates that the property should be set for all data records shown in the grid. If a specific data record is passed to the method rather than `null`, the method invocation has no effect.



: If you want to change the `Visible` or `Enabled` property of `PXUIFieldAttribute` for a button at run time, you use the corresponding static methods of `PXAction`. You usually use these methods in the `RowSelected` event handler of the primary view DAC.

## Default Values

You can set the default values to data access class (DAC) fields by using the following attributes:

- `PXDefault`: This attribute sets the default value and validates the field value when the value is saved to the database. The following attributes are derived from the `PXDefault` attribute:
  - `PXUnboundDefault` behaves in the same way as `PXDefault` does, but the default value is assigned to the field when a data record is retrieved from the database.
  - `PXDefaultValidate`.
- `PXDBDefault`: This attribute sets the default value by using the value of some source field, and updates the value if the source field value changes in the database before the data record is saved.

### PXDefault Attribute

The `PXDefault` attribute provides the default value for a DAC field. The default value is assigned to the field when the cache raises the `FieldDefaulting` event. This happens when a new row is inserted in code or through the user interface.

A value specified as a default can be a constant or the result of a BQL query. If you provide a BQL query, the attribute executes it on the `FieldDefaulting` event. You can specify both a constant and a BQL query; the attribute first executes the BQL query and then uses the constant if the BQL query returns an empty set. If you provide a DAC field as the BQL query, the attribute retrieves the value of this field from the `Current` property of the cache object. The attribute uses the cache object of the DAC type in which the field is defined.

The `PXDefault` attribute also checks that the field value is not `null` before saving a record to the database. You can adjust this behavior by using the `PersistingCheck` property. Its value indicates whether the attribute should verify that the value is not `null`, verify that the value is not `null` or a blank string, or not perform any verification.

The attribute can redirect the error that happened on the field to another field if you set the `MapErrorTo` property.

You can use the static methods of the attribute to change the attribute properties for a particular data record in the cache or for all data records in the cache.

## Differences

You usually set the default value to a DAC field by using the `PXDefault` attribute. You can set a constant as the default value or provide a BQL query to obtain a value from the database or data records from the cache. The default value is assigned to the field when a data record that includes this field is inserted into the cache.

You can use the `PXDefault` attribute just to make the field mandatory for input by using the attribute without parameters.

The `PXDefault` attribute is not suitable when the default value is retrieved from a field that can be auto-generated by the database (such as the identity field). In this case, you should use the `PXDBDefault` attribute. It updates the value assigned to the field as the default with the value generated by the database.

For example, if you implement a master-detail relationship, you should use the `PXDBDefault` attribute to bind the detail data record fields to the master data record key fields. If the master data record is new, its identity field is set to a real value by the database when the master record is saved. So if a detail data record is created before the master data record is saved, the detail data record field is set to the temporary value of the master identity field. However, the `PXDBDefault` attribute replaces the temporary value with the real value when the detail data record is saved to the database.

You can use the `PXUnboundDefault` attribute to set the default value to an unbound field. The value is assigned when a data record is retrieved from the database (on the `RowSelecting` event).

## Complex Input Controls

You can use attributes to configure complex input controls, such as drop-down lists and lookup controls.

### Drop-Down Lists

You can use the following attributes to configure a drop-down list that represents a data access class (DAC) field in the user interface:

- `PXStringList`: Configures a drop-down list from which a user can select from a fixed set of strings.
- `PXIntList`: Configures a drop-down list where a user can select from a fixed set of values. The control displays strings, while the field is assigned the integer value corresponding to the selected string.
- `PXDecimalList`: Configures a drop-down list where a user can select from a fixed set of strings converted to decimal values.
- `PXImagesList`: Configures a drop-down list where a user can select from a fixed set of images.
- `PXDBIntList`: Configures a drop-down control for an integer field. The values and labels for the drop-down control are retrieved from the specified database table.
- `PXDBStringList`: Configures a drop-down control for a string field. The values and labels for the drop-down control are retrieved from the specified database table.

### Lookup Controls

You can use the following attributes to configure a lookup control that represents a field in the user interface:

- `PXSelector`: Defines a lookup control for a DAC field that references a data record from a particular table by holding its key.
- `PXCustomSelector`: Serves as the base class to derive custom attributes used to configure lookup controls.

- **PXRestrictor**: Adds a restriction to a BQL command that selects data for a lookup control, and displays an error message when the value entered does not fit the restriction. The attribute works only with **PXSelector** and cannot be used with **PXCustomSelector**.

### Segmented Key Controls

A segmented key value is a string value that identifies a data record in the system and consists of one segment or multiple segments. A segmented key is an entity that is identified by a string (referred to as a *dimension*) and associated with segments. For each segment, you can define the list of possible values. You can create a new segment when the data records identified by the segmented key already exist in the database.

You can use the following attributes to configure a control to input a segmented key value in the user interface:

- **PXDimension**: Defines an input control that formats the input as a segmented key value and displays the list of allowed values for each key segment.
- **PXDimensionSelector**: Defines an input control that combines the functionality of the **PXDimension** attribute and the **PXSelector** attribute. A user can view the data set defined by the attribute and select a data record from this data set to assign its segmented key value to the field or to replace it with the surrogate key.
- **PXDimensionWildcard**: Behaves similarly to the **PXDimensionSelector** attribute, but also allows the ? character to be treated as a wildcard.

## Referential Integrity and Calculations

You can use the following attributes to implement referential integrity and perform calculations at run time:

- **PXParent**: Creates a reference to a parent data record. By default, when the parent data record is deleted, all child data records that reference it are also deleted. (You can change this behavior by using the **LeaveChildren** property.)
- **PXFormula**: Calculates a field from other fields of the same data record or sets an aggregation expression to calculate a parent data record field from its child data record fields. These calculations happen only when the parent field is changed or a new record is inserted.
- **PXUnboundFormula**: Calculates the value from the child data record fields and aggregates all such values computed for the child data records into the parent data record field. These calculations happen on every data reading.
- **PXDBChildIdentity**: Indicates that a DAC field references an auto-generated key field from another table, and ensures that the field value is correct after changes have been committed to the database.
- **PXLineNbr**: Generates unique line numbers that identify child data records in the parent-child relationship.

Note that all the attributes in the list above add server-side logic used at run time. The referential integrity is implemented on the server side, as are the calculations. See the [Ad Hoc SQL for Fields](#) topic for information about the attributes that enable the calculation of fields on the database side.

## Ad Hoc SQL for Fields

The following attributes define the database-side calculation of data access class (DAC) fields that are not bound to particular database columns:

- **PXDBCalced**: Defines an SQL expression that calculates an unbound field from the fields of the same DAC whose values are taken from the database.

- **PXDBScalar**: Defines an SQL subrequest that retrieves a value for an unbound DAC field. The subrequest can retrieve data from any bound field in any DAC.

The attributes add the provided expression and the subrequest into the SQL query that selects data records of the given DAC.

## Audit Fields

The following attributes are placed on data access class (DAC) fields used for data audit:

- **PXDBLastChangeDateTime**: Maps a DAC field to the database column and automatically sets the field value to the data record's last modification date and time.
- **PXDBCreatedByID**: Maps a DAC field to a database column and automatically sets the field value to the ID of the user who created the data record.
- **PXDBCreatedByScreenID**: Maps a DAC field to a database column and automatically sets the field value to the string ID of the application screen from which the data record was created.
- **PXDBCreatedDateTime**: Maps a DAC field to a database column and automatically sets the field value to the data record's creation date and time.
- **PXDBCreatedDateTimeUtc**: Maps a DAC field to a database column and automatically sets the field value to the data record's creation UTC date and time.
- **PXDBLastModifiedByID**: Maps a DAC field to a database column and automatically sets the field value to the ID of the user who last modified the data record.
- **PXDBLastModifiedByScreenID**: Maps a DAC field to a database column and automatically sets the field value to the string ID of the application screen on which the data record was last modified.
- **PXDBLastModifiedDateTime**: Maps a DAC field to a database column and automatically sets the field value to the data record's last modification date and time.
- **PXDBLastModifiedDateTimeUtc**: Maps a DAC field to a database column and automatically sets the field value to the data record's last modification date and time in UTC.

The framework binds the DAC fields to database columns and automatically assigns field values.

## Data Projection

The following attributes implement the projection of data from one table or multiple tables into a single data access class (DAC):

- **PXProjection**: Binds the DAC to a random data set. The attribute thus defines a named view, but is implemented by the server side rather than by the database.
- **PXExtraKey**: Indicates that the field implements a relationship between two tables. The use of this attribute enables the update of the referenced table when the projection is updated.

## Access Control

The group mask value indicates the access rights a user should have in order to use a data record. To be able to set access rights for particular data records, you should use the **PXDBGGroupMask** attribute to mark the data access class (DAC) field that holds the group mask value.

On a substitute form, to define the inheritance of access rights for an action that is implemented in the corresponding entry form, you can use the **PXEntryScreenRights** attribute.

## Notes

By using the **PXNote** attribute, you can give users the ability to attach text notes, files, and activity items to data records, and to search for an entity by using the full-text search index.

You should use the `PXNote` attribute in the data access class of these data records to mark the field that stores the identifier of a note in the `Note` table. Notes are used to attach text to a data record. This text is stored in the note data record in the `Note` table. Additionally, you can attach files or other entities to a data record through a note. This feature is implemented through additional tables that store the identifiers of a note and the attached entity.

The `PXNote` attribute can also be configured to save the specified table fields in a note. In this case, by using the Acumatica Framework application website search, the user will be able to search the data records by the values saved in the note.

## Report Optimization

The value of an unbound data access class (DAC) field can be calculated in the property getter. The calculation can involve other fields of the same DAC. However, when the value of the DAC field is requested, other fields are not guaranteed to be calculated or assigned their values. These situations are normal when the integration services are used, copy-paste functionality is used, or the field is used in reports.

To ensure that the fields referenced in the property getter have values when it is executed, you should use the `PXDependsOnFields` attribute.

## Attributes on DACs

You can place the following attributes on the data access class (DAC) declaration:

- `PXPrimaryGraph`: Specifies the graph that is used by default to edit a data record.
- `PXCacheName`: Specifies the user-friendly name of the DAC. The name is displayed in the user interface.
- `PXTable`: Binds a DAC that derives from another DAC to the table having the name of the derived table. Without the attribute, the derived DAC will be bound to the same table as the DAC that starts the inheritance hierarchy.
- `PXAccumulator`: Updates the values of a data record in the database according to the policies specified in the attribute parameters.
- `PXHidden`: Gives the developer the ability to hide a DAC, graph, or view from the selectors of DACs and graphs, and from generic inquiries, reports, and the web services API.

The `PXProjection` and `PXTable` attributes can also mark a DAC. See [Data Projection](#) for more details.

## Action Attributes

You use the following attributes to configure a button that will represent an action in the user interface:

- `PXButton`: Serves as the base attribute for all other attributes, which give you the ability to configure buttons. The successor attributes only set the base class properties to specific values.
- `PXSaveButton`
- `PXSaveCloseButton`
- `PXCancelButton`
- `PXCancelCloseButton`
- `PXInsertButton`
- `PXDeleteButton`
- `PXFirstButton`
- `PXPreviousButton`
- `PXNextButton`

- [\*PXLastButton\*](#)
- [\*PXSendMailButton\*](#)
- [\*PXReplyMailButton\*](#)
- [\*PXForwardMailButton\*](#)
- [\*PXTemplateMailButton\*](#)
- [\*PXLookupButton\*](#)
- [\*PXProcessButton\*](#)

Also, you can use the following attributes:

- [\*PXUIField\*](#): To configure the button layout and set access rights
- [\*PXEntryScreenRights\*](#): On a List as Entry Point screen, to define the inheritance of access rights for an action that is implemented in the appropriate entry screen

## Attributes on Data Views

You can place the following attributes on the declaration of a data view in a graph:

- [\*PXFilterable\*](#): Adds the control that makes it possible for a user to create filters and save them in the database. The control is added to the grid that uses the data view to retrieve data.
- [\*PXImport\*](#): Adds the grid toolbar button that a user clicks to load data from the file to the grid. The attribute is placed on the data view the grid uses to retrieve the data.
- [\*PXHidden\*](#): Hides the data view from the selectors of data access classes (DACs) and graphs, and from the web service API clients.
- [\*PXCopyPasteHiddenView\*](#): Indicates that the cache corresponding to the primary DAC of the data view is not copied when the copy-paste feature is used on the webpage.
- [\*PXCopyPasteHiddenFields\*](#): Indicates that the specific fields of the primary DAC of the data view are not copied when the copy-paste feature is used on the webpage.

## Working with Attachments

---

In Acumatica Framework-based and customized Acumatica ERP applications, you can attach files to the records displayed on the webpages. This chapter describes how to attach files to the records and display the attached image files on the webpages.

### In This Chapter

- [\*To Allow Attachments to a Particular Webpage\*](#)
- [\*To Display an Attached Image on the Webpage\*](#)

## To Allow Attachments to a Particular Webpage

Users of Acumatica Framework-based and Acumatica ERP applications can attach files, notes, and activities to master records displayed on webpages. They can also attach files and notes to the detail records, which are displayed in table rows on webpages. In this topic, you can find information about how to allow attachments to a particular webpage and the table rows on a specific webpage.

