This PDF is generated from authoritative online content, and is provided for convenience only. This PDF cannot be used for legal purposes. For authoritative understanding of what is and is not supported, always use the online content. To copy code samples, always use the online content.

Platform SDK Developer's Guide

Platform SDK 8.1.3

12/29/2021

Table of Contents

Welcome to the Developer's Guide!	**4**
Introductory Topics	**6**
Introducing the Platform SDK	7
Architecture of the Platform SDK	12
Connecting to a Server	14
Configuring Platform SDK Channel Encoding for String Values	25
Using the Warm Standby Application Block	27
Event Handling	45
Setting up logging in Platform SDK	56
Advanced Platform SDK Topics	**60**
Secure connections using TLS	61
Quick Start	65
Using the Platform SDK Commons Library	68
Using the Application Template Application Block	80
Configuring TLS Parameters in Configuration Manager	86
Using and Configuring Security Providers	106
OpenSSL Configuration File	117
Use Cases	124
Lazy Parsing of Message Attributes	127
Server-Specific Overviews	**131**
Telephony (T-Server)	132
List of TLib Functions	143
List of TLib Datatypes	144
List of TLib Unstructured Data	145
Configuration	146
Connecting Using UTF-8 Character Encoding	157
Change Password On Next Login	160
Getting the Last Login Info	161
Using the Configuration Object Model Application Block	163
Introduction to the Configuration Layer Objects	200
List of Configuration Layer Objects	203
List of Configuration Layer Enumerations	204
Stat Server	205
Interaction Server	219
Universal Contact Server	224

Creating an E-Mail	232
Chat	245
Outbound	252
Management Layer	254
LCA Hang-Up Detection Support	268
Handle Application "Graceful Stop" with the LCA Protocol	274
Routing Server	276
Component Overviews	**288**
Using the Log Library	289
Using the Switch Policy Library	306
Legacy Topics	**324**
Using the Message Broker Application Block	325
Event Handling Using the Message Broker Application Block	337
Using the Protocol Manager Application Block	353
Connecting to a Server Using the Protocol Manager Application Block	374
Explicitly Choosing a Netty or Mina Connection Layer	383
Platform SDK Resources	**385**

Welcome to the Developer's Guide!

This guide offers a collection of articles that will help you get started with Platform SDK development.

For detailed information about the Platform SDKs, please refer to the Platform SDK API Reference for your specific release.

Introductory Topics

The following articles give information about common Platform SDK functionality and protocol usage that all developers should be aware of:

- Introducing the Platform SDK
- Architecture of the Platform SDK
- Connecting to a Server
- Configuring Platform SDK Channel Encoding for String Values
- Using the Warm Standby Application Block
- Event Handling
- Setting Up Logging in Platform SDK

Advanced Platform SDK Topics

The following articles provide details about advanced Platform SDK features you may want to take advantage of:

- Secure Connections Using TLS
- Lazy Parsing of Message Attributes

Server-Specific Overviews

- Telephony (T-Server)
  - List of TLib Functions
  - List of TLib Datatypes
  - List of TLib Unstructured Data
- Configuration
  - Connecting Using the UTF-8 Enconding
  - Change Password On Next Login
  - Getting the Last Login Info
  - Using the Configuration Object Model Application Block

    - Introduction to Configuration Layer Objects
    - List of Configuration Layer Objects
    - List of Configuration Layer Enumerations
- Stat Server
- Interaction Server
- Universal Contact Server
    - Creating an E-Mail
- Chat
- Outbound
- Management Layer
    - LCA Hang-Up Detection Support
    - Handle Application "Graceful Stop" with the LCA Protocol
- Routing Server

Component Overviews

- Using the Log Library
- Using the Switch Policy Library

Legacy Topics

Topics in this section are no longer applicable for new development, but are maintained here for backwards compatibility.

- Using the Message Broker Application Block
- Event Handling Using the Message Broker Application Block
- Using the Protocol Manager Application Block
- Connecting to a Server Using the Protocol Manager Application Block
- Explicitly Choosing a Netty or Mina Connection Layer

Additional Resources

The following page contains reference materials that may be useful when developing applications with Platform SDK, including links to related documents and downloadable code samples.

- Platform SDK Resources

Introductory Topics

The following articles give information about common Platform SDK functionality and protocol usage that all developers should be aware of:

- Introducing the Platform SDK
- Architecture of the Platform SDK
- Connecting to a Server
- Configure Platform SDK Channel Encoding for String Values
- Using the Warm Standby Application Block
- Event Handling
- Setting Up Logging in Platform SDK

Introducing the Platform SDK

The Platform SDK exposes the protocols of Genesys servers as an API. This means you can write .NET and Java applications that communicate with these servers in their native protocols.

You can use the Platform SDK to do two main things:

- Establish and maintain a connection to each Genesys server used by your application
- Send and receive messages to and from each of these Genesys servers

In addition to enabling these two basic functions, the Platform SDK ships with *application blocks*, which have been built on top of the Platform SDK in order to provide simple yet high-performance ways to do things like configuring warm standby settings for your connections and working with configuration objects.

The following image shows the relationship between the Platform SDK protocol objects and the servers each of them connects with.

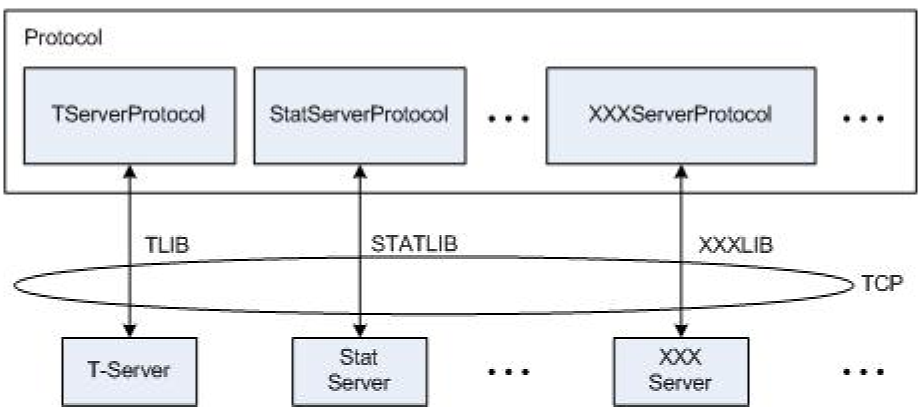

Each protocol object subclasses `ClientChannel`, which in turn subclasses `DuplexChannel` and implements the `Protocol` interface. This means they all share a common interface to the Genesys servers. The protocol objects communicate with the corresponding Genesys servers over a TCP connection, with each one using the native protocol of the server it connects with. For example, the `TServerProtocol` object communicates over TCP with a T-Server, using the TLIB protocol that is native to the T-Server.

As mentioned above, the Platform SDK also includes reusable production-quality application blocks that can be dropped into your code to provide simple yet high-performance ways to carry out important functions that are commonly needed by applications that communicate with Genesys servers.

As shown below, there are two main types of application blocks.

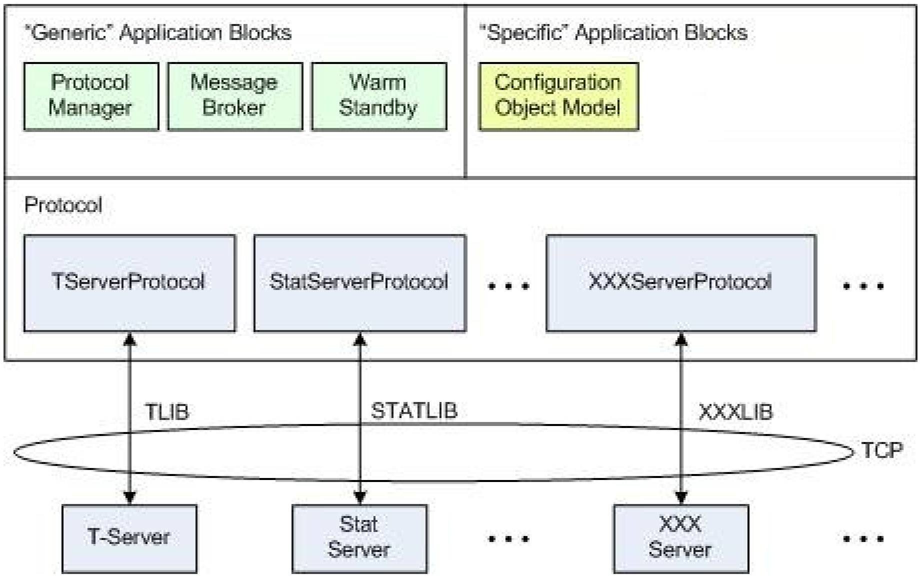

The generic Warm Standby Application Block can be used to great effect either in conjunction with Hot Standby or on its own. Generic Protocol Manager and Message Broker application blocks have been marked as legacy items in the 8.1.1 release of Platform SDK. Documentation for legacy application blocks is still included to support backwards compatibility, but improvements to the Platform SDK code mean they are no longer recommended for use in new development.

The Configuration Object Model Application Block provides more specific functionality, and will only be beneficial for specific types of applications. This application block makes it easy to work with objects in the Genesys Configuration Layer, and is not required unless you are writing an application that requires this functionality.

Finally, the Platform SDK includes additional components designed to make development of custom applications easier. These components offer support for useful features such as customized logging or switch abstraction.

The Protocols

The Platform SDK is divided into separate “protocols.” Each component works with one or more of Genesys servers.

The following table shows the servers each of the Platform SDK protocols connects with, and gives the names of the native protocols that are used to communicate with each server.