For more information about how a user can attach files and notes to webpages, see [\*To Attach a File to a Record\*](#) and [\*To Attach a Note to a Record\*](#). For information about activities, see [\*Tasks, Events, and Activities\*](#).



: The extensions of the files that can be uploaded to webpages must be registered on the [File Upload Preferences](#) (SM202550) form. If the required file types are not registered on this form, you have to add them and save your changes. On this form, you can also define the maximum size of an uploaded file (in kilobytes).

### To Allow Attachments to a Particular Webpage

1. In the data access class (DAC) that provides data for the webpage, add the `NoteID` field, as the following code shows.

```
#region NoteID
public abstract class noteID : PX.Data.IBqlField { }

[PXNote]
public virtual Guid? NoteID { get; set; }
#endregion
```



: The database table that corresponds to the DAC must contain the `NoteID` column with the `uniqueidentifier` data type.

2. Rebuild the project.

Once you have added the `NoteID` field to the DAC and rebuilt the project, the following buttons appear on the title bar of the webpage:

- **Notes**, which users click to attach notes to the webpage
- **Files**, which users click to attach files to the webpage

3. Optional: To change whether each of these elements is displayed on the title bar, in the ASPX code of the webpage, specify the values of the following properties of the `PXFormView` control:

- `NoteIndicator`: Indicates (if set to `True`) that the **Notes** button is displayed on the title bar.
- `FilesIndicator`: Indicates (if set to `True`) that the **Files** button is displayed on the title bar.
- `ActivityIndicator`: Indicates (if set to `True`) that the **Activities** button is displayed on the title bar. This button, which users click to attach activities to the webpage, is not displayed by default.



: The `LinkIndicator` property, which controlled whether the **Help > Get Link** element was displayed on the title bar, is now obsolete. **Help > Get Link** is always displayed on the title bar.

### To Allow Attachments to Table Rows on a Webpage

1. In the data access class (DAC) that provides data for the table rows, add the `NoteID` field, as the following code shows.

```
#region NoteID
public abstract class noteID : PX.Data.IBqlField { }



[PXNote]
public virtual Guid? NoteID { get; set; }
#endregion
```



: The database table that corresponds to the DAC must contain the `NoteID` column with the `uniqueidentifier` data type.

2. Rebuild the project.

Once you have added the `NoteID` field to the DAC and rebuilt the project, the following columns appear in the table:

- : The Notes column, which users click to attach notes to the webpage
  - : The Files column, which users click to attach files to the webpage
3. Optional: To change whether each of these columns is displayed in the table, in the ASPX code of the webpage, specify the values of the following properties of the `PXGrid` control that corresponds to the table:
- `NoteIndicator`: Indicates (if set to `True`) that the Notes column is displayed in the table
  - `FilesIndicator`: Indicates (if set to `True`) that the Files column is displayed in the table

## To Display an Attached Image on the Webpage

In this topic, you will learn how to display an attached image file on the webpage. In Acumatica ERP, you can find an example of an image file displayed in the UI on the **Attributes** tab of the [Stock Items](#) (IN202500) form.

### To Display the Attached Image on the Webpage

1. Add the `NoteID` field and the field that stores the path to the image to the data access class (DAC) that provides data for the webpage on which you want to display the image, as shown in the following code.

```
#region NoteID
public abstract class noteID : PX.Data.IBqlField { }

[PXNote]
public virtual Guid? NoteID { get; set; }
#endregion

#region ImageUrl
public abstract class imageUrl : PX.Data.IBqlField { }

[PXDBString(255)]
[PXUIField(DisplayName = "Image")]
public virtual string ImageUrl { get; set; }
#endregion
```



: The database table that provides data for the webpage on which you want to display the image must contain the following columns:

- `NoteID` with the `uniqueidentifier` data type, to make it possible to attach images
  - The field (in this example, `ImageUrl`) with the `varchar(255)` data type, to store the internal path to the attached image
2. In the ASPX code of the webpage that works with this DAC, add the `PXImageUploader` control linked to the `ImageUrl` data field, as shown in the following code.

```
<px:PXImageUploader Height="320px" Width="430px"
    ID="imgUploader" runat="server" DataField="ImageUrl"
    AllowUpload="true" ShowComment="true"
/>
```

3. Rebuild the project.



## Configuring the UI from the Back End

In this chapter, you can find information about configuration of the user interface of an Acumatica Framework-based application that involves not only the changing of the ASPX code of the webpage (as described in [Configuring ASPX Webpages and Reports](#)) but also the implementation of business logic in the corresponding graph.

### In This Chapter

- [To Configure an Input Mask and a Display Mask for a Field](#)
- [To Display a Dialog Box](#)

## To Configure an Input Mask and a Display Mask for a Field

In this topic, you can learn how to create a field on a webpage of an Acumatica Framework-based application so that its value is displayed in a specific format and how to govern what a user can enter as a value of this field. You can specify input and display masks for fields of the string and date and time types. For fields of the string type, you specify only the input mask, which defines both the format in which the user enters the value and the way the value is then displayed. For the date and time fields, you can specify different input and display masks.

### To Specify an Input Mask and a Display Mask for a String Field

1. In the data access class (DAC), add a new field or modify an existing `string` field as follows:
  - a. Add one of the string attributes (`PXDBString` or `PXString`) to the property field.
  - b. Specify the value of the `InputMask` property of the attribute. Use the following conventions to define the mask:
    - C or &: The user can enter any symbol.
    - A or a: The user can enter any letter or digit.
    - L or ?: The user can enter only a letter.
    - #, 0, or 9: The user can enter only a digit.
    - >: All of the characters that follow this symbol should be in uppercase.
    - <: All of the characters that follow this symbol should be in lowercase.

The following example shows the use of the `InputMask` property.

```
//Users can enter only digits.
//If a user enters "1234567890", the value is displayed as "(123)
456-7890".
[PXDBString(10, InputMask = "(###) ###-####")]
[PXUIField(DisplayName = "Parameter 1")]
public virtual string Parameter1 { get; set; }
```



: The value is stored in the database without any formatting characters. That is, for the code example above, if a user enters 1234567890, the field in the database for the corresponding record will contain the same value (1234567890).

2. In the ASPX code of the webpage, add a new `PXMaskEdit` control or modify the control that corresponds to the field so that it has the `PXMaskEdit` type, as shown in the following code example.

```
<px:PXMaskEdit ID="edParameter1" runat="server" DataField="Parameter1"/>
```

## To Specify an Input and a Display Mask for a String Field at Run Time

To specify the input mask for a string field at run time, use a `SetInputMask` method of the [PXDBString](#) or [PXString](#) attribute. You use the same conventions to define the mask as those described for the `InputMask` property in [To Specify an Input Mask and a Display Mask for a String Field](#).

In the following example, the input mask of the `AccountMask` field is changed at run time.

```
protected virtual void GLBudgetTree_IsGroup_FieldSelecting(PXCache sender,
    PXFieldSelectingEventArgs e)
{
    PXStringState strState = (PXStringState)sender.GetStateExt(
        null, typeof(GLBudgetTree.accountID).Name);
    PXDBStringAttribute.SetInputMask(sender,
        typeof(GLBudgetTree.accountMask).Name,
        strState.InputMask.Replace('#', 'C'));
}
```

## To Specify an Input or a Display Mask for a Date and Time Field

1. In the data access class (DAC), add a new field or modify an existing data and time field as follows:
  - a. Add one of the date and time attributes ([PXDate](#), [PXDateAndTime](#), [PXDBDate](#), [PXDBTime](#), or [PXDBDateAndTime](#)) to the property field.
  - b. Specify the value of the `InputMask` or `DisplayMask` property of the attribute. Use the [standard](#) and [custom](#) date and time format strings.

The following example shows the use of the `InputMask` and `DisplayMask` properties.

```
[PXDateAndTime(DisplayMask = "D", InputMask = "d")]
[PXUIField(DisplayName = "Parameter 1")]
public virtual DateTime? Parameter1 { get; set; }
```

2. In the ASPX code of the webpage, add a new `PXDateTimeEdit` control or modify the control that corresponds to the field so that it has the `PXDateTimeEdit` type, as shown in the following code example.

```
<px:PXDateTimeEdit ID="edParameter1" runat="server" DataField="Parameter1"/>
```



: You can change how the `PXDateTimeEdit` control is displayed (whether the control shows a calendar selector or a drop-down list with time values) by specifying the value of the `TimeMode` property. The following example causes the system to display the list of time values.

```
<px:PXDateTimeEdit ID="edParameter1" runat="server" DataField="Parameter1"
    TimeMode="true"/>
```

## To Display a Dialog Box

When a user clicks a button on a webpage of an Acumatica Framework-based application, you may need to display a dialog box that displays the settings related to the action to be performed. For example, on the [Companies](#) (CS101500) form, if you click **Create Ledger**, the system opens the **Create Ledger** dialog box, where a user can specify the setting related to the action.

### To Display a Dialog Box

1. In the graph that corresponds to the webpage, add the action and the delegate for the button that opens the dialog box, as shown in the following example.

```
public PXAction<MainDAC> openDialogBox;
```

```
[PXUIField(DisplayName = "Open Dialog Box",
    MapEnableRights = PXCacheRights.Update,
    MapViewRights = PXCacheRights.Update)]
[PXButton]
public virtual IEnumerable OpenDialogBox(PXAdapter adapter)
{
    return adapter.Get();
}
```

In this example, the `MainDAC` DAC is the main DAC of the primary view of the graph; therefore, the action is added to the toolbar of the form by default.

2. If you want to specify additional properties of the action (such as `CommitChanges`), in the ASPX code of the webpage, add the `PXDSCallbackCommand` tag for the action and specify the necessary properties of the tag, as shown in the following example.

```
<px:PXDataSource ID="ds" runat="server" ...>
    <CallbackCommands>
        ...
        <px:PXDSCallbackCommand Name="CreateLedger" CommitChanges="true"/>
    </CallbackCommands>
</px:PXDataSource>
```

3. In the graph that corresponds to the webpage, add the data view for the dialog box, as shown in the following code.

```
[Serializable]
public class DialogBoxParameters : IBqlTable
{
    public abstract class parameter1 : IBqlField { }
    [PXString(10, IsUnicode = true)]
    public virtual string Parameter1 { get; set; }

    public abstract class parameter2 : IBqlField { }
    [PXString(10, IsUnicode = true)]
    public virtual string Parameter2 { get; set; }
}

PXFilter<DialogBoxParameters> OpenDialogBoxView;
```

In this code, we add a simple DAC with two unbound fields and use a data view based on the `PXFilter` class.

4. In the ASPX code of the webpage, add the `PXSmartPanel` container with the `Key` property equal to the name of the data view you created for the dialog box, as shown in the following code. In the `PXPanel` container inside `PXSmartPanel`, add the commit buttons of the dialog box (such as **OK**, **Cancel**).

```
<px:PXSmartPanel ID="pnlOpenDialogBox" runat="server"
    Style="z-index: 108;" Caption="Open Dialog Box" CaptionVisible="True"
    Key="OpenDialogBoxView" LoadOnDemand="true" AutoCallBack-Command="Refresh"
    AutoCallBack-Target="formOpenDialogBox" AutoCallBack-Enabled="true"
    AcceptButtonID="cbOk" CancelButtonID="cbCancel">
    <px:PXFormView ID="formOpenDialogBox" runat="server"
        DataSourceID="ds" DataMember="OpenDialogBoxView"
        SkinID="Transparent">
        <ContentStyle BorderWidth="0px"></ContentStyle>
        <Template>
            <px:PXLayoutRule runat="server" StartColumn="True"
                LabelsWidth="SM" ControlSize="M" />
            <px:PXTextEdit ID="edParameter1" runat="server"
                DataField="Parameter1" CommitChanges="True"/>
            <px:PXTextEdit ID="edParameter2" runat="server"
                DataField="Parameter2" CommitChanges="True" />
        </Template>
    </px:PXFormView>
```

```

<px:PXPanel ID="PXPanel3" runat="server" SkinID="Buttons">
  <px:PXButton ID="cbOK" runat="server" Text="OK"
    CommandSourceID="ds" DialogResult="OK" />
  <px:PXButton ID="cbCancel" runat="server" Text="Cancel"
    DialogResult="Cancel" />
</px:PXPanel>
</px:PXSmartPanel>

```

5. In the button delegate, perform a call to an [AskExt](#) method of the view specified in the `Key` property of the `PXSmartPanel` container.

```

public virtual IEnumerable OpenDialogBox(PXAdapter adapter)
{
    if (OpenDialogBoxView.AskExt() == WebDialogResult.OK &&
        string.IsNullOrEmpty(OpenDialogBoxView.Current.Parameter1)
        == false)
    {
        ...
    }
    return adapter.Get();
}

```

When the user clicks the button on the webpage, the execution interrupts on the `AskExt` call and the dialog box is displayed. After the user clicks a button in the dialog box, the `AskExt` method returns the dialog box result.

6. Rebuild the project.

## Executing Code Asynchronously

In this chapter, you can find information about how to process operations whose execution takes long time. You execute these operations asynchronously in separate threads.

### In This Chapter

- [Asynchronous Execution](#)
- [To Add a Button to the Processing Dialog Box](#)
- [To Not Display the Processing Dialog Box](#)

## Asynchronous Execution

An instance of a graph is created on each round trip to process a request created by the user on the appropriate form. After the request is processed, the graph instance must be cleared from the memory of the Acumatica ERP server. If you implement code that might require a long time to execute an action or to process a document or data, you should execute this code asynchronously in a separate thread.