Platform SDK Protocol Name	Genesys Servers	Native Protocols
Configuration Platform SDK	Configuration Server	CFGLIB
Contacts Platform SDK	Universal Contact Server	UCS Protocol
Management Platform SDK	• Message Server	• GMESSAGELIB

Platform SDK Protocol Name	Genesys Servers	Native Protocols
	• Solution Control Server • Local Control Agent	• SCSLIB • LCALIB
Open Media Platform SDK	Interaction Server	ITX, ESP
Outbound Contact Platform SDK	Outbound Contact Server	• CMLIB • OCS-Desktop Protocol
Routing Platform SDK	• Universal Routing Server • Custom Server	• URS Protocol • Custom Server Protocol
Statistics Platform SDK	Stat Server	STATLIB
Voice Platform SDK	T-Servers	• TLIB • Preview Interaction Protocol
Web Media Platform SDK	• Chat Server • E-Mail Server Java • Callback Server	• MCR Chat Lib • MCR E-Mail Lib • MCR Callback Lib • ESP E-Mail Lib

Configuration Platform SDK

The Configuration Platform SDK enables you to build applications that use the services of the Genesys Configuration Server. This allows these applications to either query on objects in the Configuration Layer of your Genesys environment or to add, modify, and delete information about those objects, while taking advantage of an environment in which Configuration Server carries out several important administrative functions.

Contacts Platform SDK

The Contacts Platform SDK allows you to build applications that view, or interact with, the contact information for your contact center. This SDK accesses information directly from Universal Contact Server, allowing you to design applications that access contact information when dealing with multimedia interactions such as chat or e-mail, for example.

Management Platform SDK

The Management Platform SDK enables you to write applications that interact with Message Server, Solution Control Server, and Local Control Agents.

Open Media Platform SDK

With the Open Media Platform SDK, you can build client applications that feed open media interactions into your Genesys environment, or applications that act as custom media servers to perform external service processing (ESP) on interactions that have already entered it.

Outbound Contact Platform SDK

The Outbound Contact Platform SDK can be used to build applications that allow you to manage outbound campaigns.

Routing Platform SDK

The Routing Platform SDK allows you to write Java applications that combine logic from your custom application with the router-based logic of URS, in order to solve many common interaction-related tasks.

Statistics Platform SDK

With the Statistics Platform SDK, you can build applications that use the services of Stat Server in order to solicit and monitor statistics from your Genesys environment.

Stat Server tracks information about customer interaction networks (contact center, enterprise-wide, or multi-enterprise telephony and computer networks). It also converts the data accumulated for directory numbers (DNs), agents, agent groups, and non-telephony–specific object types, such as e-mail and chat sessions, into statistically useful information.

Voice Platform SDK

The Voice Platform SDK enables you to design applications that monitor and handle voice interactions from a traditional or IP-based telephony device.

Web Media Platform SDK

The Web Media Platform SDK can be used to build applications that interact with Chat Server, E-Mail Server Java, and Callback Server through a web server interface.

The Application Blocks

Genesys application blocks are reusable production-quality components that provide specific functionality needed by a broad range of Genesys customers. They have been designed using industry best practices and provided with source code so they can be used “as is”, extended, or tailored if you need to. Please see the License Agreement for details.

> **Tip**
>
> If you have questions or suggestions about the application blocks, please contact us in the Genesys Forums.

Configuration Object Model Application Block

The Configuration Object Model (COM) Application Block provides a consistent and intuitive object model for applications that need to work with Configuration Server objects. Use the COM Application Block when you need to create, update, or delete Configuration Layer Objects.

Warm Standby Application Block

You can use the Warm Standby Application Block to switch to a backup server in case your primary server fails, in cases where you do not need to guarantee the integrity of existing interactions.

Message Broker Application Block (deprecated)

The Message Broker Application Block makes it easy for your applications to handle events in an efficient way.

> Important
>
> This application block is considered a legacy product starting with release 8.1.1. Documentation is provided for backwards compatibility, but new development should consider using the improved method of message handling.

Protocol Manager Application Block (deprecated)

The Protocol Manager Application Block allows for simplified communication with more than one server. It takes care of opening and closing connections to many different servers, as well as handling the reconfiguration of high availability connections.

> Important
>
> This application block is considered a legacy product starting with release 8.1.1. Documentation is provided for backwards compatibility, but new development should consider using the improved method of connecting to servers.

The Components

Additional components are included to provide useful functionality for creating custom applications with the Platform SDK, even if that doesn't necessarily involve communicating with Genesys servers.

Platform SDK Switch Policy Library

The Platform SDK Switch Policy Library can be used in applications that need to perform agent-related switch activity with a variety of T-Servers, without knowing beforehand what kinds of T-Servers will be used.

Platform SDK Log Library

The Platform SDK Log Library for .NET presents an easy-to-use API for logging messages in custom-built applications.

Architecture of the Platform SDK

The Platform SDKs enable you to write client or server applications that use messages to communicate with Genesys servers.

Each SDK has one or more `Protocol` objects that you can use in your client applications to establish communication with the appropriate server. These objects use the native protocols of the Genesys servers they are designed to work with.

From a conceptual standpoint, your application's `Protocol` object, will be communicating directly with the appropriate server using the server's protocol running on TCP/IP, as shown below.

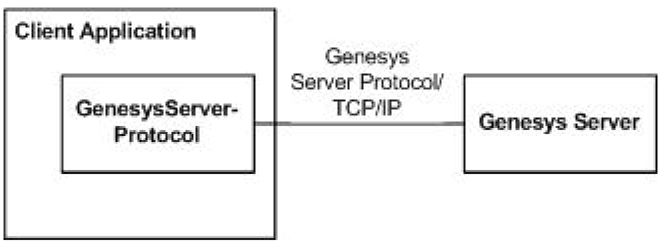

Once you have opened a connection to the server, you are ready to send and receive messages. The Platform SDK supports two message exchange patterns. In some cases, you will need to follow the Request/Response pattern. That is, you will send a message and wait for a response, as shown below.

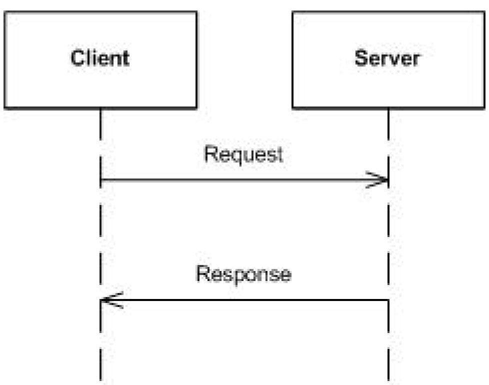

At other times, following the Unsolicited Event pattern, you simply need to wait for unsolicited messages of a certain type.

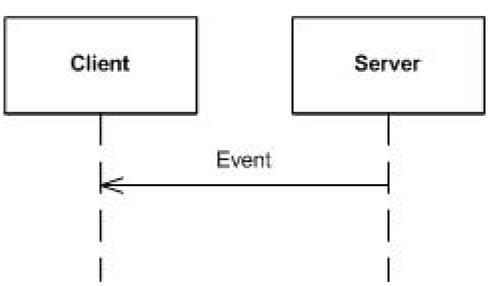

The messages you send will be in the form of `Request` classes, such as `RequestAgentLogin` or `RequestAnswerCall`. The messages you receive, whether solicited or not, will be in the form of `Event` classes, such as `EventAck` or `EventDialing`.

As you can see, the architecture of the Platform SDKs is fairly simple — but you can use it to do some powerful things.

Connecting to a Server

Java

The applications you write with the Platform SDK need to communicate with one or more Genesys servers, so the first thing you need to do is create connections with these servers. You will have to reference libraries and add `import` statements to your project for each specific protocol you are working with. These steps are not explicitly described here because the files and packages required will vary depending on which protocols you plan to use.

> **Important**
> Starting with release 8.1.1, the Platform SDK uses Netty by default for the implementation of its transport layer. Therefore, your project will need to reference Netty as well.

Once you have connected to a server, you use that connection to exchange messages with the server. For details about sending and recieving messages to and from a server, refer to the event handling article.

Creating a Protocol Object

To connect to a Genesys server, you create an instance of the associated protocol class. As an example, this article will describe a conntection to a Genesys T-Server using the `TServerProtocol` class. (For different applications, please use this API Reference to check protocol details for the specific server that you wish to connect to.)

In order to create a protocol object, you will first need to create an Endpoint object which acts as a container for generic connection parameters. An Endpoint object contains, at a minimum, a server name, the host name where the server is running, and the port on which the server is listening. The server name will appear in logs but does not affect protocol behavior; it may be any name that is significant to you.

```
[Java]

Endpoint tserverEndpoint = new Endpoint("T-Server", TSERVER_HOST, TSERVER_PORT);
TServerProtocol tserverProtocol = new TServerProtocol(tserverEndpoint);
```

After creating your protocol object, you need to specify some connection parameters that are specific to that protocol. These parameters will differ depending on which server you are connecting to. Please check to the sections specific to the server that you wish to connect to for more information.

Once configuration is complete, you can open the connection to your server.

Opening a Synchronous Connection

The easiest way to open a connection to your server is to do it synchronously, which means that the method will block any additional processing until the server connection has either opened successfully or failed definitively. This is commonly used for non-interactive, batch applications. In this case, you can add code for using the protocol directly after opening. In the case of failure, the open method will throw an exception that should be caught and handled.

```
[Java]

tserverProtocol.open();
// You can start sending requests here.
```

Opening an Asynchronous Connection

You may prefer to open a connection using asynchronous (non-blocking) methods. This is usually preferred for user-interactive applications, in order to avoid blocking the GUI thread so that the application does not appear "frozen" to the user.

> ### Important
> Be careful when using this method. Connecting asynchronously means that you need to be sure that the `Opened` event is received *before* you send any requests. Otherwise, you might be trying to use a connection that is not yet open.

```
[Java]

tserverProtocol.beginOpen();
// Watch for an Opened event before trying to send or recieve messages.
```

> ### Important
> When using the `BeginOpen()` method, make sure that your code waits for the `Opened` event to fire before attempting to send or receive messages.

Closing a Connection

When you have finished communicating with your servers, you can close the connection. Similar to how a connection is opened, you can also choose to close a connection either synchronrously or asynchronously by using one of the following methods:

```
[Java]

// Synchronous
tserverProtocol.close();
```

Or:

```
[Java]
```

```
// Asyncronous
tserverProtocol.beginClose();
```

You may want to set up a connection to more than one server. To do that, you will need to repeat the steps outlined above for every server you connect to.

Configuring ADDP

The Advanced Disconnection Detection Protocol (ADDP) is a Genesys proprietary add-on to the TCP/IP stack. It implements a periodic poll when no actual activity occurs over a given connection. If a configurable timeout expires without a response from the opposite process, the connection is considered lost.

To enable ADDP, use the configuration options of your Endpoint object. Set the `UseAddp` property to `true` and configure the rest of the properties based on your desired performance. For a description of all ADDP-related options, please refer to the API Reference.

```
[Java]

PropertyConfiguration tserverConfig = new PropertyConfiguration();
tserverConfig.setUseAddp(true);
tserverConfig.setAddpServerTimeout(10);
tserverConfig.setAddpClientTimeout(10);
tserverConfig.setAddpTrace("both");

Endpoint tserverEndpoint = new Endpoint("T-Server", TSERVER_HOST, TSERVER_PORT,
tserverConfig);
TServerProtocol tserverProtocol = new TServerProtocol(tserverEndpoint);
```

Configuring Warm Standby

The WarmStandby Application Block will help you connect or reconnect to your Genesys servers. You will benefit by using the WarmStandby for every application that needs to maintain open connections to Genesys servers, whether you use hot standby or you are only connecting to a single server with no backup redundancy configured.

If you use hot standby, use the WarmStandby Application Block when retrying the connection to your primary or backup server until success, or for reconnecting after both the primary and backup servers are unavailable.

If you are connecting to a single server, use the WarmStandby Application Block to retry the first connection or to reconnect after that server has been unavailable. In this case, configure the `WarmStandbyService` to use the same Endpoint as primary and backup.

Activating the WarmStandby Application Block

To activate the WarmStandby Application Block, you create, configure and start a `WarmStandbyService` object. Two Endpoint objects must be defined: one with parameters for connecting to your primary server and one for connecting to your backup server. You must also remember to start the `WarmStandbyService` before opening the protocol.

```
[Java]

Endpoint tserverEndpoint = new Endpoint("T-Server", TSERVER_HOST, TSERVER_PORT,
tserverConfig);
Endpoint tserverBackupEndpoint = new Endpoint("T-Server", TSERVER_BACKUP_HOST,
TSERVER_BACKUP_PORT, tserverConfig);

TServerProtocol tserverProtocol = new TServerProtocol(tserverEndpoint);

WarmStandbyConfiguration warmStandbyConfig = new WarmStandbyConfiguration(tserverEndpoint,
tserverBackupEndpoint);
warmStandbyConfig.setTimeout(5000);
warmStandbyConfig.setAttempts((short)2);

WarmStandbyService warmStandby = new WarmStandbyService(tserverProtocol);
warmStandby.applyConfiguration(warmStandbyConfig);
warmStandby.start();

tserverProtocol.open();
```

Stopping the WarmStandby Application Block

Stop the `WarmStandbyService` object when your application does not need to maintain the connection with the server any longer. This is typically done at the end of your program.

```
[Java]

warmStandby.stop();
tserverProtocol.close();
```

For more information about how the WarmStandby Application Block works, please refer to the WarmStandby Application Block documentation.

AsyncInvokers

AsyncInvokers are an important aspect of the Platform SDK protocols. They encapsulate the way a piece of code is executed. By using invokers, you can customize what thread executes protocol channel events and handles protocol events. You can also use a thread-pool for parsing protocol messages.

For GUI applications, you normally want most of the logic to happen in the context of the GUI thread. That will enable you to update GUI elements directly, and will simplify your code because you will not have to care about multithreading.

For instance, if you are working with a Swing application, you can use the following `AsyncInvoker` implementation:

```
[Java]

public class SwingInvoker implements AsyncInvoker {

        @Override
        public void invoke(Runnable target) {
```

```
                SwingUtilities.invokeLater(target);
        }

        @Override
        public void dispose() {}

}
```

Assigning a Protocol Invoker

The protocol invoker is in charge of executing channel events (such as channel closed and channel opened) and protocol events (received messages from the server). Usually, when developing a GUI application, you will want to use the GUI thread for handling all kinds of protocol events. By using the AsyncInvoker class described in the section before, you can assign a protocol invoker like this:

```
[Java]

TServerProtocol tserverProtocol = new TServerProtocol(tserverEndpoint);
tserverProtocol.setInvoker(new SwingInvoker());
```

The protocol invoker is of utmost importance for your application. If you do not explicitly set an invoker, then a default internal Platform SDK thread is used, and you will need to use care with possible multithreading issues.

Enabling ADDP on Platform SDK Protocol Client Connections

ADDP is enabled as part of the configuration process for a particular protocol connection instance, and can either be initialized before the connection is open or reconfigured on already opened connection.

The `ConnectionConfiguration` interface describes all connection properties (including details about ADDP, TLS, and the channel character set encoding) for a single instance. For example, if a connection has already configured TLS and later needs to add or change ADDP options then a new `ConnectionConfiguration` should be initialized with the previously set TLS options (along with and other values which are to be preserved) and then have new ADDP options added. The latest configuration applied with overwrite previously set properties.

> **Tip**
> Changing the configuration immediately after a connection is opened, or from the channel event handlers, is not recommended. Some connection configuration options (including ADDP) can be changed on the fly, however the channel configuration is not expected to change often or quickly - options are not treated as if they are dynamic values.

Platform SDK connections have the following ADDP configuration options available:

- set the `protocol` option value to addp to enable ADDP;
- `addp timeout` - specifies how often the client will send ADDP ping requests and wait for responses;

- `addp remote timeout` - specifies how often the server will send ADDP ping requests and wait for responses;
- `addp tracing enable` - this option is used to enable logging of ADDP activities on both the client and server; can be set to "none", "local", "remote", "full" (or "both").

Here is an initialization code sample:

```
protocol = <SomePsdkProtocol>(<Endpoint>);
protocol.set<ProtoSpecificOptions>(<val>);
PropertyConfiguration conf = new PropertyConfiguration();
conf.setOption(AddpInterceptor.PROTOCOL_NAME_KEY, AddpInterceptor.NAME);
conf.setOption(AddpInterceptor.TIMEOUT_KEY, "10");
conf.setOption(AddpInterceptor.REMOTE_TIMEOUT_KEY, "11.5");
conf.setOption(AddpInterceptor.TRACE_KEY, "full");
protocol.configure(conf);
protocol.open();
```

Note that timeout values are stored as strings and parsed as "Float". So, it is ok to have:

```
conf.setOption(AddpInterceptor.TIMEOUT_KEY, "10");
conf.setInteger(AddpInterceptor.TIMEOUT_KEY, 10);    // its the same value
conf.setOption(AddpInterceptor.TIMEOUT_KEY, "11.5"); // = is treated as 11500 ms
```

Also note that in `conf.setOption(AddpInterceptor.TRACE_KEY, "full")`, the `conf.setOption(..)` method accepts string values for the option and it does not know `CfgTraceMode` enumeration - it is defined in configuration protocol which is out of the Platform SDK common libraries.

In release 8.1.0 of Platform SDK for Java, this property handling logic was improved with truncation of the "CFGTM" prefix. Platform SDK for .NET includes this feature starting from release 8.1.1.

So, if you use latest Platform SDK 8.1.0 version for Java, writing `CfgTraceMode.CFGTMBoth.toString()` is acceptable, but earlier versions of Platform SDK or Platform SDK for .NET require that you translate the enumeration values to the corresponding string values.

Possible values can be ["full", "local", "remote", "none"]:

- "local" means logging of ADDP activities locally on client side.
- "remote" for Platform SDK means sending of special initialization bit in ADDP initialization message to server side to ask server to write own ADDP tracing records to server side log.
- "full" means "local" + "remote".

Unknown values are treated as “none”.

Note that comparison is case-insensitive for option values, so "FULL" == "Full" == "full".

.NET

The applications you write with the Platform SDK need to communicate with one or more Genesys servers, so the first thing you need to do is create connections with these servers. You will have to reference libraries and add using statements to your project for each specific protocol you are working with. These steps are not explicitly described here because the files and packages required will vary depending on which protocols you plan to use. Once you have connected to a server, you use that connection to exchange messages with the server. For details about sending and receiving messages to and from a server, refer to the Event Handling article.

Creating a Protocol Object

To connect to a Genesys server, you create an instance of the associated protocol class. As an example, this article will describe a conntection to a Genesys T-Server using the TServerProtocol class. (For different applications, please use this API Reference to check protocol details for the specific server that you wish to connect to.)

In order to create a protocol object, you will first need to create an Endpoint object which acts as a container for generic connection parameters. An Endpoint object contains, at a minimum, a server name, the host name where the server is running, and the port on which the server is listening. The server name will appear in logs but does not affect protocol behavior; it may be any name that is significant to you.

```
[C#]

var tserverEndpoint = new Endpoint("T-Server", TServerHost, TServerPort);
var tserverProtocol = new TServerProtocol(tserverEndpoint);
```

After creating your protocol object, you need to specify some connection parameters that are specific to that protocol. These parameters will differ depending on which server you are connecting to. Please check to the sections specific to the server that you wish to connect to for more information.

Once configuration is complete, you can open the connection to your server.

Opening a Synchronous Connection

The easiest way to open a connection to your server is to do it synchronously, which means that the method will block any additional processing until the server connection has either opened successfully or failed definitively. This is commonly used for non-interactive, batch applications. In this case, you can add code for using the protocol directly after opening. In the case of failure, the open method will throw an exception that should be caught and handled.

```
[C#]

tserverProtocol.Open();
// You can start sending requests here.
```

Opening an Asynchronous Connection

You may prefer to open a connection using asynchronous (non-blocking) methods. This is usually preferred for user-interactive applications, in order to avoid blocking the GUI thread so that the application does not appear "frozen" to the user.

Important: Be careful when using this method. Connecting asynchronously means that you need to be sure that the Opened event is received *before* you send any requests. Otherwise, you might be trying to use a connection that is not yet open.

```
[C#]

tserverProtocol.BeginOpen();
// Watch for an Opened event before trying to send or recieve messages.
```

> **Important**
>
> When using the `BeginOpen()` method, make sure that your code waits for the Opened event to fire before attempting to send or receive messages.

Closing a Connection

When you have finished communicating with your servers, you can close the connection. Similar to how a connection is opened, you can also choose to close a connection either synchronously or asynchronously by using one of the following methods:

```
[C#]

tserverProtocol.Close();
```

Or:

```
[C#]

tserverProtocol.BeginClose();
```

Or:

```
[C#]

tserverProtocol.Dispose();
```

You may want to set up a connection to more than one server. To do that, you will need to repeat the steps outlined above for every server.