### Using the `PXLongOperation` Class

To make the system invoke the method in a separate thread, you can use the `PXLongOperation.StartOperation` method. Within the method that you pass to `StartOperation`, you can, for example, create a new instance of a graph and invoke a processing method on that instance. The following code snippet demonstrates how you can execute code asynchronously as a long-running operation in a method of a graph.

```

public class MyGraph : PXGraph
{
    ...
    public void MyMethod()
    {
        ...
    }
}

```

```

    PXLongOperation.StartOperation(this, delegate()
    {
        // insert the delegate method code here
        ...
        GraphName graph = PXGraph.CreateInstance<GraphName>();
        foreach (... in ...)
        {
            ...
        }
        ...
    });
    ...
}
    ...
}

```

If you need to start a long-running operation in a method of a graph extension, you have to use the `Base` property instead of the `this` keyword in the first parameter of the `StartOperation` method, as shown in the following code snippet.

```


public class MyGraph_Extension : PXGraphExtension<MyGraph>
{
    ...
    public void MyMethod()
    {
        ...
        PXLongOperation.StartOperation(Base, delegate()
        {
            // insert the delegate method code here
            ...
        });
        ...
    }
    ...
}

```

### Using the Custom Information Dictionary

In the delegate method of a long-running operation, you can store a data object in the `_CustomInfo` dictionary of the long-running operation and get the list of records processed by the method. You can add to the dictionary any data object needed for a long-running operation by using a `SetCustomInfo` method.

The following diagram shows that each long-running operation includes the `_CustomInfo` dictionary, which can contain multiple key-value pairs with custom data.



**Figure: Location of custom data in the memory of the Acumatica ERP server**

For a processing operation, the system stores the `PXProcessingMessagesCollection<TTable>` list of messages in the dictionary. Each message in the list is of the `PXProcessingMessage` type, which includes a string message and an error level that is of the `PXErrorLevel` type.

See *New way to work with CustomInfo of PXLONGOperation* at <http://asiablog.acumatica.com> for more information about the use of the dictionary.

### Executing a Processing Operation as a Long-Running Operation

When a user clicks an action button on a form to start a processing operation, the data source control of the form generates a request to the Acumatica ERP server to execute the action delegate defined for the button. The server creates an instance of the graph, which provides the business logic for the form and invokes the action delegate method.

Because a processing operation is a long-running operation, in the action delegate method, the data processing code must be included in the `PXLONGOperation.StartOperation` method call as the definition of the long-running operation delegate. When the action delegate method is executed, the `StartOperation` method creates an instance of the `PXAsyncResult` class to hold the data and state of the long-running operation; the method also initiates the execution of the long-running operation delegate asynchronously in a separate thread.

If the duration of the long-running operation is longer than five seconds, the server releases the graph instance. Then the form, which is still opened in the browser, generates requests to the server to get the results of the long-running operation every five seconds. For such a request, the server uses the ID of the long-running operation to check the operation status. If the operation has completed, the server creates an instance of the graph and restores the graph state and the cache data to finish processing the action delegate and to return results to the form.

The following diagram shows how the server executes an action asynchronously and how it returns the results of the action to the form.

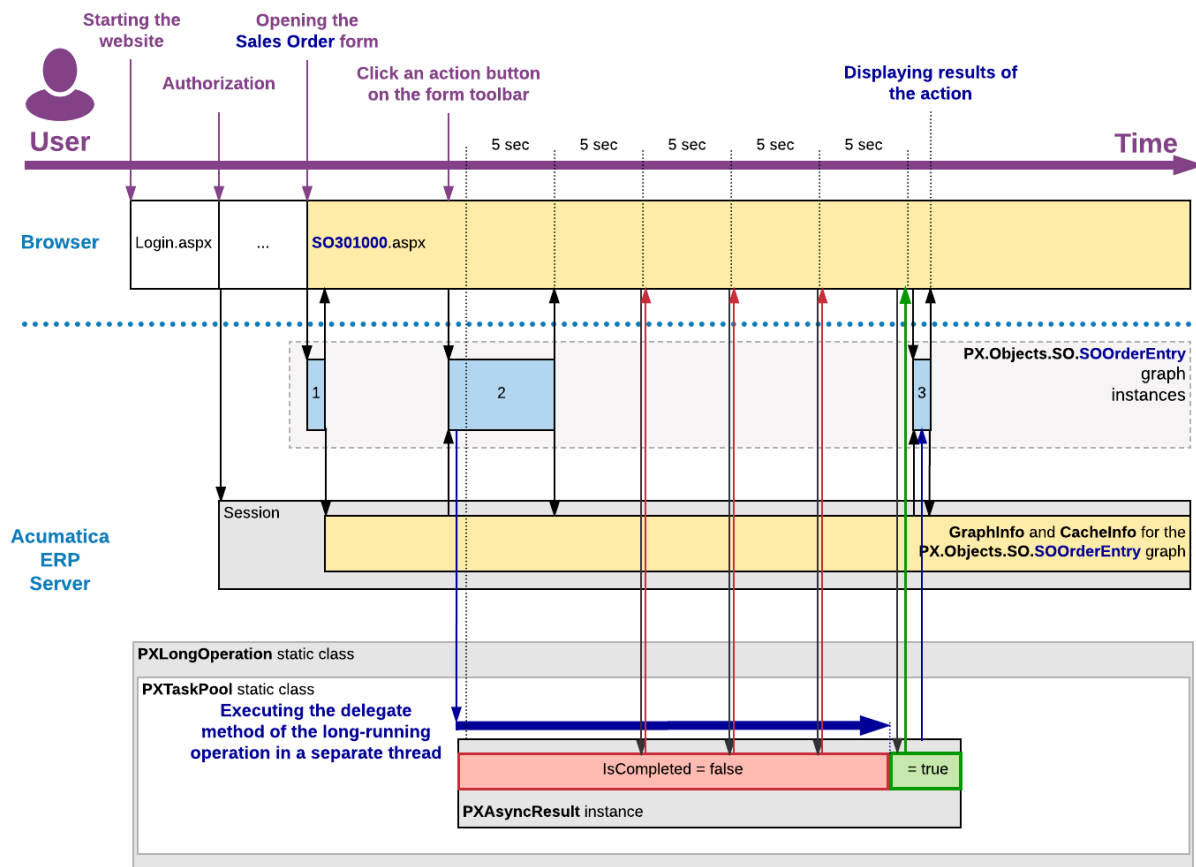


Figure: Execution of an action that uses a long-running operation

### Processing a Report as a Long-Running Operation

When the user launches a report, either from the report form or by clicking an action button on the maintenance or entry form, the system redirects the user to the report launcher form (`ReportLauncher.aspx`), which is designed to automatically run a report for the received parameters. The ASPX page for this form contains the `PXReportViewer` control, whose Java Scripts objects and functions are designed to get the report data and display the data on the form.

To run the report, the report launcher creates a request to the `PX.Web.UI.PXReportViewer` control on the server. To process the request, the server instantiates the `PX.Reports.Web.WebReport` class and invokes its `Render` method, which launches the report generation as a long-running operation in a separate thread.

The resulting report data is an object of the `PX.Reports.Data.ReportNode` type stored in the `_CustomInfo` dictionary of the current long-running operation under the `DEFAULT_CUSTOM_INFO_KEY` key. To provide quick access to the report data when the user views different pages of the report, the system saves the report data in the session as an object of the `PX.Reports.Web.WebReport` type.

After the long-running operation has completed, the `PXReportViewer` control gets the report data from the dictionary and displays the report on the report launcher form. For details on how the report is displayed and how the report data is retrieved, see [Display of Reports](#).

The following diagram shows how the server generates a report asynchronously and how it returns the resulting report data to the report launcher form.

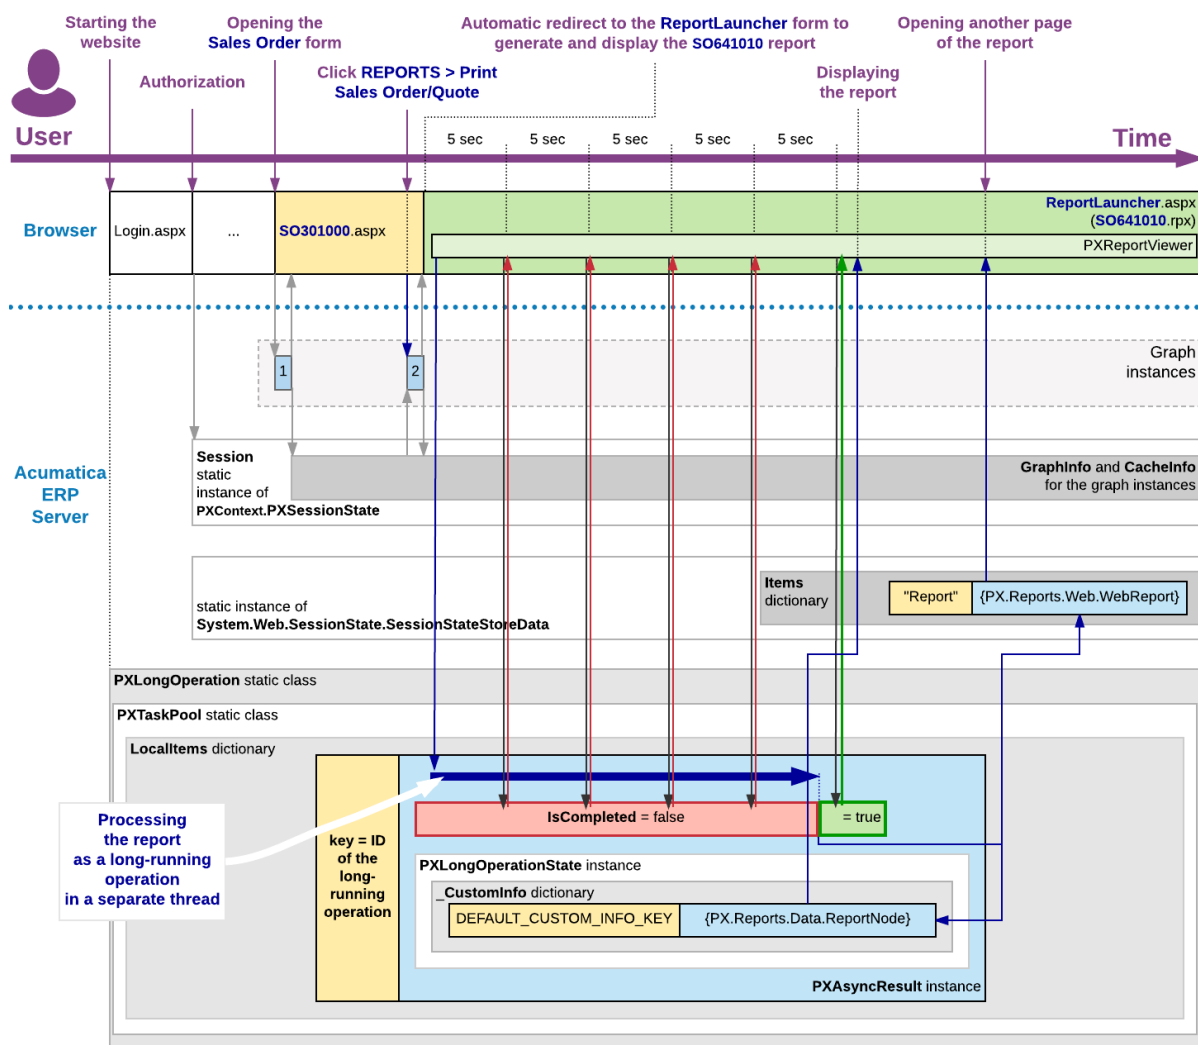


Figure: Execution of an action that launches the generation of a report

### Executing a Long-Running Operation in a Cluster

If Acumatica ERP is configured to run in a cluster of application servers behind a load balancer, it is not possible to predict which application server will receive the next request from the client. In this model, session-specific data is serialized and stored in a high-performance remote server, such as Redis or MS SQL, to be shared between the application servers.

When the user clicks an action button on a form to start a processing operation, the load balancer forwards the request to an Acumatica ERP server to execute the action delegate defined for the button. The server creates an instance of the graph, which provides the business logic for the form, and invokes the action delegate method.

When the action delegate method is executed, the `StartOperation` method creates an instance of the `PXAsyncResult` class to hold the data and state of the long-running operation, initiates the execution of the long-running operation delegate asynchronously in a separate thread, and stores the serialized data of the operation in the remote storage.

If the duration of the long-running operation is longer than five seconds, the server releases the graph instance, stores the serialized data of the graph in the remote storage, and continues processing the long-running operation in a separate thread. When the operation has completed, this server sets the operation status to `PXLongRunStatus.Completed` and updates the operation data in the remote session storage.



Until the form that is still opened in the browsers obtains the request results, it generates requests to the site URL every five seconds to get these results. On every such request, the load balancer selects a server to be used to process the request and forwards the request to the server. The server uses the long-running operation ID, which is usually equal to the graph UID, to check the operation's status. If the operation is completed, the server creates an instance of the graph to finish processing the action delegate and to return results to the form.

The following diagram shows how the data of a user session and of a long-running operation are stored in the remote session storage of a cluster.

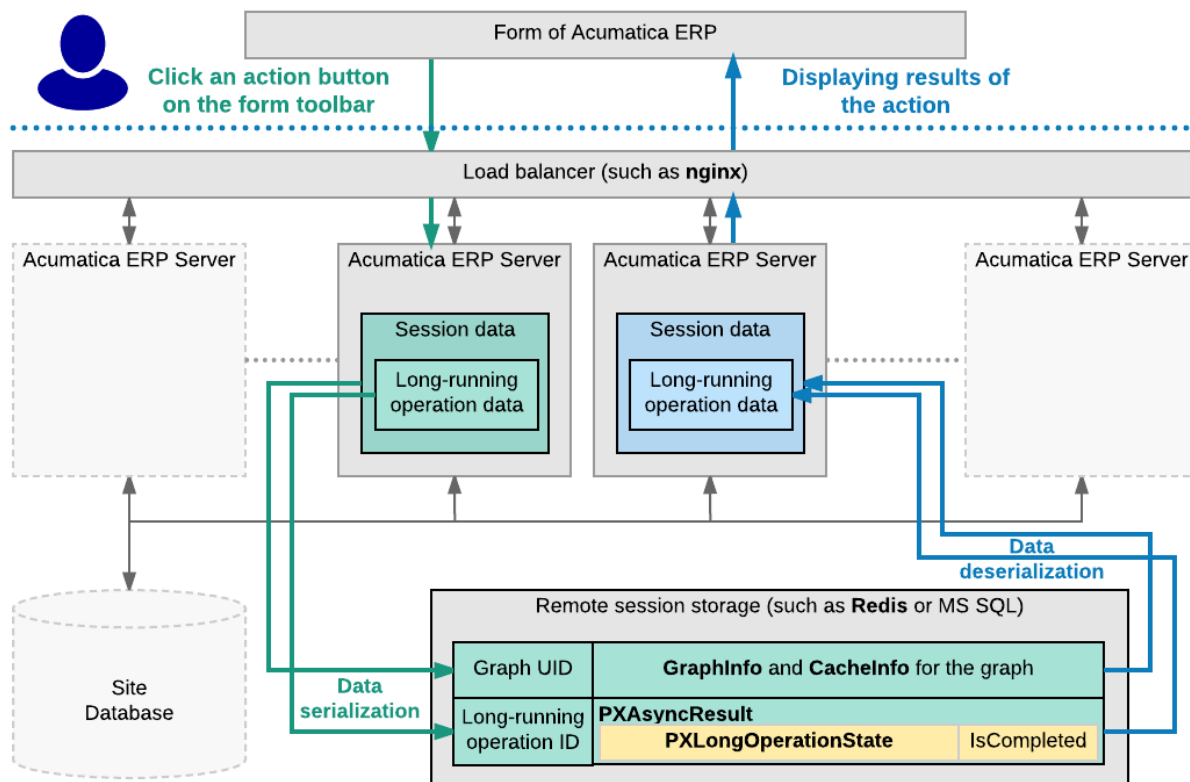


Figure: Execution of an action that uses a long-running operation in a cluster

## To Add a Button to the Processing Dialog Box

When a user starts a processing operation on a processing form, such as the [Release AR Documents](#) (AR501000) form of Acumatica ERP, the **Processing** dialog box opens, which displays the status of the processing. When a processing operation is started, all elements of the processing form become unavailable. If you need to make a button from the processing form available during processing, you have to add this button to the processing dialog box, as described in this topic.

### To Add a Button to the Processing Dialog Box

To add a button to the processing dialog box, you can use one of the following approaches:

- For the action that corresponds to the button, in the graph, set the value of the `VisibleOnProcessingResults` property of `PXButtonAttribute` or its descendant to `true`, as shown in the following code example.

```
[PXUIField(DisplayName = Messages.ShowDocuments)]
[PXButton(VisibleOnProcessingResults = true)]
public virtual IEnumerable showDocuments(PXAdapter adapter)
{
    ShowOpenDocuments(SelectedItems);
    return adapter.Get();
}
```

```
}
```

- In the ASPX file that corresponds to the form, set the value of the `VisibleOnProcessingResults` property of `PXDSCallbackCommand` to `True`, as shown in the following example.

```
<px:PXDataSource ID = "ds" Width="100%"
  runat="server" Visible="True" PrimaryView="Filter"
  TypeName="PX.Objects.FA.FAClosingProcess" >
  <CallbackCommands>
    <px:PXDSCallbackCommand Name = "showDocuments"
      VisibleOnProcessingResults="True"/>
  </CallbackCommands>
</px:PXDataSource>
```

## To Not Display the Processing Dialog Box

When a user starts a processing operation on a processing form, such as the [Release AR Documents](#) (AR501000) form of Acumatica ERP, the **Processing** dialog box opens, which displays the status of the processing. You can turn off the displaying of the processing dialog box and instead display the progress and the result of the processing on the form toolbar.

### To Turn Off the Displaying of the Processing Dialog Box

To not display the processing dialog box on a processing form, you can do one of the following:

- To not display the processing dialog box for a custom page, override the `IsProcessing` property of the graph that corresponds to the form, as shown in the following code.

```
public override bool IsProcessing
{
    get { return false; }
    set { }
}
```

- To not display the processing dialog box for a customized Acumatica ERP form, configure the `IsProcessing` property of the graph that corresponds to the form in a graph extension as shown in the following code.

```
public class AllocationProcess_Extension : PXGraphExtension<AllocationProcess>
{
    public override void Initialize()
    {
        Base.IsProcessing = false;
    }
}
```

- To not display the processing dialog box for all processing forms, add the `ProcessingProgressDialog` key in the `appSettings` section of the `web.config` file of the application set to `False`, as shown in the following example.

```
<add key="ProcessingProgressDialog" value="false" />
```

## Localizing Applications

Acumatica Framework provides built-in localization tools that you can use to translate the user interface and application messages to different languages. This chapter provides guidelines on how to prepare the Acumatica Framework application for localization efforts.

To get the application ready for localization, you must prepare data access classes (DACs) and the application code.

### In This Chapter

- [Localization](#)
- [Strings That Can Be Localized](#)
- [To Prepare DACs for Localization](#)
- [To Localize Application Messages](#)
- [To Work with Multi-Language Fields](#)
- [To Optimize Memory Consumption of Localized Data](#)

## Localization

Applications created with Acumatica Framework can be localized on the presentation, business logic, and database level owing to:

- Standard Microsoft.NET localization mechanism is implemented for localizing the presentation layer.
- All messages returned from the business logic layer can be localized through the dictionary mechanism. For details about how to implement message localization in your code, see [To Localize Application Messages](#).
- The runtime environment of Acumatica Framework supports the Unicode standard to store and operate with data in a non-ANSI format.
- Information like addresses or product descriptions can be stored in special, language-specific, database fields and presented in the user selected language. For details on implementation of such fields, see [To Work with Multi-Language Fields](#).

Acumatica Framework provides a built-in utility that enables localization of the product by the end user. Once localization is entered and applied, the application does not require any recompilation or re-installation. Also, localization can also be exported, imported, and merged. For more information about how to use the built-in localization mechanism, see [Translation Process](#) in the System Administration Guide.

## Strings That Can Be Localized

By using the [Translation Dictionaries](#) (SM200540) form, you can add translations for the string constants that are collected from the code of the application, and save them to the database. When a user signs in with a specific language, the systems loads the translations and displays the translated strings to the user. For more information on localization, see [Locales and Languages](#).

The system collects for localization the string constants that are specified in the following items of the application:

- The `DisplayName` property of the `PXUIField` attribute of the fields of data access classes (DACs)
- The `DisplayName` property of the `PXUIField` attribute of fields and actions of a business logic controller (BLC) object, which override the attributes of the fields and actions of a DAC
- The `AllowedLabels` property of the `PXStringList` and `PXIntList` attributes
- The `Values` property of the classes that implement the `ILocalizableValues` interface
- Captions of form, grid, and panel controls and labels of input controls, which are specified in ASPX
- Titles of all nodes in the site map
- Report elements (such as text box labels and diagram agendas)
- `public const string` fields of the classes marked with the `PXLocalizable` attribute

You can also translate user input to multiple languages and store translations in the database. For more information on the localization of user input, see [To Work with Multi-Language Fields](#).

## To Prepare DACs for Localization

When the system localizes the fields of the data access classes (DACs) and DAC names, it collects the string constants that are specified in the following code elements:

- The `DisplayName` property of the `PXUIField` attribute of the fields of DACs
- The `AllowedLabels` property of the `PXStringList` attribute or `PXIntList` attribute of the fields of DACs

To prepare each DAC for localization, you need to perform the steps that are described in this topic.

### To Prepare Each DAC for Localization

1. Make sure the `DisplayName` parameter of the `PXUIField` attribute is specified for each visible field in the DAC, as shown in the following example.



: If you change the `DisplayName` value of the `PXUIField` attribute on the fly (by creating your own `PXFieldState`), you should localize the string independently.

```
public new abstract class docType : PX.Data.IBqlField{}
[PXDBString(3, IsKey = true, IsFixed = true)]
[PXDefault()]
[PXUIField(DisplayName = "Document Type")]
public override string DocType { get; set; }
```

2. Specify the values that should be displayed in drop-down lists by using the `PXStringList` attribute, as shown in the following example.

```
public abstract class lineSource : PX.Data.IBqlField{}
[PXString(1, IsFixed = true)]
[PXStringList(
    new string[] { "D", "R" },
    new string[] { "Draft", "Request" })]
[PXUIField(DisplayName = "Line Source")]
public virtual string LineSource { get; set; }
```

## To Localize Application Messages

For localization of the messages in the source code, the system collects the strings from the classes that are marked with the `PXLocalizable` attribute.

To make your application display localized messages, you need to perform the steps, which are described in this topic.

### To Display Localized Messages in Your Application

1. Move all strings that should be translated to the public static `Messages` class and specify the `PXLocalizable` attribute for this class, as shown in the following code.



: The exceptions to this requirement are field descriptions and list attributes in the data access classes, which are handled separately. For details on how to make field descriptions and list attributes localizable, see [To Prepare DACs for Localization](#).

```
using System;
using PX.Data;

[PXLocalizable()]
public static class Messages
{
```

```

public const string FieldNotFound = "The field is not found.";
public const string InvalidAddress = "The address is not valid.";
public const string AdditionalData = "Author's title: {0}; author's name:
{1}."
}

```



: No hyphenation is provided by the system. During the acquisition process of localizable data, all the new-line symbols (`\n\r`) are to be removed. You can use the reserved symbol (`~`) to cause the insertion of a new line.

2. If the message from the `Messages` class is used in an error or warning message, which is displayed when an exception of the `PXException` type or of a type derived from `PXException` is thrown, provide a non-localized message, as shown in the following example. The system displays the localized message automatically if there is a translation for this message in the database.

```

if (field == null)
{
    throw new PXException(Messages.FieldNotFound);
}

```

3. If you need to receive the translation of a message from the `Messages` class within the application code (for example, if the message is displayed in the confirmation dialog box, which is displayed if you use the `Ask()` method of a data view in the code), use one of the following methods:

- `PXMessages.Localize()`: The method searches for the translation of the provided string in the database and returns the first translation found.

```

string msg = PXMessages.Localize(Messages.FieldNotFound);

```

- `PXMessages.LocalizeFormat()`: The method searches for the translation of the provided string, which includes placeholders (such as `{0}` or `{1}`), in the database and returns the first translation found.
- `PXLocalizer.Localize()`: The method returns the translation with the given key, which you specify in the second parameter. A string may have multiple translations; one translation for each occurrence of the string in the application. For each of the occurrences, a key value is created. For example, if the string is declared in a class marked with the `PXLocalizable` attribute, the full qualified name of the class is the key, as the following code shows.

```

string localizedMsg = PXLocalizer.Localize(
    ActionsMessages.ChangesWillBeSaved,
    typeof(ActionsMessages).ToString());

```

## To Work with Multi-Language Fields

With Acumatica Framework, you can create fields into which a user can type values in multiple languages if multiple locales are configured in the applicable Acumatica Framework application. For example, in Acumatica ERP, if an instance works with English and French locales, you can specify the value of the **Description** box on the [Stock Items](#) (IN202500) form in English and French. For details on multi-language fields on Acumatica ERP forms, see [Managing Locales and Languages](#).

### To Configure a Field to Have Values in Multiple Languages

1. In the data access class (DAC) that you want to contain a multi-language field, define the `NoteID` field with the `PXNote` attribute, as follows.

```

public abstract class noteID : IBqlField { }

```

```
[PXNote]
public virtual Guid? NoteID { get; set; }
```

2. If you want to configure a field to have values in multiple languages, annotate this field with the `PXDBLocalizableString` attribute. The following code shows an example of the use of the `PXDBLocalizableString` attribute.

```
[PXDBLocalizableString(60, IsUnicode = true)]
```

The `PXDBLocalizableString` attribute works similarly to the `PXDBString` attribute, but unlike the `PXDBString` attribute, the `PXDBLocalizableString` attribute can be used instead of the `PXDBText` and `PXString` attributes.