Configuring ADDP

The Advanced Disconnection Detection Protocol (ADDP) is a Genesys proprietary add-on to the TCP/IP stack. It implements a periodic poll when no actual activity occurs over a given connection. If a configurable timeout expires without a response from the opposite process, the connection is considered lost.

To enable ADDP, use the configuration options of your Endpoint object. Set the `UseAddp` property to true and configure the rest of the properties based on your desired performance. For a description of all ADDP-related options, please refer to the API Reference.

```
[C#]

var tserverConfig = new PropertyConfiguration();
```

```
tserverConfig.UseAddp = true;
tserverConfig.AddpServerTimeout = 10;
tserverConfig.AddpClientTimeout = 10;
tserverConfig.AddpTrace = "both";

var tserverEndpoint = new Endpoint("T-Server", TServerHost, TServerPort, tserverConfig);
var tserverProtocol = new TServerProtocol(tserverEndpoint);
```

Configuring Warm Standby

The Warm Standby Application Block will help you connect or reconnect to your Genesys servers. You will benefit by using the Warm Standby for every application that needs to maintain open connections to Genesys servers, whether you use hot standby or you are only connecting to a single server with no backup redundancy configured.

If you use hot standby, use the Warm Standby Application Block when retrying the connection to your primary or backup server until success, or for reconnecting after both the primary and backup servers are unavailable.

If you are connecting to a single server, use the Warm Standby Application Block to retry the first connection or to reconnect after that server has been unavailable. In this case, configure the `WarmStandbyService` to use the same `Endpoint` as primary and backup.

Activating the WarmStandby Application Block

To activate the Warm Standby Application Block, you create, configure and start a `WarmStandbyService` object. Two `Endpoint` objects must be defined: one with parameters for connecting to your primary server and one for connecting to your backup server. You must also remember to start the `WarmStandbyService` before opening the protocol.

```
[C#]

var tserverEndpoint = new Endpoint("T-Server", TServerHost, TServerPort, tserverConfig);
var tserverBackupEndpoint = new Endpoint("T-Server", TServerBackupHost, TServerBackupPort,
tserverConfig);

var tserverProtocol = new TServerProtocol(tserverEndpoint);

var warmStandbyConfig = new WarmStandbyConfiguration(tserverEndpoint, tserverBackupEndpoint);
warmStandbyConfig.Timeout = 5000;
warmStandbyConfig.Attempts = 2;

var warmStandby = new WarmStandbyService(tserverProtocol);
warmStandby.ApplyConfiguration(warmStandbyConfig);
warmStandby.Start();

tserverProtocol.Open();
```

Stopping the WarmStandby Application Block

Stop the `WarmStandbyService` object when your application does not need to maintain the connection with the server any longer. This is typically done at the end of your program.

```
[C#]
```

```
warmStandby.Stop();
tserverProtocol.Dispose();
```

For more information about how the Warm Standby Application Block works, please refer to the Warm Standby Application Block documentation.

AsyncInvokers

AsyncInvokers are an important aspect of the Platform SDK protocols. They encapsulate the way a piece of code is executed. By using invokers, you can customize what thread executes protocol channel events and handles protocol events. You can also use a thread-pool for parsing protocol messages.

For GUI applications, you normally want most of the logic to happen in the context of the GUI thread. That will enable you to update GUI elements directly, and will simplify your code because you will not have to care about multi-threading.

For instance, if you are working with a Windows Forms or WPF application,, you can use the following `IAsyncInvoker` implementation:

```
[C#]

public class SyncContextInvoker : IAsyncInvoker
{
	private readonly SynchronizationContext syncContext;

	public SyncContextInvoker()
	{
		this.syncContext = SynchronizationContext.Current;
	}

	public void Invoke(Delegate d, params object[] args)
	{
		syncContext.Post(s => { d.DynamicInvoke(args); }, null);
	}

	public void Invoke(WaitCallback callback, object state)
	{
		syncContext.Post(s => { callback(state); }, null);
	}

	public void Invoke(EventHandler handler, object sender, EventArgs args)
	{
		syncContext.Post(s => { handler(sender, args); }, null);
	}
}
```

The Protocol Invoker

The protocol invoker is in charge of executing channel events (such as channel closed and channel opened) and protocol events (received messages from the server). Usually, when developing a GUI application, you will want to use the GUI thread for handling all kinds of protocol events. By using the class implemented in the section before, you can assign a protocol invoker like this:

[C#]

```
var tserverProtocol = new TServerProtocol(tserverEndpoint);
tserverProtocol.Invoker = new SyncContextInvoker();
```

The protocol invoker is of utmost importance for your application. If you do not explicitly set an invoker, then a default internal Platform SDK thread is used, and you will need to use care with multi-threading issues.

Advanced: Multithreaded Message Parsing

> **Tip**
> Please apply this section only if your application is suffering from performance problems because of large message parsing. You should identify the bottleneck using profiling techniques, and should measure the effect after making these changes by using the same profiling techniques.

Take into account that the technique described here can affect the correctness of your application, since concurrently parsing messages can affect the order in which those messages are received. So use this technique only selectively and in places where order of received messages is not relevant.

Every message you receive from a Genesys server is formatted in some way. Most Genesys servers use binary protocols, while some use XML-based protocols. When your application receives one of these messages, it parses the message and places it in the message queue for the appropriate protocol.

By default, the Platform SDK uses a single thread for parsing all messages. This parsing can be time-consuming in certain cases, and some applications may face performance issues. For example, some applications may receive lots of large binary-format messages, such as some of the statistics messages generated by Stat Server, while others might need to parse messages in non-binary formats, such as the XML format used to communicate with Genesys Multimedia (or e-Services) servers.

If message parsing becomes a bottleneck for your application, you can try to enable multi-threaded message parsing. This is done by setting the protocol connection invoker to an invoker that dispatches work to a pool of threads. One such invoker is provided out-of-the-box:

[C#]

```
statServerProtocol.SetConnectionInvoker(DefaultInvoker.InvokerSingleton);
```

Configuring Platform SDK Channel Encoding for String Values

While sending string attributes/values to a server (or to the other side of any connection), Platform SDK packs strings to their binary representation. The binary representation depends on actual charset encoding, so it is important that this data will be unpacked with correct encoding when received on the other side of the connection.

Genesys protocols do not allow client and server sides to synchronize (that is, exchange) the encoding being used<ref>**Exception:** A Configuration Server 8.1.2+ deployment that is configured as UTF-8 multi-lingual can automatically synchronize UTF-8 encoding with Platform SDK 8.1.3 or later. For details, see Connecting Using UTF-8 Character Encoding.</ref>, so application developers may need to handle this configuration manually. The most common situation requiring this type of configuration occurs when a Genesys server and the application using Platform SDK to connect with that server have different localization settings, causing default encoding to be different on both sides.

There are two possible solutions for synchronizing the client side encoding with that of the server side:

1. (Java only) Change default jvm encoding with the jvm argument: `java -Dfile.encoding=...`
   This changes the charset encoding for the entire jvm, so will affect the main application and any Platform SDK connections to other servers. It may affect the client application relation with other components on the client host.
2. (Java only) Starting with Platform SDK 8.1.3, the new `com.genesyslab.platform.defaultcharset` system property can be used to set default charset encoding for Platform SDK connections without the need to change default encoding for whole jvm.
   Platform SDK checks this property once before opening the first connection, and if a value is specified then it will be used as the default encoding for all Platform SDK connections (instead of the value defined for the jvm).
3. Configure a particular Platform SDK connection to use the server side encoding with following connection configuration option (added in Platform SDK 8.0.1):

```
[Java]

protocol = ...;

PropertyConfiguration conf = new PropertyConfiguration();
conf.setOption(Connection.STR_ATTR_ENCODING_NAME_KEY, "windows-1252");

protocol.configure(conf);
protocol.open();

[C#]

protocol = ...;

PropertyConfiguration conf = new PropertyConfiguration();
conf.SetOption(ConnectionBase.StringAttributeEncodingKey, "windows-1252");

protocol.Configure(conf);
protocol.Open();
```

<references/>

Using the Warm Standby Application Block

The Warm Standby Application Block is a reusable production-quality component that enables developers to switch to a backup server in case their primary server fails, without needing to guarantee the integrity of existing interactions. It has been designed using industry best practices and provided with source code so it can be used "as is," extended, or tailored if you need to. Please see the License Agreement for details.

This article examines the architecture and design of the Warm Standby Application Block, as well as giving details about how to setup the QuickStart application that ships with this application block.

Java

Architecture and Design

Many contact center environments require redundant backup servers that are able to take over quickly if a primary server fails. In this situation, the primary server operates in active mode, accepting connections and exchanging messages with clients. The backup server, on the other hand, is in standby mode. If the primary server fails, the backup server switches to active mode, assuming the role and behavior of the primary server.

There are two standby modes: *warm standby* and *hot standby*. The main difference between them is that warm standby mode does not ensure the continuation of interactions in progress when a failure occurs, while hot standby mode does.

The Client Channel Architecture

Since the Warm Standby Application Block is designed to be used in the context of a Client Channel architecture, it is important to understand that architecture before talking about the application block itself.

To start with, this architecture consists of three functional components:

- A connection
- A client channel
- A protocol channel

These components are shown in the following figure.

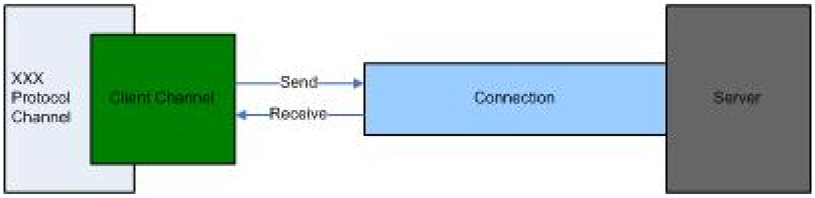

The *connection* controls all necessary TCP/IP connection activities, while the client channel contains the protocol- and server-independent channel functionality that is common for a protocol channel. Finally, the *protocol channel* controls all of the client channel activities that are dependent on the protocol and the server.

Client Channel State

The state of a client channel is based on the state of the corresponding connection. There are four major states:

- Opening (Registration)
- Opened
- Closing
- Closed

The figure below shows a detailed client channel state diagram.

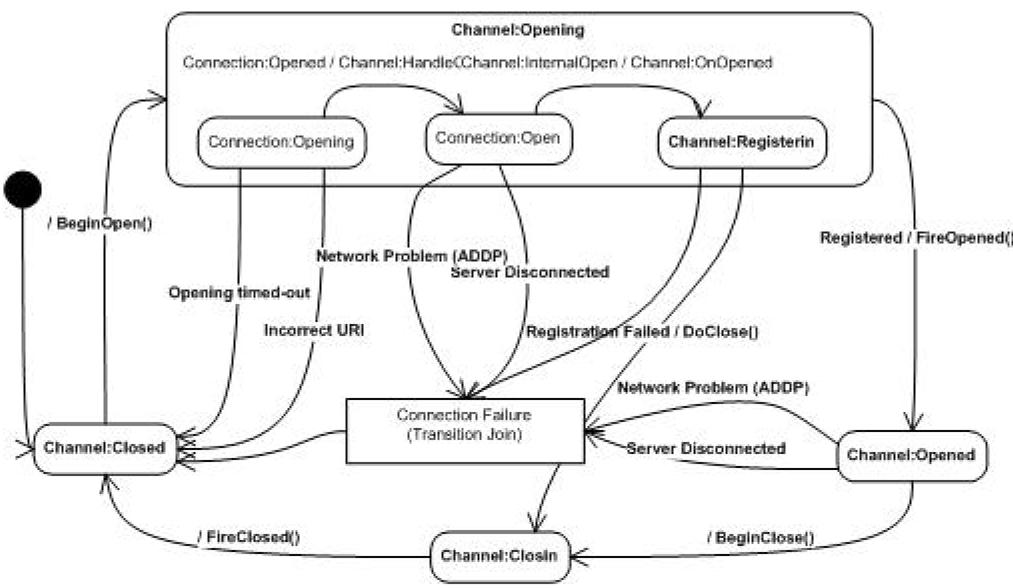

In addition to establishing a TCP/IP connection, several activities may take place when a client channel opens. These activities can include things like:

- A preliminary exchange of messages with the server, which is known as registration
- Reading the client channel's locally stored configuration information

You can often determine the cause of a client channel failure by checking the state of the client channel just before it closed. There are exceptions to this rule, however, such as a registration failure, which is protocol-specific.

Client Channel Failure Scenarios

There are several common client channel failure scenarios:

Client Channel Failure Scenarios

Scenario	Description	Source States	Condition	Target State	Protocol-Dependent
Opening Timed Out	Channel tries to open connection to non-existing URI	Opening	Connection opening timeout	Closed	No
Wrong URI	Channel tries to open connection to non-existing URI	Opening	Incorrect URI exception	Closed	No
Connection Problem	Channel connection detects a connection problem	Opened Opening	Server disconnected	Closed	No
Network Problem (ADDP)	Channel connection detects a network problem (ADDP)	Opened Opening	Network problem (ADDP)	Closed	No
Wrong Server or Protocol	Channel tries to open connection with an incorrect server or protocol	Opening	Registration Failed/ ProtocolException	Closing	Yes
Registration Failure	One of the channel registration steps failed	Opening	Registration Failed/ ProtocolException	Closing	Yes

Note that the first four scenarios, *Opening timed-out*, *Wrong URI*, *Connection Problem*, and *Network Problem* happen with the connection (TCP/IP) component. They do not involve protocol- or server-specific elements, whether in terms of failure-specific data or in terms of channel recovery actions and data.

The *Wrong Server* or *Protocol* and *Registration Failure* scenarios are protocol- or server-dependent and can be different for each type of protocol channel.

Application Block Architecture

The Warm Standby Application Block’s functionality is based on intercepting the channel's transition from a non-closed state to the Closed state. As you can see in the following figure, the application block is able to pick up this information because it sits between the client and protocol channels.

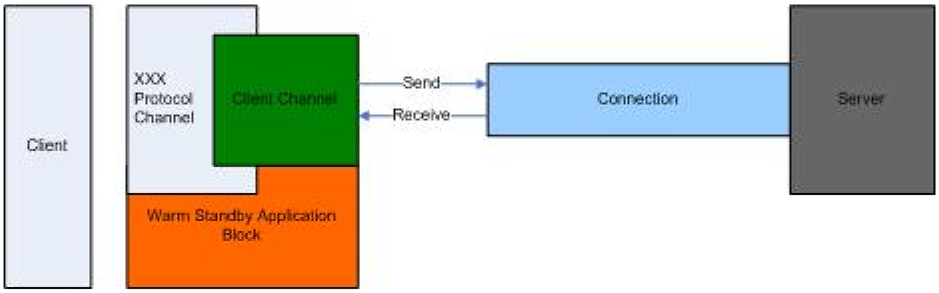

Upon receiving the channel’s `Closed` event, the application block uses diagnostic information to determine why the channel has closed. This diagnostic information is necessary to determine what actions, if any, the application block should take to restore the channel’s connectivity to the server.

The Warm Standby Application Block can take several different steps to recover channel connectivity. These steps are:

- Do nothing (close the channel by request of the user application)
- Attempt to open the channel without switching over its connectivity configuration from primary to backup
- Attempt to open the channel, switching its connectivity configuration from primary to backup
- Deactivate, in case of a fatal failure

Any application block activity will be followed by a corresponding event generated by the application block. These events will provide user applications with the opportunity to monitor and react to all of the application block’s activities and failures

To control channel connectivity with a warm standby mechanism, the user application should activate the Warm Standby Application Block instance that is responsible for handling the particular channel's connectivity failure and recovery.

Warm Standby Application Block Algorithm

The Warm Standby Application Block has 4 states, as shown below.

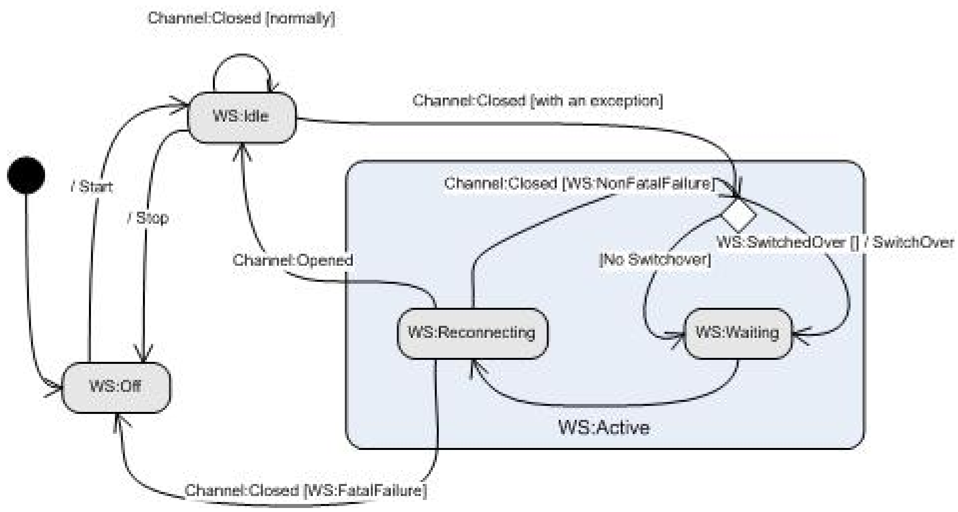

As soon as a channel's Warm Standby Application Block is activated, it goes into the idle state, waiting for the channel's Closed event. When the channel issues a `Closed` event, the application block checks to see if the channel was closed due to a connectivity failure. If so, the application block instance starts the channel connectivity recovery procedure, as shown below.

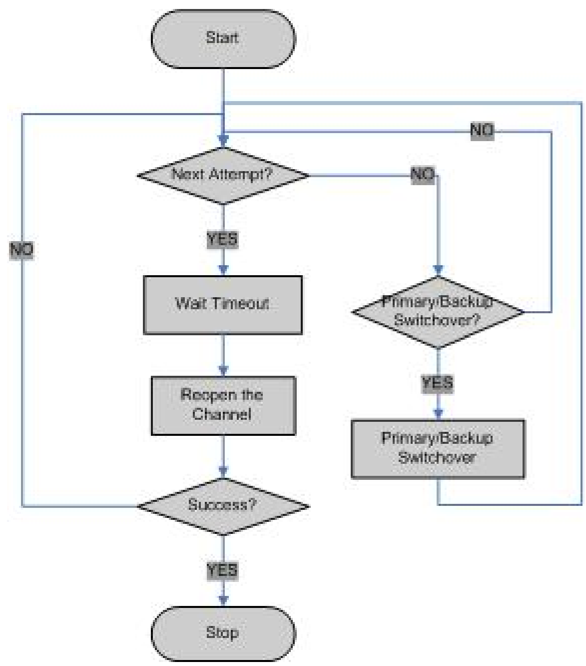

Here is the procedure for the Warm Standby Application Block:

- The user should activate the Warm Standby Application Block for every channel he or she intends to work with.
- In the active state, the application block waits for the channel's `Closed` event.
- On receiving the channel's `Closed` event, the application block activates the channel connectivity recovery procedure.

Application Block Components

The Warm Standby Application Block distribution consists of two main components:

1. The application block itself, which provides an interface that you can use to integrate it into different GUI applications.
2. A sample application, the *WarmStandbyQuickStart* application, which is built on the Warm Standby Application Block

As shown below, the application block itself runs on top of the Platform SDK, while the QuickStart application runs on top of the application block.

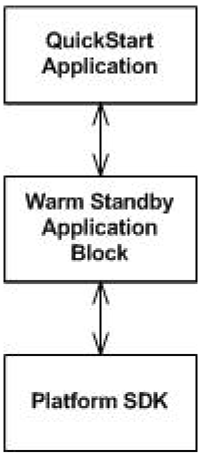

The Warm Standby Application Block Interface

The Warm Standby Application Block consists of the following classes:

- `WarmStandbyService`
- `WarmStandbyConfiguration`

The `WarmStandbyService` class monitors and controls the connectivity of the channel it is responsible for, while the `WarmStandbyConfiguration` class handles all the parameters that are needed for the proper functioning of the warm standby process.