3. If you need to give values in multiple languages to a field with the `PXDBText` attribute, replace this attribute with the `PXDBLocalizableString` attribute and do not specify the length parameter, as shown in the following example.

```
[PXDBLocalizableString(IsUnicode = true)]
```

4. If you need to configure a field that has the `PXString` attribute, which is used in conjunction with the `PXDBCalced` attribute, replace the `PXString` attribute with the `PXDBLocalizableString` attribute and set the value of the `NonDB` parameter to `true`, as shown in the following example.

```
[PXDBLocalizableString(255, IsUnicode = true, NonDB = true,
    BqlField = typeof(PaymentMethod.descr))]
[PXDBCalced(typeof(Switch<Case<Where<PaymentMethod.descr, IsNotNull>,
    PaymentMethod.descr>, CustomerPaymentMethod.descr>), typeof(string))]
```

### To Configure the Default Value of a Multi-Language Field

If you want a multi-language field to have a default value in a specific language, use the `PXLocalizableDefault` attribute instead of the `PXDefault` attribute and specify in its second parameter either a BQL field or a BQL select that provides language selection.

For example, in Acumatica ERP, the `SOLine` line description defaulted to the appropriate `InventoryItem` description based on the language that is set for a customer. The `TransactionDesr` field of the `SOLine` DAC has the `PXLocalizableDefault` attribute with a second parameter that specifies the language as follows: `typeof(Customer.languageName)`. See the following example of the use of the `PXLocalizableDefault` attribute.

```
[PXLocalizableDefault(typeof(Search<InventoryItem.descr,
    Where<InventoryItem.inventoryID,
    Equal<Current<SOLine.inventoryID>>>>),
    typeof(Customer.languageName),
    PersistingCheck = PXPersistingCheck.Nothing)]
```

### To Obtain the Value of a Multi-Language Field in the Current Locale

If you want to obtain the value of a multi-language field in the current locale, use the `PXDatabase.SelectSingle()` or `PXDatabase.SelectMulti()` method, and pass to it the return value of the `PXDBLocalizableStringAttribute.GetValueSelect()` static method instead of passing a new `PXDataField` object to it. (The `PXDBLocalizableStringAttribute.GetValueSelect()` method takes three input parameters: the table name, the field name, and a Boolean flag that indicates whether strings should be returned as text with unlimited length.)

The following code shows an example of the use of the `PXDBLocalizableStringAttribute.GetValueSelect()` method.

```
foreach (PXDataRecord record in PXDatabase.SelectMulti<Numbering>(
    newPXDataField<Numbering.numberingID>(),
```

```
PXDBLocalizableStringAttribute.GetValueSelect("Numbering",
    "NewSymbol", false),
newPXDataField<Numbering.userNumbering>()))
{
    ...
}
```



: Generally, you use the `PXDatabase.SelectSingle()` and `PXDatabase.SelectMulti()` methods for retrieving data within the `Prefetch()` method of a database slot. Don't forget to add language code to the slot key when you obtain a slot, as shown in the following example, because with the use of `PXDBLocalizableStringAttribute`, the data becomes language-specific. Therefore, you need different slot instances for different languages.

```
Numberings items = PXDatabase.GetSlot<Numberings>(
    typeof(Numberings).Name + currentLanguage, typeof(Numbering));
```

### To Obtain the Value of a Multi-Language Field in a Specific Language

If you want to obtain the value of a multi-language field in a specific language, use the `PXDBLocalizableStringAttribute.GetTranslation()` method. Pass to the method as input parameters a DAC cache, a DAC instance, a field name, and the ISO code of the language.

The following code shows an example of use of the `PXDBLocalizableStringAttribute.GetTranslation()` method.

```
tran.TranDesc =
    PXDBLocalizableStringAttribute.GetTranslation(
        Caches[typeof(InventoryItem)], item, typeof(InventoryItem.descr).Name,
        customer.Current?.LanguageName);
```

## To Optimize Memory Consumption of Localized Data

To optimize the memory consumption of static data, you can move the localization data from all customer application instances to centralized storage. By default, the localization data is kept in the database of every Acumatica ERP instance, and the total size of this data therefore equals the number of instances times the size of the data. If you move the localization data to centralized storage, there is only one copy of this data.

Alternatively, you can optimize the consumption of memory by disabling localization.

Whether you set up centralized storage of localization data or disable localization, you should perform the following steps:

1. Implement a custom translation provider. Follow the instruction in [To Implement a Custom Translation Provider](#) or [To Disable Localization](#) in this topic depending on which way of optimization of memory consumption you select.
2. Place the assembly file with the new provider in the `Bin` directory of the Acumatica ERP instance, and add the assembly to the customization project as a *File* element.
3. Register the new provider in the `pxtranslate` element of the `web.config` file, as described in [To Register the New Provider in Web.config](#) in this topic.

### To Implement a Custom Translation Provider

To implement a custom translation provider, derive a class from the `PXTranslationProvider` class and override the `LoadCultureDictionary()` method, as the following example shows.

```
public class DemoTranslationProvider : PXTranslationProvider
{
    public override PXCultureDictionary LoadCultureDictionary(
        string locale, bool includeObsolete, bool escapeStrings)
    {
    }
```

```

        PXCultureDictionary dictionary = new PXCultureDictionary();
        ...
        // Adding a general translation for some string
        dictionary.Append(
            valueToTranslate,
            new PXCultureValue(locale, translation));
        // Adding a special translation for some string
        dictionary.AppendException(
            valueToTranslate,
            new PXCultureEx(resourceID, locale, translation));
        ...
        return dictionary;
    }
}

```

The `LoadCultureDictionary()` method returns an instance of the `PXCultureDictionary` type. Values are added to objects of this type through the `Append()` and `AppendException()` methods. `Append()` adds a general translation for a string. `AppendException()` adds a translation for a special case (exception) identified by the resource key.

The code below defines a custom translation provider that loads the localization data from an external Acumatica ERP database by using ADO.NET tools.

```

using System;
using System.Collections.Generic;
using System.Data.SqlClient;
using System.Linq;
using System.Text;
using PX.Data;
using PX.Translation;

namespace Demo.Translation
{
    public class DemoTranslationProvider : PXTranslationProvider
    {
        private struct TranslationKey
        {
            public Guid id;
            public string resKey;
            public string locale;
        }

        // The connection string for the database that stores localization
        // data
        // Specify a specific value of the connection string
        private const string connectionString = "";

        // Overriding the method that returns the dictionary of
        // localization data
        public override PXCultureDictionary LoadCultureDictionary(
            string locale, bool includeObsolete, bool escapeStrings)
        {
            string localizationValueSelect;
            string localizationTranslationSelect;
            InitializeSelectCommand(locale, includeObsolete,
                                   out localizationValueSelect,
                                   out localizationTranslationSelect);

            Dictionary<Guid, string> localizationValue;
            Dictionary<TranslationKey, string> localizationTranslation;
            SelectLocalizationValues(localizationValueSelect,
                                   localizationTranslationSelect,
                                   out localizationValue,
                                   out localizationTranslation);

            return CreateCultureDictionary(escapeStrings, localizationValue,
                                           localizationTranslation);
        }
    }
}

```



```

// Builds the SQL statement for selecting localization data
private void InitializeSelectCommand(
    string locale, bool includeObsolete,
    out string localizationValueSelect,
    out string localizationTranslationSelect)
{
    StringBuilder localizationValueSelectBld =
        new StringBuilder("Select IDlv, NeutralValue" +
            "From LocalizationValue");
    if (!includeObsolete)
    {
        localizationValueSelectBld.Append(" Where IsObsolete = 0");
    }
    localizationValueSelect = localizationValueSelectBld.ToString();

    StringBuilder localizationTranslationSelectBld =
        new StringBuilder("Select IDlt, ResKey, Value, Locale" +
            "From LocalizationTranslation");
    if (!string.IsNullOrEmpty(locale))
    {
        localizationTranslationSelectBld.AppendFormat(
            " Where Locale = '{0}'", locale);
    }
    localizationTranslationSelect =
        localizationTranslationSelectBld.ToString();
}

// Retrieves localization data from the database by using the provided
// SQL statement
private void SelectLocalizationValues(
    string localizationValueSelect,
    string localizationTranslationSelect,
    out Dictionary<Guid, string> localizationValue,
    out Dictionary<TranslationKey, string> localizationTranslation)
{
    localizationValue = new Dictionary<Guid, string>();
    localizationTranslation =
        new Dictionary<TranslationKey, string>();

    using (SqlConnection connection =
        new SqlConnection(connectionString))
    {
        connection.Open();

        SqlCommand command = new SqlCommand(localizationValueSelect,
            connection);
        using (SqlDataReader reader = command.ExecuteReader())
        {
            while (reader.Read())
            {
                localizationValue.Add(reader.GetGuid(0),
                    reader.GetString(1));
            }
        }

        command.CommandText = localizationTranslationSelect;
        using (SqlDataReader reader = command.ExecuteReader())
        {
            while (reader.Read())
            {
                TranslationKey newTranslationKey = new TranslationKey()
                {
                    id = reader.GetGuid(0),
                    resKey = reader.GetString(1),
                    locale = reader.GetString(3)
                };
                localizationTranslation.Add(newTranslationKey,
                    reader.GetString(2));
            }
        }
    }
}

```

```

    }
}

// Fills the PXCultureDictionary object with localization data by using
// the provided dictionaries of values to translate and the
// corresponding translations
private PXCultureDictionary CreateCultureDictionary(
    bool escapeStrings,
    Dictionary<Guid, string> localizationValue,
    Dictionary<TranslationKey, string> localizationTranslation)
{
    PXCultureDictionary dictionary = new PXCultureDictionary();

    if (localizationTranslation.Count != 0)
    {
        foreach (Guid id in localizationValue.Keys)
        {
            IEnumerable<TranslationKey> localizationTranslationKeys =
                from translationRowKey in localizationTranslation.Keys
                where translationRowKey.id == id
                select translationRowKey;
            foreach (TranslationKey key in localizationTranslationKeys)
            {
                string translationResKey = key.resKey;
                string translationLocale = key.locale;
                string translationValue = localizationTranslation[key];
                string value = escapeStrings ?
                    PXLocalizer.EscapeString(translationValue) :
                    translationValue;

                if (string.IsNullOrEmpty(translationResKey))
                {
                    dictionary.Append(
                        localizationValue[id],
                        new PXCultureValue(translationLocale, value));
                }
                else
                {
                    dictionary.AppendException(
                        localizationValue[id],
                        new PXCultureEx(translationResKey,
                            translationLocale, value));
                }
            }
        }
    }
    return dictionary;
}
}
}

```

### To Disable Localization

To disable localization, implement a custom translation provider with the `LoadCultureDictionary()` method that returns `null`, as the following code shows.

```

public class DemoTranslationProvider : PXTranslationProvider
{
    public override PXCultureDictionary LoadCultureDictionary(
        string locale, bool includeObsolete, bool escapeStrings)
    {
        return null;
    }
}

```

### To Register the New Provider in Web.config

Once the provider class is defined, register it in the `web.config` file by adding a new key to the `providers` collection of the `pxtranslate` element and specifying the new key in the `defaultProvider` property of `pxtranslate`. Use the `add` element to register the provider. Set the `name` attribute to the key, which can be any unique value, and specify the type of the custom provider in the `type` attribute.

The following code shows the configuration of `DemoTranslationProvider`, introduced in the example above, in the `pxtranslate` element of the `web.config` file.

```
<pxtranslate defaultProvider="DemoTranslationProvider">
  <providers>
    <!--The default translation provider-->
    <remove name="PXDBTranslatonProvider" />
    <add name="PXDBTranslatonProvider"
        type="PX.Data.PXDBTranslatonProvider, PX.Data" />

    <!--The custom translation provider-->
    <remove name="DemoTranslationProvider" />
    <add name="DemoTranslationProvider"
        type="Demo.Translation.DemoTranslationProvider, TranslationProvider"
        applicationName="/" />
  </providers>
</pxtranslate>
```

## Reusing Business Logic

In an Acumatica ERP application or an Acumatica Framework-based application, you may need to use the same business logic in multiple places. For example, Acumatica ERP supports the calculation of amounts in multiple currencies. Therefore, the business logic containers (also called *graphs*) that implement the multi-currency logic are included in different modules of the application.

With the ability to reuse business logic in Acumatica ERP or Acumatica Framework, you can include the main business logic of particular functionality (such as multi-currency processing) in reusable generic graph extensions and use this logic whenever you need to. If you need to adjust this logic for the specifics of a particular module, you can override this business logic in the implementation of this module. For example, you can assign different names for the UI elements that are linked to the same fields of a data access class in different modules.

In this section, you can find information about generic graph extensions.

### In This Chapter

- [Reusable Business Logic Implementation](#)
- [Mapped Cache Extensions and the Application Database](#)
- [Reusable Business Logic and the Application Website](#)
- [Use of Generic Graph Extensions by the System](#)
- [Generic Graph Extensions Declared in Acumatica ERP](#)
- [To Insert Reusable Business Logic That Has Already Been Declared](#)
- [To Sort Multiple Generic Graph Extensions](#)
- [To Implement Reusable Business Logic](#)

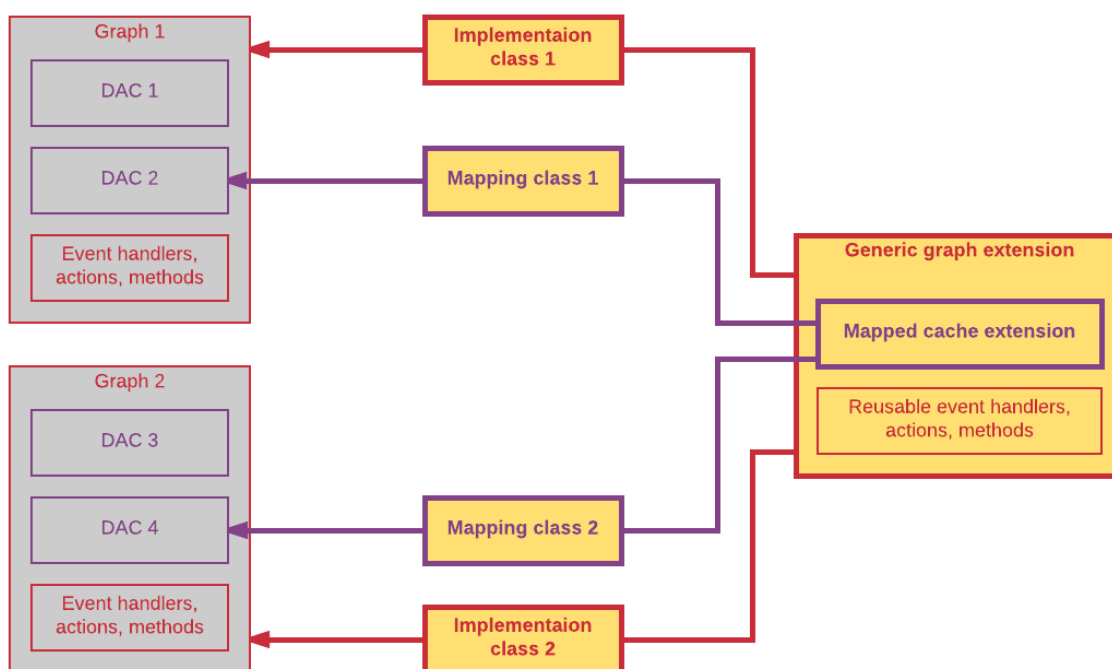
## Reusable Business Logic Implementation

Suppose that you want to use the same business logic in multiple places in your application. That is, you have at least two graphs in which you need to insert the logic. The graphs operate with data in the

data access classes (DACs) and implement business logic through event handlers, actions, and other methods.

You encapsulate the business logic that you want to reuse in a *generic graph extension*, which is a graph extension that does not relate to any particular graph and can be used with any base graph. The generic graph extension operates with data by using the *mapped cache extensions*, which are cache extensions that are not bound to any particular DAC and can extend any DAC.

To connect the mapped cache extensions to a particular DAC, you use a *mapping class*, which maps the fields of a mapped cache extension to the fields of a DAC. To connect the generic graph extension to a particular base graph, in the base graph, you define an *implementation class*, which inherits the generic graph extension. The following diagram shows in yellow rectangles the classes that you need to implement to reuse the business logic.



**Figure: Required classes**

These classes are described in detail in the sections below. (For details on the implementation of the classes, see [To Implement Reusable Business Logic](#) and [To Insert Reusable Business Logic That Has Already Been Declared](#).)

### Mapped Cache Extension

A mapped cache extension is an analog of a data access class (DAC) for a generic graph extension. In the mapped cache extension, you include the main fields that are used in the reusable business logic implementation. You map the fields of a mapped cache extension to the fields of a base DAC by using the mapping class, which is described in the following section. The mapped cache extension can also include the fields that are not mapped to any base DAC fields. (For details, see [Mapped Cache Extensions and the Application Database](#).)

The class of a mapped cache extension inherits from the `PXMappedCacheExtension` abstract class, which derives from `PXCacheExtension` and `IBqlTable`.

The declaration of a field of a mapped cache extension includes the following two required members, which are the same as the required members of a DAC field:

- A public abstract class (which is also referred to as *class field*).

You derive the class from the `IBqlField` interface and assign it a name that starts with a lowercase letter.

- A public virtual property (which is also referred to as *property field*).

You assign the property a name that starts with an uppercase letter. The system assigns the `PXMergeAttributes` attribute with `MergeMethod.Merge` to each field of a mapped cache extension automatically. If you define the `PXMergeAttributes` attribute for a field of a mapped cache extension explicitly, the explicitly defined attribute overrides the automatically defined. You can also define any other attributes for the property field of the mapped cache extension, or not define the attributes at all.

The following code shows an example of a mapped cache extension.

```
//Mapped cache extension
public class Document : PXMappedCacheExtension
{
    //BAccountID field
    public abstract class bAccountID : IBqlField
    {
    }
    protected Int32? _BAccountID;

    public virtual Int32? BAccountID
    {
        get
        {
            return _BAccountID;
        }
        set
        {
            _BAccountID = value;
        }
    }

    //CuryID field
    public abstract class curyID : IBqlField
    {
    }
    protected String _CuryID;

    public virtual String CuryID
    {
        get
        {
            return _CuryID;
        }
        set
        {
            _CuryID = value;
        }
    }

    ...
}
```

## Mapping Class

A mapping class is a protected class that defines the mapping between the fields of a mapped cache extension and the fields of a DAC. In a generic graph extension, you declare a mapping class for each mapped cache extension that you need to use in the reusable logic implementation.

A mapping class implements the `IBqlMapping` interface, which has the following two properties:

- **Extension:** The mapped cache extension

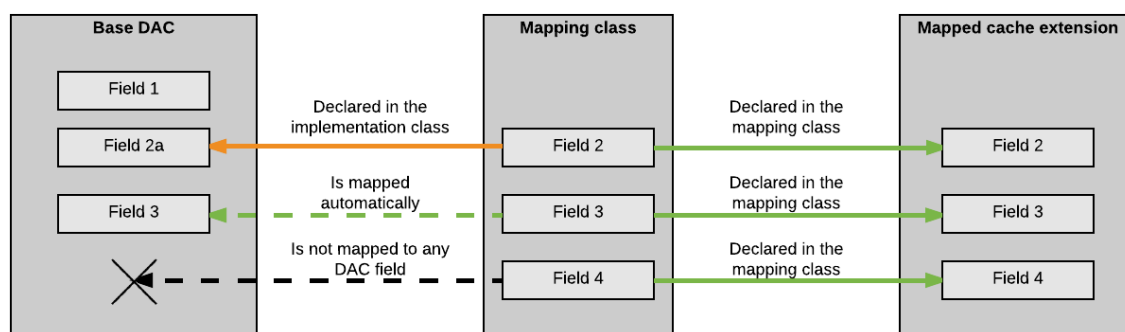
- **Table:** The DAC to which the extension is mapped

In the declaration of the mapping class, you also include declarations of the properties for each field of the mapped cache extension that you want to map to a field of the DAC, as the following code shows.

```
//A mapping class
protected class DocumentMapping : IBqlMapping
{
    public Type Extension => typeof(Document);
    protected Type _table;
    public Type Table => _table;

    public DocumentMapping(Type table)
    {
        _table = table;
    }
    public Type BAccountID = typeof(Document.bAccountID);
    public Type CuryInfoID = typeof(Document.curyInfoID);
    public Type CuryID = typeof(Document.curyID);
    public Type DocumentDate = typeof(Document.documentDate);
}
```

If the name of a property field of the DAC is the same as the name of the mapping class property, the DAC field will be automatically mapped to the field of the mapped cache extension by the implementation class (which is described below). If the name of a property field of the DAC differs from the name of the mapping class field, you redefine the mapping manually in the implementation class. If no field in the DAC has the name of the mapping class field, and no mapping is defined in the implementation class, the field of the mapped cache extension is not mapped to any base DAC field, as shown in the following diagram.



**Figure: Mapping declaration**

### Generic Graph Extension

A generic graph extension is a `public abstract` class that encapsulates business logic that can be used in multiple places of an Acumatica ERP application or an Acumatica Framework-based application. The class inherits from the `PXGraphExtension<TGraph>` class, as the following code shows.

```
public abstract class MultiCurrencyGraph<TGraph, TPrimary> :
    PXGraphExtension<TGraph>
    where TGraph : PXGraph
    where TPrimary : class, IBqlTable, new()
{
}
```

In the generic graph extension, you declare the following items:

- The mapping classes.

- The `protected` abstract methods that return the mapping classes. You have to override these methods in the implementation class.
- The views that can have either mapping-based declaration or standard declaration. You declare a mapping-based view by using the `PXSelectExtension<Table>` class.
- The event handlers, actions, and other methods.

### Implementation Class

An implementation class defines the implementation of a generic graph extension for a particular graph. You declare the implementation class as a class that derives from the generic graph extension class with the following type parameters:

- The base graph to which you add reusable logic
- The main DAC of the primary data view of the base graph

In this class, you can override the mapping defined by the mapping class, override other the methods of the base class, and insert your own views, methods, and event handlers, as the following code shows.

```
public class MultiCurrency : MultiCurrencyGraph<OpportunityMaint, CROpportunity>
{
    protected override DocumentMapping GetDocumentMapping()
    {
        return new DocumentMapping(typeof(CROpportunity))
        {
            DocumentDate = typeof(CROpportunity.closeDate)
        };
    }

    protected override CurySourceMapping GetCurySourceMapping()
    {
        return new CurySourceMapping(typeof(Customer));
    }

    public PXSelect<CRSetup> crCurrency;
    protected PXSelectExtension<CurySource> SourceSetup =>
        new PXSelectExtension<CurySource>(crCurrency);

    protected virtual CurySourceMapping GetSourceSetupMapping()
    {
        return new CurySourceMapping(typeof(CRSetup))
        {
            CuryID = typeof(CRSetup.defaultCuryID),
            CuryRateTypeID = typeof(CRSetup.defaultRateTypeID)
        };
    }

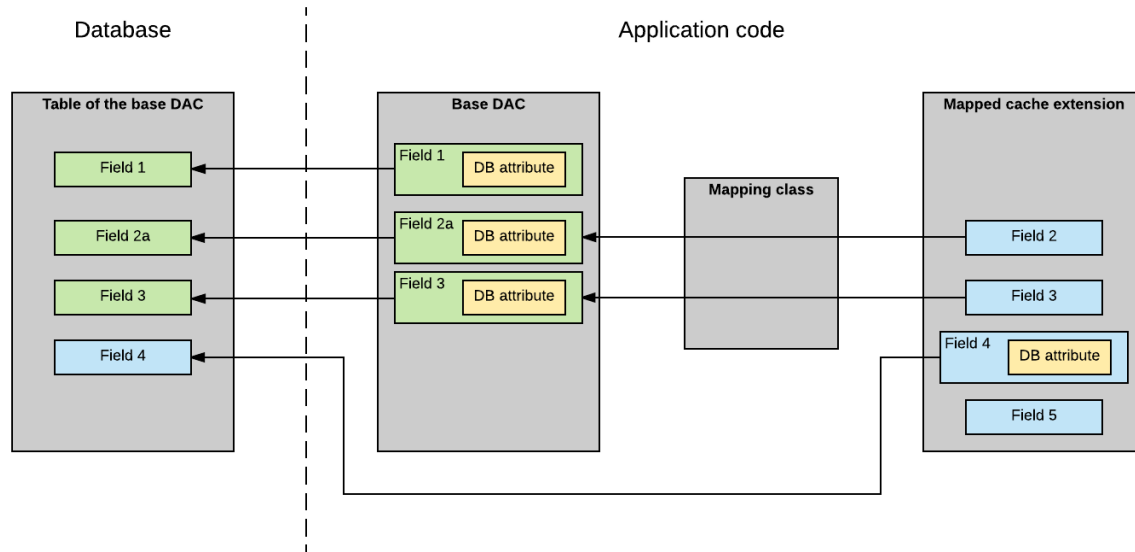
    protected override CurySource CurrentSourceSelect()
    {
        ...
    }
}
```

## Mapped Cache Extensions and the Application Database

The fields of a mapped cache extension that are mapped to the fields of a base data access class (DAC) are used by the system to work with the database columns to which the fields of the base DAC are bound.

If a mapped cache extension includes fields that are bound to database columns (with the type attributes that are derived from the `PXDBFieldAttribute` class, such as `PXDBString`), the database table that corresponds to the base DAC must contain these fields of the mapped cache extension. That

is, the database table must include the fields bound to a database column that are defined both in the base DAC and the mapped cache extension, as shown in the following diagram.