Starting with release 8.1.1, default behavior for the `WarmStandbyService` connection restoration includes the following improvements to provide improved performance:

- Following a switchover or the first reconnection attempt, `WarmStandbyService` no longer waits for a timeout to occur.
- Check backup server availability by performing a fast first switchover.

User applications can subscribe to the controlled channel’s `Closed` and `Opened` events in order to monitor and handle channel connectivity.

`WarmStandbyService`'s `StateChanged` event is fired on any change of state in `WarmStandby`, providing the means for a user application to monitor state changes and to control the application block's activities.

Using the Warm Standby Application Block

Installing the Warm Standby Application Block

Before you install the Warm Standby Application Block, it is important to review the software requirements and the structure of the software distribution.

Software Requirements

To work with the Warm Standby Application Block, you must ensure that your system meets the software requirements established in the Genesys Supported Operating Environment Reference Manual, as well as meeting the following minimum software requirements:

- JDK 1.6 or higher

Building the Warm Standby Application Block

To build the Warm Standby Application Block:

1. Open the *<Platform SDK Folder>\applicationblocks\warmstandby* folder.
2. Run either *build.bat* or *build.sh*, depending on your platform.

You may need to edit the path specified in the quickstart file by adding quotation marks if your ANT_HOME environment variable contains spaces.

This build file will create the *warmstandbyappblock.jar* file, located within the *<Platform SDK Folder>\applicationblocks\warmstandby\dist\lib* directory.

Now you are ready to add the appropriate import statements to your source code and start using the Warm Standby Application Block:

```
[Java]
import com.genesyslab.platform.applicationblocks.warmstandby.*;
```

Using the QuickStart Application

The easiest way to start using the Warm Standby Application Block is to use the bundled QuickStart application. This application ships in the same folder as the application block.

To run the QuickStart application:

1. Open the *\ApplicationBlocks\WarmStandby\quickstart* folder.
2. Run either *quickstart.bat* or *quickstart.sh*, depending on your platform.

You may need to edit the path specified in the quickstart file by adding quotation marks if your

ANT_HOME environment variable contains spaces.

After you start the application, you will see the user interface shown below.

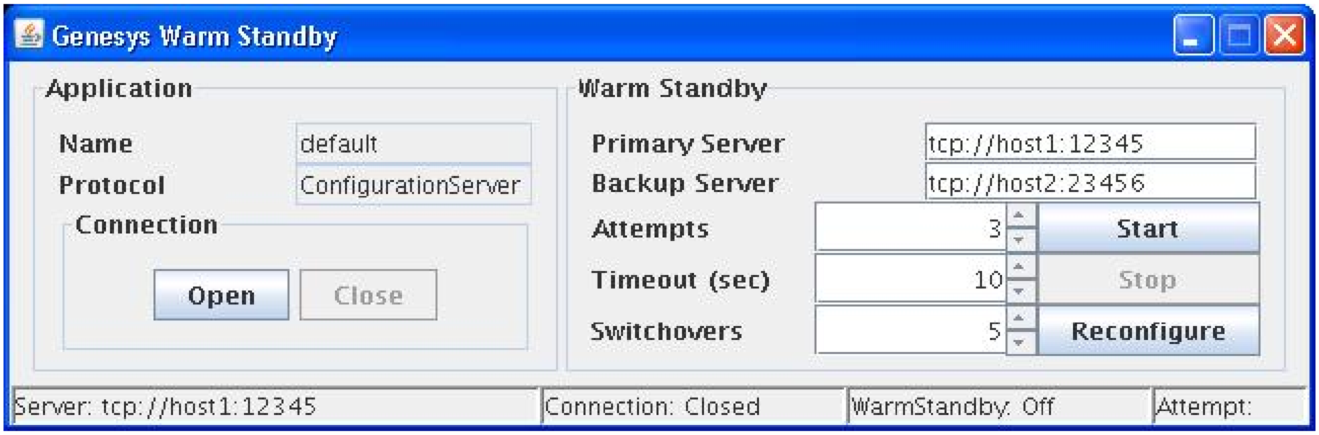

On startup, the QuickStart application uses values specified by the *quickstart.properties* configuration file. You can change these values either by editing that file or by overwriting them after running the user interface.

This form has two main sections. The left side enables you to set up a connection for the application indicated in the Name field, using the protocol specified in the Protocol field. To open the connection, press the *Open* button. Press the *Close* button to close it.

The right side of the form lets you specify primary and backup servers. It also lets you specify the number of times the warm standby mechanism will try to contact the primary server, and what the timeout value should be for each attempt.

Once you have the desired values, you can press the *Start* button to turn on the warm standby feature. If you would like to change the configuration after warm standby is turned on, simply modify the configuration information and press the Reconfigure button. The warm standby configuration will be changed dynamically.

.NET

Architecture and Design

Many contact center environments require redundant backup servers that are able to take over quickly if a primary server fails. In this situation, the primary server operates in active mode, accepting connections and exchanging messages with clients. The backup server, on the other hand, is in standby mode. If the primary server fails, the backup server switches to active mode, assuming the role and behavior of the primary server.

There are two standby modes: *warm standby* and *hot standby*. The main difference between them is that warm standby mode does not ensure the continuation of interactions in progress when a failure occurs, while hot standby mode does.

The Client Channel Architecture

Since the Warm Standby Application Block is designed to be used in the context of a Client Channel architecture, it is important to understand that architecture before talking about the application block itself.

To start with, this architecture consists of three functional components:

- A connection
- A client channel
- A protocol channel

These components are shown in the following figure.

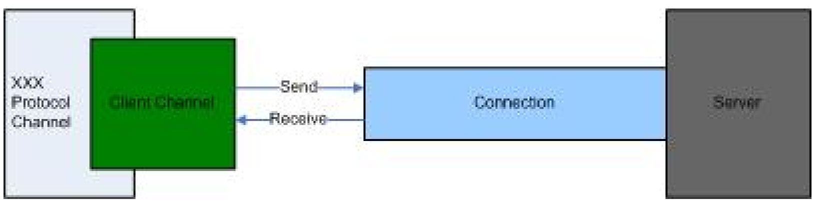

The *connection* controls all necessary TCP/IP connection activities, while the client channel contains the protocol- and server-independent channel functionality that is common for a protocol channel. Finally, the *protocol channel* controls all of the client channel activities that are dependent on the protocol and the server.

Client Channel State

The state of a client channel is based on the state of the corresponding connection. There are four major states:

- Opening (Registration)
- Opened
- Closing
- Closed

The figure below shows a detailed client channel state diagram.

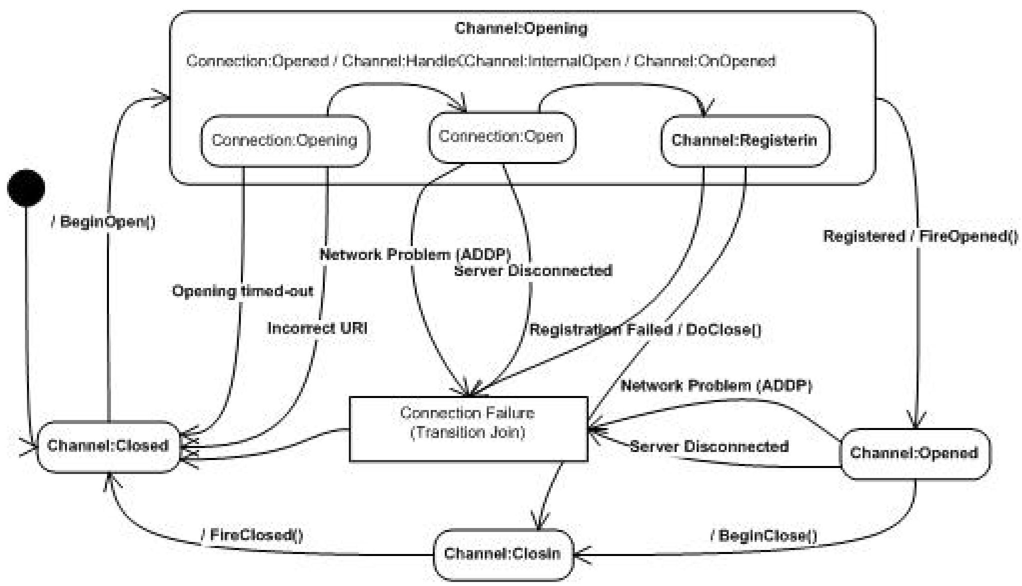

In addition to establishing a TCP/IP connection, several activities may take place when a client channel opens. These activities can include things like:

- A preliminary exchange of messages with the server, which is known as registration
- Reading the client channel's locally stored configuration information

You can often determine the cause of a client channel failure by checking the state of the client channel just before it closed. There are exceptions to this rule, however, such as a registration failure, which is protocol-specific.

Client Channel Failure Scenarios

There are several common client channel failure scenarios:

Client Channel Failure Scenarios

Scenario	Description	Source States	Condition	Target State	Protocol-Dependent
Opening Timed Out	Channel tries to open connection to non-existing URI	Opening	Connection opening timeout	Closed	No
Wrong URI	Channel tries to open connection to non-existing URI	Opening	Incorrect URI exception	Closed	No
Connection Problem	Channel connection detects a connection	Opened Opening	Server disconnected	Closed	No

Scenario	Description	Source States	Condition	Target State	Protocol-Dependent
	problem				
Network Problem (ADDP)	Channel connection detects a network problem (ADDP)	Opened Opening	Network problem (ADDP)	Closed	No
Wrong Server or Protocol	Channel tries to open connection with an incorrect server or protocol	Opening	Registration Failed/ ProtocolException	Closing	Yes
Registration Failure	One of the channel registration steps failed	Opening	Registration Failed/ ProtocolException	Closing	Yes

Note that the first four scenarios, *Opening timed-out, Wrong URI, Connection Problem,* and *Network Problem* happen with the connection (TCP/IP) component. They do not involve protocol- or server-specific elements, whether in terms of failure-specific data or in terms of channel recovery actions and data.

The *Wrong Server* or *Protocol* and *Registration Failure* scenarios are protocol- or server-dependent and can be different for each type of protocol channel.