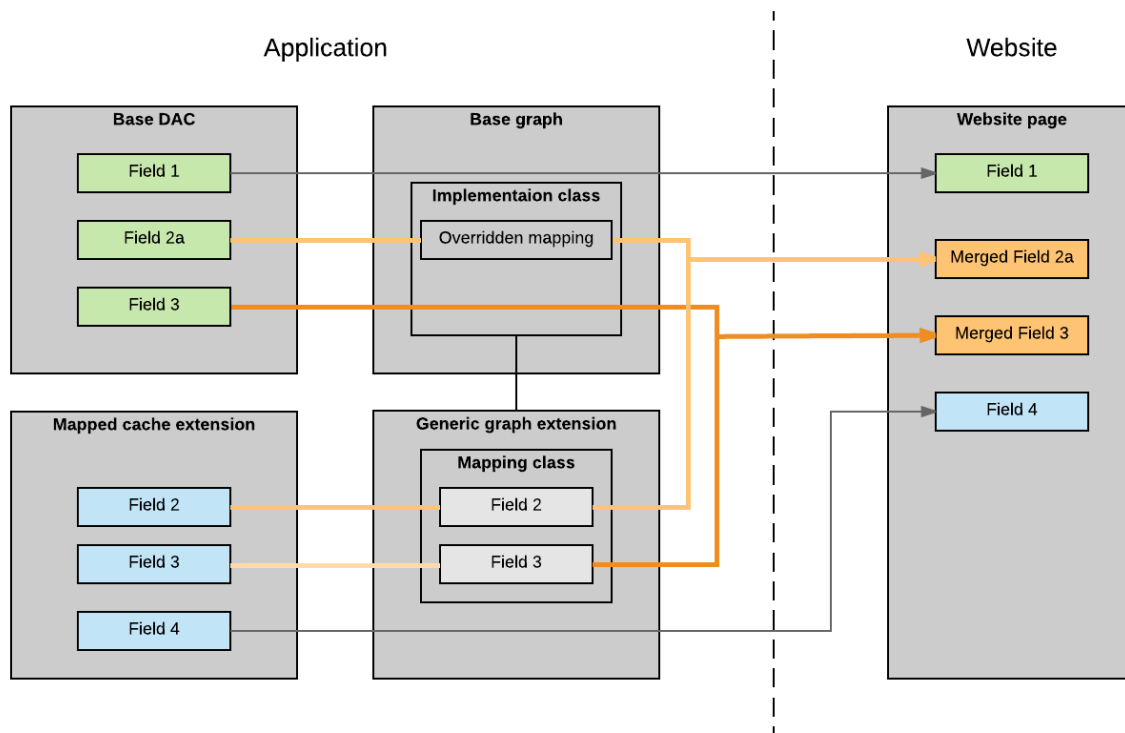


**Figure: Database schema for a DAC and mapped cache extension**

## Reusable Business Logic and the Application Website

If the field of a mapped cache extension is mapped to a field of the base data access class (DAC), you can use the merged field (that is, the base field that has merged attributes of the base field and the mapped cache extension field) to configure the UI elements of the website page. If the field of the mapped cache extension is not mapped to a field of the base DAC, you can use the field of the mapped cache extension in the website page, as shown in the following diagram.





**Figure: Fields on the website page**

The actions that are defined in the base graph, generic graph extension, and implementation class are automatically added by the system on the website page, as shown in the following diagram. The implementation class can override the actions declared in the generic graph extension.

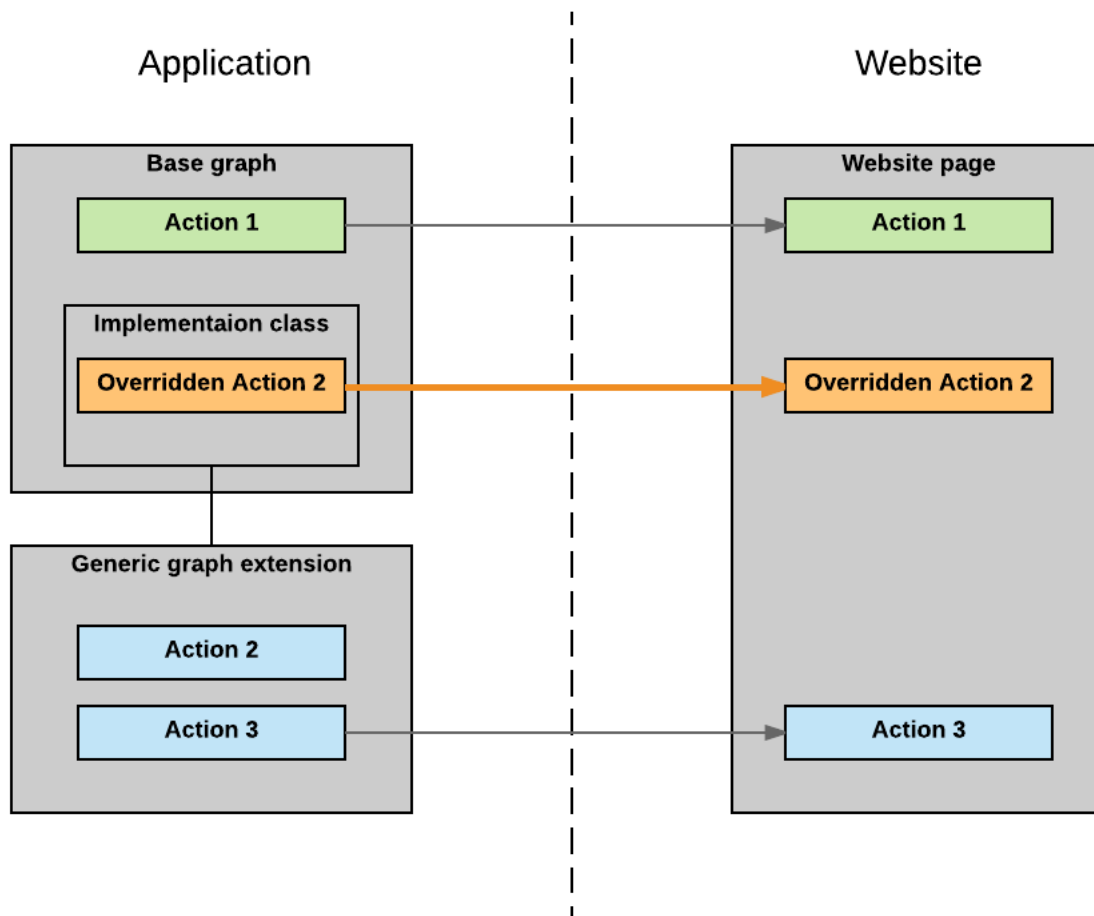


Figure: Actions on a website page

## Use of Generic Graph Extensions by the System

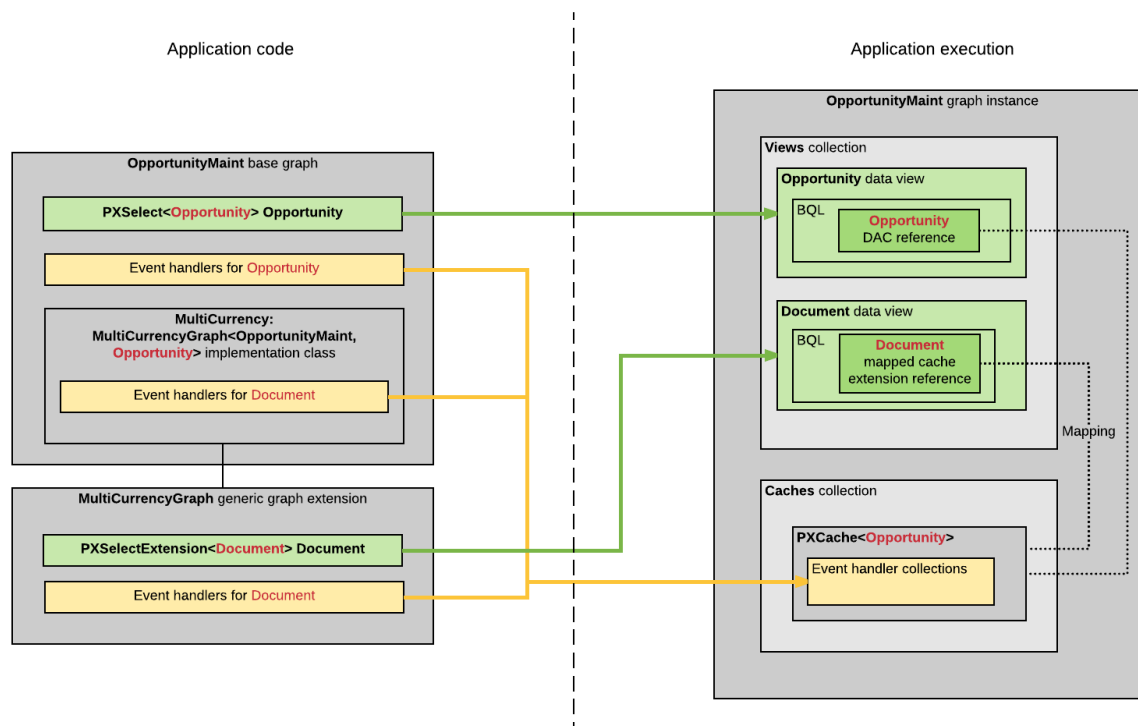
In this topic, you will learn how Acumatica ERP or an Acumatica Framework-based application works with generic graph extensions.

### Initialization of a Graph Instance That Includes Reusable Logic

During the initialization of a graph instance that includes reusable logic, the system adds to the collection of data views the data views that are declared in the base graph, including those that are declared in the implementation class or the generic graph extension from which the implementation class inherits. The mapping classes define the `PXCache<Table>` objects in which the mapping-based views keep data records.

Event handlers that are declared in the base graph, implementation class, and generic graph extension are added by the system to the collections of event handlers in the corresponding `PXCache` object. The event handlers defined for the fields or rows of the mapped cache extension are added to the `PXCache` object of the base DAC type to which the mapped cache extension is mapped.

The following diagram illustrates the initialization of a sample `OpportunityMaint` graph instance that includes reusable logic.



**Figure: Initialization of a graph instance**

## Generic Graph Extensions Declared in Acumatica ERP

The source code of Acumatica ERP includes the definitions of the generic graph extensions, which are described in detail in the following sections. These graph extensions are declared in the `PX.Objects.Extensions` namespace. You can use these generic graph extensions if you want to include the implemented functionality in the forms of your application. For details on how to include this functionality in your application, see [To Insert Reusable Business Logic That Has Already Been Declared](#).

### Multi-Currency Extension

If you need to work with multiple currencies on a form, you can insert an implementation of the `MultiCurrencyGraph<TGraph, TPrimary>` abstract class in the graph that provides business logic for the form. For more information on the use of multiple currencies in the system, see [Currency Management](#) in the Financial Management Guide.

The `MultiCurrencyGraph<TGraph, TPrimary>` class works with the following mapped cache extensions:

- **Document:** Represents a document that supports multi-currency
- **CurySource:** Contains the information on the currency source

For more information on these classes, see the API Reference.

### Sales Price Extension

If you need to work with multiple price lists on a form, you can insert an implementation of the `SalesPriceGraph<TGraph, TPrimary>` abstract class in the graph that provides business logic for the form. For more information on sales prices, see [Managing Sales Prices](#) in the Financial Management Guide.

The `SalesPriceGraph<TGraph, TPrimary>` class works with the following mapped cache extensions:

- **Document:** Represents a document that supports multiple price lists
- **Detail:** Represents a detail line of the document
- **PriceClassSource:** Provides information about the source of the price class

For more information on these classes, see the API Reference.

### Discount Extension

If you need to work with discounts on a form, you can insert an implementation of the `DiscountGraph<TGraph, TPrimary>` abstract class in the graph that provides business logic for the form. For more information on discounts, see [Managing Customer Discounts](#) in the Financial Management Guide.

The `DiscountGraph<TGraph, TPrimary>` class works with the following mapped cache extensions:

- **Document:** Represents a document that supports discounts
- **Detail:** Represents a detail line of the document
- **Discount:** Provides information about the discount

For details on these classes, see the API Reference.

### Sales Tax Extension

If you need to apply sales taxes to amounts of a form, you can insert an implementation of the `TaxGraph<TGraph, TPrimary>` abstract class in the graph that provides business logic for the form. For more information on taxes in the system, see [Taxes](#) in the Financial Management Guide.

The `TaxGraph<TGraph, TPrimary>` class works with the following mapped cache extensions:

- **Document:** Represents a document that supports sales taxes.
- **Detail:** Represents a detail line of the document.
- **TaxTotal:** Represents the tax total amount
- **TaxDetail:** Represents a tax detail line

For detailed descriptions of the classes, see the API Reference.

## To Insert Reusable Business Logic That Has Already Been Declared

In this topic, you can find information about how to insert an already-declared generic graph extension in the application code.

### To Add a Generic Graph Extension to a Graph

1. Review the generic graph extension that provides the business logic that you want to reuse as follows:
  - a. Identify the mapped cache extensions the generic graph extension works with and the list of their fields, and decide whether the default mapping (which is defined by the mapping class of the generic graph extension) is suitable for the base data access class (DAC) that you are going to use.
  - b. Identify the fields of the mapped cache extension that are bound to columns of a database table, and make sure the database table that corresponds to the base DAC includes the columns to store the data from the mapped cache extension.
2. In the code of the graph you need to add the reusable business logic to, add the public implementation class that derives from the generic graph extension of the needed type. Use the following types in the type parameters of the generic graph extension:

- The base graph to which you add reusable logic
- The main DAC of the primary data view of the base graph

In the following code, the `MultiCurrency` class extends the `OpportunityMaint` graph. The `MultiCurrency` class derives from public abstract class `MultiCurrencyGraph<TGraph, TPrimary> : PXGraphExtension<TGraph>`.

```
public class MultiCurrency : MultiCurrencyGraph<OpportunityMaint,
    CROpportunity>
{
}
```

3. In the added class, override the abstract methods of the generic graph extension as follows:
  - In the overridden methods of the generic graph extension that return the mapping classes, either use the default mapping of the fields of the mapped cache extension to the fields of the base DAC or adjust the mapping.
  - In the other overridden methods, implement the required business logic. For details on the implementation of the methods in the generic graph extension declared in Acumatica ERP, see API Reference.
4. In the added class, adjust the reused business logic by doing any of the following:
  - Override other methods of the base class.
  - Add your own views, methods, and event handlers.