Application Block Architecture

The Warm Standby Application Block’s functionality is based on intercepting the channel's transition from a non-closed state to the Closed state. As you can see in the following figure, the application block is able to pick up this information because it sits between the client and protocol channels.

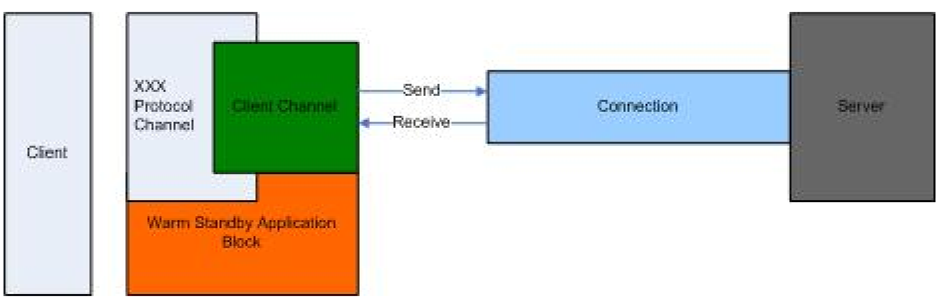

Upon receiving the channel’s `Closed` event, the application block uses diagnostic information to determine why the channel has closed. This diagnostic information is necessary to determine what actions, if any, the application block should take to restore the channel’s connectivity to the server.

The Warm Standby Application Block can take several different steps to recover channel connectivity. These steps are:

- Do nothing (close the channel by request of the user application)
- Attempt to open the channel without switching over its connectivity configuration from primary to backup
- Attempt to open the channel, switching its connectivity configuration from primary to backup
- Deactivate, in case of a fatal failure

Any application block activity will be followed by a corresponding event generated by the application block. These events will provide user applications with the opportunity to monitor and react to all of the application block’s activities and failures

To control channel connectivity with a warm standby mechanism, the user application should activate the Warm Standby Application Block instance that is responsible for handling the particular channel's connectivity failure and recovery.

Warm Standby Application Block Algorithm

The Warm Standby Application Block has 4 states, as shown below.

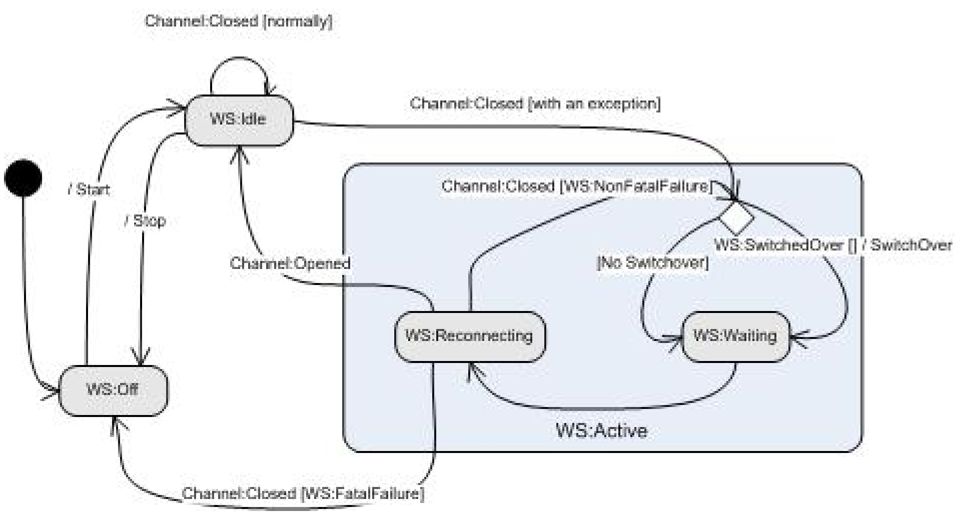

As soon as a channel’s Warm Standby Application Block is activated, it goes into the idle state, waiting for the channel’s Closed event. When the channel issues a `Closed` event, the application block checks to see if the channel was closed due to a connectivity failure. If so, the application block instance starts the channel connectivity recovery procedure, as shown below.

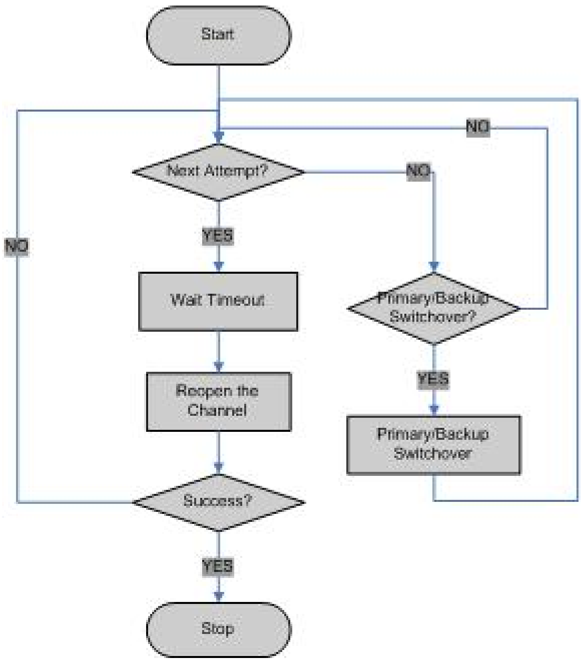

Here is the procedure for the Warm Standby Application Block:

- The user should activate the Warm Standby Application Block for every channel he or she intends to work with.
- In the active state, the application block waits for the channel's `Closed` event.
- On receiving the channel's `Closed` event, the application block activates the channel connectivity recovery procedure.

Application Block Components

The Warm Standby Application Block distribution consists of two main components:

1. The application block itself, which provides an interface that you can use to integrate it into different GUI applications.
2. A sample application, the *WarmStandbyQuickStart* application, which is built on the Warm Standby Application Block

As shown below, the application block itself runs on top of the Platform SDK, while the QuickStart application runs on top of the application block.

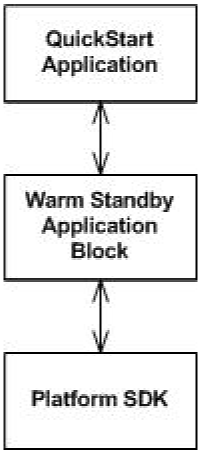

The Warm Standby Application Block Interface

The Warm Standby Application Block consists of the following classes:

- `WarmStandbyService`
- `WarmStandbyConfiguration`

These classes are shown in greater detail below.

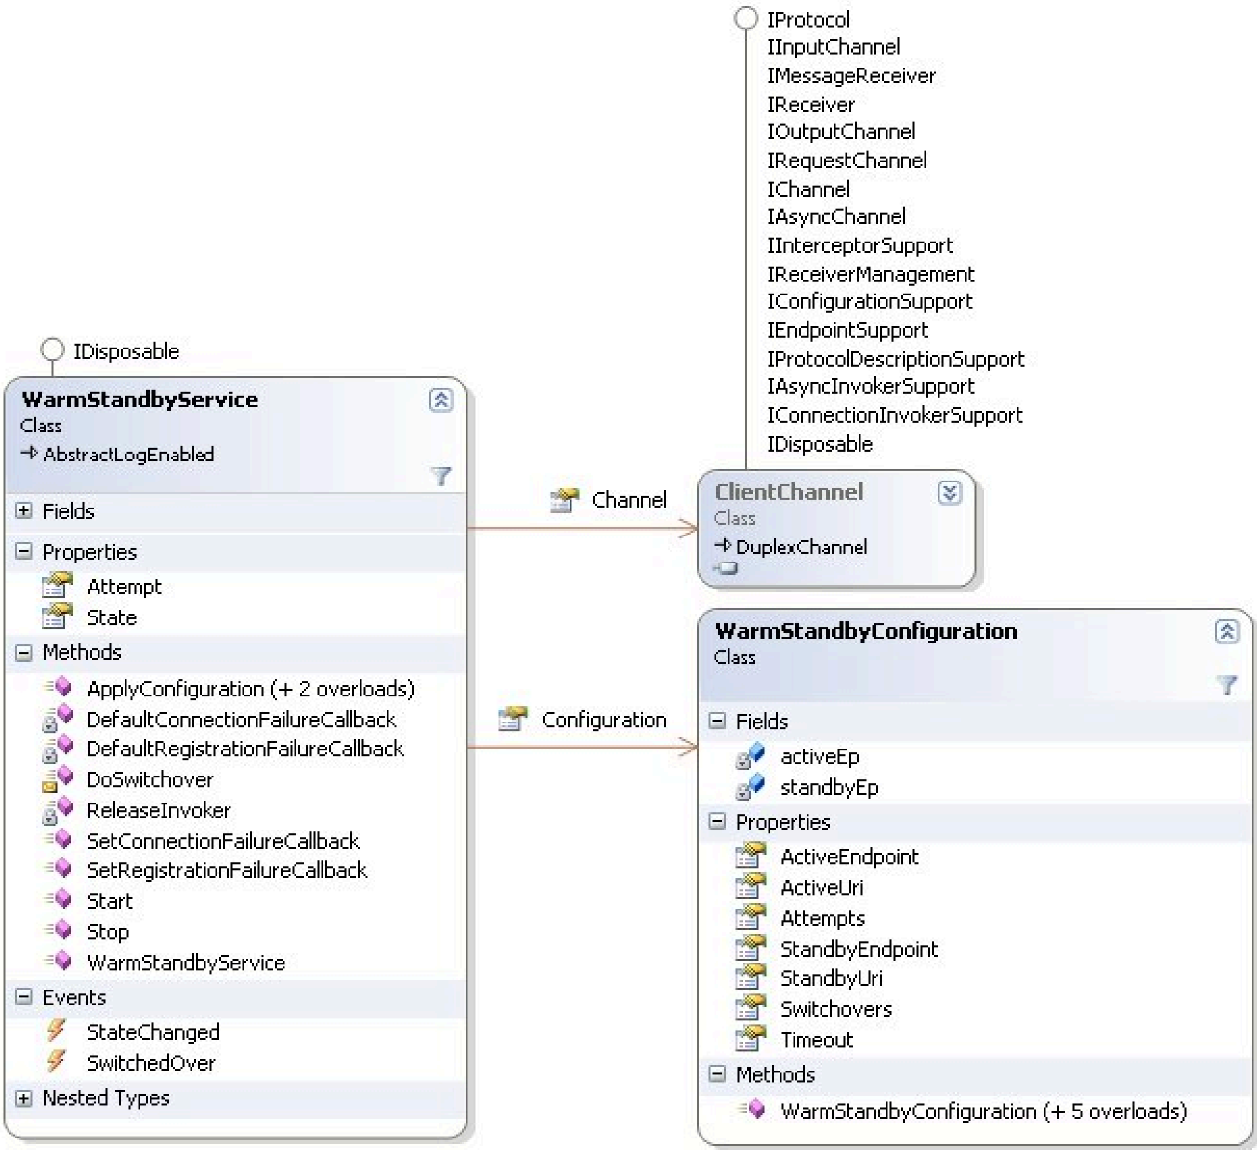

The `WarmStandbyService` class monitors and controls the connectivity of the channel it is responsible for, while the `WarmStandbyConfiguration` class handles all the parameters that are needed for the proper functioning of the warm standby process.