The following code shows a sample implementation of the `MultiCurrency` class, which reuses the multi-currency business logic defined in the `MultiCurrencyGraph` generic graph extension.

```
public class MultiCurrency : MultiCurrencyGraph<OpportunityMaint,
    CROpportunity>
{
    protected override DocumentMapping GetDocumentMapping()
    {
        return new DocumentMapping(typeof(CROpportunity))
        {
            DocumentDate = typeof(CROpportunity.closeDate)
        };
    }

    protected override CurySourceMapping GetCurySourceMapping()
    {
        return new CurySourceMapping(typeof(Customer));
    }

    public PXSelect<CRSetup> crCurrency;
    protected PXSelectExtension<CurySource> SourceSetup =>
        new PXSelectExtension<CurySource>(crCurrency);

    protected virtual CurySourceMapping GetSourceSetupMapping()
    {
        return new CurySourceMapping(typeof(CRSetup))
        {
            CuryID = typeof(CRSetup.defaultCuryID),
            CuryRateTypeID = typeof(CRSetup.defaultRateTypeID)
        };
    }

    protected override CurySource CurrentSourceSelect()
    {
        ...
    }
}
```

## To Sort Multiple Generic Graph Extensions

If you need to add multiple generic graph extensions to a graph, you need to define the order in which the extensions are applied.

To define the order in which the generic graph extensions are applied, add the class inherited from the `Autofac.Module` class and implement the sorting of the generic class extensions, as the following code shows.

```
public class ServiceRegistration : Module
{
    protected override void Load(ContainerBuilder builder)
    {
        builder.ActivateOnApplicationStart<ExtensionSorting>();
    }
    private class ExtensionSorting
    {
        private static readonly Dictionary<Type, int> _order = new Dictionary<Type, int>
        {
            {typeof(MultiCurrency), 4},
            {typeof(SalesPrice), 3},
            {typeof(Discount), 2},
            {typeof(SalesTax), 1},
        };
        public ExtensionSorting()
        {
            ...
        }
    }
}
```

## To Implement Reusable Business Logic

If you need to use the same business logic in multiple modules of your Acumatica ERP application or an Acumatica Framework-based application and this logic is not included in the source code of Acumatica ERP, you can define your own generic graph extensions, as described in this topic. For the list of the predefined generic graph extensions, see [Generic Graph Extensions Declared in Acumatica ERP](#).

### To Create a Generic Graph Extension

1. In the code of your application, define the mapped cache extensions, which inherit from the `PXMappedCacheExtension` abstract class. For details on the mapped cache extensions, see [Mapped Cache Extension](#).

```
//Mapped cache extension
public class Document : PXMappedCacheExtension
{
    //BAccountID field
    public abstract class bAccountID : IBqlField
    {
    }
    protected Int32? _BAccountID;

    public virtual Int32? BAccountID
    {
        get
        {
            return _BAccountID;
        }
        set
        {
            _BAccountID = value;
        }
    }

    //CuryID field
```

```

public abstract class curyID : IBqlField
{
}
protected String _CuryID;

public virtual String CuryID
{
    get
    {
        return _CuryID;
    }
    set
    {
        _CuryID = value;
    }
}

...
}

```

2. In the code of your application, define the generic graph extension as follows:

- a. Define a class inherited from the `PXGraphExtension<TGraph>` class. The following code shows a declaration of a generic graph extension.

```

public abstract class MultiCurrencyGraph<TGraph, TPrimary> :
    PXGraphExtension<TGraph>
    where TGraph : PXGraph
    where TPrimary : class, IBqlTable, new()
{
}

```

- b. In the generic graph extension, for each mapped cache extension that you defined in the first step, declare the `protected` mapping class, as shown in the following code. For details on the mapping classes, see [Mapping Class](#).

```

//A mapping class
protected class DocumentMapping : IBqlMapping
{
    public Type Extension => typeof(Document);
    protected Type _table;
    public Type Table => _table;

    public DocumentMapping(Type table)
    {
        _table = table;
    }
    public Type BAccountID = typeof(Document.bAccountID);
    public Type CuryInfoID = typeof(Document.curyInfoID);
    public Type CuryID = typeof(Document.curyID);
    public Type DocumentDate = typeof(Document.documentDate);
}

```

- c. In the generic graph extension, for each mapping class, declare the `protected` abstract method that returns the mapping class, as shown in the following code.

```

protected abstract DocumentMapping GetDocumentMapping();

```

- d. In the generic graph extension, define the views that use the mapped cache extensions, as the following code shows. To define each view, you use the `PXSelectExtension<Table> : PXSelectBase<Table>` class, where `Table` is a mapped cache extension.

```

//A view that uses the mapped cache extension
public PXSelectExtension<Document> Documents;

```



: In the generic graph extension, you can define standard views as well as the views that use the mapped cache extensions.

- e. In the generic graph extension, define the reusable event handlers, as the following code shows.

```
protected virtual void _(
    Events.FieldUpdated2<Document.documentDate, Document> e)
{
    if (e.Row == null) return;
    CurrencyInfoAttribute.SetEffectiveDate<Document.documentDate>(
        Documents.Cache,
        new PXFieldUpdatedEventArgs(e.Row, e.OldValue, e.ExternalCall));
}
```

- f. In the generic graph extension, implement any other business logic that you want to reuse, such as filters and actions.

Once you have defined the mapped cache extensions and the generic graph extension, you can insert reusable business logic to any part of your application, as described in [To Insert Reusable Business Logic That Has Already Been Declared](#).



# Troubleshooting Acumatica Framework-Based Applications

---

In this part of the guide, you can find information about troubleshooting the applications based on Acumatica Framework.

## In This Part

- [To Debug Acumatica Framework-Based Applications](#)

## To Debug Acumatica Framework-Based Applications

---

This topic describes how to link an Acumatica Framework-based application site to the database and start the Acumatica Framework application in debug mode.

### To Debug an Application from Visual Studio

1. In Visual Studio, open the solution of your Acumatica Framework-based application.
2. In the `Site` folder of the solution, open the `web.config` file.
3. In the `connectionStrings` section of the file, modify the connection string by specifying the credentials to your development database as follows:
  - For a locally installed Microsoft SQL Server that uses SQL Server authentication (line breaks are for display purposes only)

```
<connectionStrings>
  <remove name="ProjectX" />
  <add name="ProjectX" providerName="System.Data.SqlClient"
    connectionString="Data Source=(local);Initial
    Catalog=Project_Catalog;
    User Id=User_ID; Password=User_Password"
</connectionStrings>
```

- For a locally installed Microsoft SQL Server that uses Windows authentication (line breaks are for display purposes only)

```
<connectionStrings>
  <remove name="ProjectX" />
  <add name="ProjectX" providerName="System.Data.SqlClient"
    connectionString="Data Source=(local);Initial
    Catalog=Project_Catalog;
    Integrated Security=True"/>
</connectionStrings>
```

- For a remote Microsoft SQL Server that uses SQL Server authentication (line breaks are for display purposes only)

```
connectionStrings>
  <remove name="ProjectX" />
  <add name="ProjectX" providerName="System.Data.SqlClient"
    connectionString="Data Source=Server_Name; Initial
    Catalog=Project_Catalog;
    User Id=User_ID; Password=User_Password"
</connectionStrings>
```

4. In the `system.web` section of the file, set the `debug` attribute of the `compilation` element to `True`, as shown in the following example.

```
<compilation debug="True" defaultLanguage="c#"
    numRecompilesBeforeAppRestart="9999" targetFramework="4.7.1">
```

5. In Visual Studio, right-click the `Site` folder of the solution, and click **Set as StartUp Project**.
6. Right-click the `Main.aspx` file in the `Site` folder, and click **Set as Start Page**.
7. Optional: If you need to debug a server error that throws an exception, do the following:
  - a. On the toolbar, click **Debug > Windows > Exception Settings**.
  - b. In the **Exception Settings** panel, which opens, expand **Common Language Runtime Exceptions**, and select the check box for the exception that is thrown (such as **System.ArgumentOutOfRangeException**).
8. Run the solution in the Debug mode.

# Glossary

---

The following table contains definitions of the basic terms used in Acumatica Framework.

Term	Definition
<b>Action</b>	An interface for executing a specific operation with data that is implemented in a business logic controller (BLC). An action is represented by the corresponding button on the user interface.
<b>Acumatica Cloud xRP Platform</b>	The platform for the development of cloud ERP applications (such as Acumatica ERP and customizations of it), the mobile application for Acumatica ERP, and applications integrated with Acumatica ERP by means of the web services API.
<b>Acumatica Customization Platform</b>	The part of the Acumatica Cloud xRP Platform that provides customization tools for the development of applications embedded in Acumatica ERP (also called customizations of Acumatica ERP).
<b>Acumatica Framework</b>	The part of the Acumatica Cloud xRP Platform that provides the platform API, web controls, and other tools for building ERP applications.
<b>Acumatica Framework Templates</b>	A set of Visual Studio templates provided as a part of Acumatica Framework for creating application pages and business logic controllers.
<b>Acumatica Framework-based application</b>	An application created by means of Acumatica Framework tools.
<b>Analytical report</b>	A report created with the Analytical Report Manager. For details on this tool, see <a href="#">Analytical Report Manager</a> .
<b>Bound field</b>	A data field that represents a column from a database table. Compare to <b>Unbound field</b> .
<b>BQL statement</b>	A generic BQL class specialization that represents a specific query to the database. The type parameters specified in the BQL statement are BQL operator classes and DACs.
<b>Business logic controller (BLC)</b>	A stateless controller class that is intended for the execution of business logic on a particular application page. A BLC (also called a <i>graph</i> ) is derived from the <code>PXGraph</code> generic class.
<b>Business query language (BQL)</b>	A set of generic classes for querying data records from the database.
<b>Cache</b>	A collection of modified data records from the same table stored in the user session and shared between requests.
<b>Custom report</b>	A report created on the custom report form.
<b>Customization of Acumatica ERP</b>	A modification of the user interface, business logic, and the database scheme without recompilation and re-installation of Acumatica ERP. This modification is packed in a customization project.
<b>Customization project</b>	A container that holds the changes you have made during a particular customization of Acumatica ERP.
<b>DAC field</b>	See <b>Field</b> .
<b>Data access class (DAC)</b>	A class that represents a database table.
<b>Data entry form</b>	An application webpage that is used for the input of business documents.
<b>Data member</b>	A data view specified as the data source for a container of UI controls (a form, a tab, or a grid).

<b>Data record</b>	A specific record retrieved from the database or created in code and wrapped in a DAC instance.
<b>Data view</b>	A BQL statement that the BLC uses to access and manipulate data. A developer defines a data view in code by using <a href="#">PXSelect</a> classes.
<b>Datasource control</b>	A service control on a page that is used to bind the page to a particular BLC. This control represents the page toolbar, which contains action buttons.
<b>Embedded application</b>	See <b>Customization of Acumatica ERP</b> .
<b>Event</b>	A way to provide notifications from Acumatica Framework to the application. Most business logic is implemented in event handlers.
<b>Event handler</b>	A method that is invoked by Acumatica Framework when the corresponding event is raised.
<b>Field (DAC field)</b>	A part of the DAC definition that typically represents a database column. A DAC field consists of an abstract class that is used to refer to the field in BQL and a property holding the actual field value.
<b>Form</b>	See <b>Webpage</b> .
<b>Graph</b>	See <b>Business logic controller</b> .
<b>Inquiry form</b>	An application webpage that displays a list of data records selected by the specified filter.
<b>Integrated application</b>	A third-party application integrated with Acumatica ERP by means of web services API.
<b>Maintenance form</b>	An application helper webpage that is used for the input of data on the data entry and processing pages.
<b>Mobile API</b>	The API that is used for customization of the Acumatica mobile application. For a description of the API, see <a href="#">Mobile Site Map Reference</a> .
<b>Multitenant application</b>	An application in which multiple tenants use the same Acumatica Framework-based application. For each tenant, the website looks identical and provides the same business logic. However, each tenant has exclusive access to the tenant's individual data and can have restricted access to the data of other tenants.
<b>Page template</b>	A Visual Studio template that is provided by Acumatica Framework and used for creating application pages.
<b>Platform API</b>	The API that is used to develop Acumatica Framework-based applications and customizations of Acumatica ERP.
<b>Primary BLC</b>	The BLC that corresponds to the default editing page of the data record. This BLC is specified in the <a href="#">PXPrimaryGraph</a> attribute.
<b>Primary DAC</b>	The first data access class specified in a BQL statement.
<b>Primary data view</b>	The first data view defined in a business logic controller.
<b>Processing form</b>	An application webpage that provides mass processing operations.
<b>Report Designer</b>	A visual editor for creating report forms and printable pages.
<b>Report form</b>	An RPX page created in Report Designer that defines the form used for generating reports in the application.
<b>Screen</b>	See <b>Webpage</b> .
<b>Setup form</b>	An application webpage that provides the configuration parameters for the application.

<b>Unbound field</b>	A data field that exists only on the model level in a DAC definition, and that is not bound to a column of the database table. Compare to <b>Bound field</b> .
<b>Webpage</b>	A page that provides the UI of the application. Each webpage used in the application is a declarative ASPX page created from one of the Acumatica Framework Templates or a report form.
<b>Web services API</b>	The API for development of applications integrated with Acumatica ERP through SOAP or REST.