Starting with release 8.1.1, default behavior for the `WarmStandbyService` connection restoration includes the following improvements to provide improved performance:

- Following a switchover or the first reconnection attempt, `WarmStandbyService` no longer waits for a timeout to occur.
- Check backup server availability by performing a fast first switchover.

User applications can subscribe to the controlled channel’s `Closed` and `Opened` events in order to monitor and handle channel connectivity.

WarmStandbyService's StateChanged event is fired on any change of state in WarmStandby, providing the means for a user application to monitor state changes and to control the application block's activities.

Using the Warm Standby Application Block

Installing the Warm Standby Application Block

Before you install the Warm Standby Application Block, it is important to review the software requirements and the structure of the software distribution.

Software Requirements

To work with the Warm Standby Application Block, you must ensure that your system meets the software requirements established in the Genesys Supported Operating Environment Reference Manual.

Configuring the Warm Standby Application Block

In order to use the QuickStart application, you need to set up the XML configuration file that comes with the application block. This file is located at *Quickstart\app.config*. This is what the contents look like:

```
<?xml version="1.0" encoding="utf-8" ?>
<configuration>
        <configSections>

        </configSections>
        <WarmStandbyQuickStart>
                <Channel
                        ClientType="19"
                        ProtocolName="ConfigurationServer"
                        ClientName="default"
                />
                <WarmStandby
                        PrimaryServer="tcp://hostname:9999"
                        BackupServer="tcp://hostname:9999"
                        Attempts="3"
                        Timeout="10"
                        Switchovers="3"
                />
                <ConfServer
                        UserName="default"
                        UserPassword="password"
                />
        </WarmStandbyQuickStart>

</configuration>
```

Follow the instructions in the comments and save the file.

Building the Warm Standby Application Block

The Platform SDK distribution includes a *Genesyslab.Platform.ApplicationBlocks.WarmStandby.dll* file that you can use as is. This file is located in the bin directory at the root level of the Platform SDK directory. To build your own copy of this application block, follow the instructions below:

To build the Warm Standby Application Block:

1. Open the *<Platform SDK Folder>\ApplicationBlocks\WarmStandby* folder.
2. Double-click *WarmStandby.sln*.
3. Build the solution.

Using the QuickStart Application

The easiest way to start using the Warm Standby Application Block is to use the bundled QuickStart application. This application ships in the same folder as the application block.

To run the QuickStart application:

1. Open the *<Platform SDK Folder>\ApplicationBlocks\WarmStandby* folder.
2. Double-click *WarmStandbyQuickStart.sln*.
3. Build the solution.
4. Find the executable for the QuickStart application, which will be at *<Platform SDK Folder>\ApplicationBlocks\WarmStandby\QuickStart\bin\Debug\WarmStandbyQuickStart.exe*.
5. Double-click *WarmStandbyQuickStart.exe*.

After you start the application, you will see the user interface shown below.

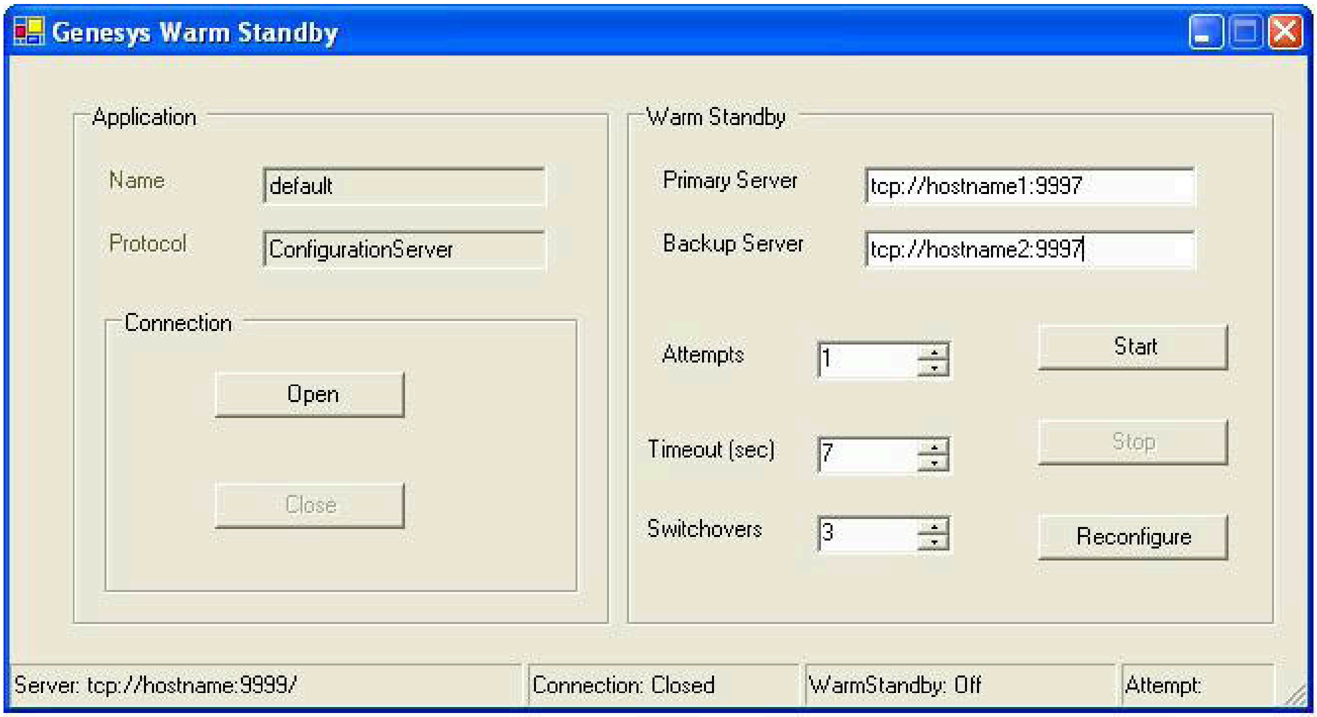

This form has two main sections. The left side enables you to set up a connection for the application indicated in the *Name* field, using the protocol specified in the *Protocol* field. To open the connection, press the *Open* button. Press the *Close* button to close it.

The right side of the form lets you specify primary and backup servers. It also lets you specify the number of times the warm standby mechanism will try to contact the primary server, and what the timeout value should be for each attempt. On startup, these values are picked up from the configuration file, but you can change them in the user interface.

Once you have the desired values, you can press the *Start* button to turn on the warm standby feature. If you would like to change the configuration after warm standby is turned on, simply modify the configuration information and press the Reconfigure button. The warm standby configuration will be changed dynamically.

Event Handling

Java

Once you have connected to a server, much of the work for your application will involves sending messages to that server and handling the events you receive from the server. This article describes how to send and receive messages from a server.

Messages: Overview of Events and Requests

Messages you send to a server are called requests, while messages you receive are called events. An event that is received from a server as the result of executing a request is called a response. In summary, messages can be classified by using the following taxonomy:

- Requests: sent to the server
- Events: received from the server
    - Responses: received as the result of a request
    - Unsolicited events: not a direct result of a request

> **Tip**
> On this page, we will use the more general term "message" instead of "event", in order to avoid confusion between protocol events and programming events.

For example, you may send a request to log in an agent or to gather statistics. You might also send a request to update a configuration object, or to shut down an application.

In each of these cases, the server will respond with an event message, as shown below.

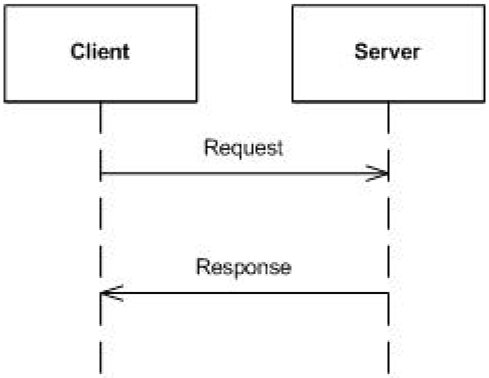

You may also get unsolicited events from the server. That means receiving events that are not a response to a specific request. For example, `EventRinging` will notify you of a call ringing on an extension that you are currently monitoring.

Receiving Messages

With the Platform SDK, you can receive messages synchronously or asynchronously. It is important that you define the way your application will work in this aspect. In general, you will probably use only one type or the other in the same application.

Interactive applications normally use asynchronous message handling, because that will prevent the UI thread from being blocked, which could make the application appear "frozen" to a user. On the other hand, non-interactive batch applications commonly use synchronous response handling, as that allows writing easy code that performs step-by-step.

Receiving Messages Asynchronously

Most Platform SDK applications need to handle unsolicited events. This is particularly true for applications that monitor the status of contact center resources, such as extensions.

You receive server messages by implementing a `MessageHandler` that contains the event-handling logic:

```
[Java]

MessageHandler tserverMessageHandler = new MessageHandler() {
        @Override
        public void onMessage(Message message) {
                // your event-handling code goes here
        }
};
```

Then you set your implementation as the protocol `MessageHandler`.

```
[Java]

tserverProtocol.setMessageHandler(tserverMessageHandler);
```

Important

You need to know that your event-handling logic will be executed by using the protocol invoker. Please set the invoker appropriate for your application needs. For more information about the protocol invoker and how to set it, refer to Connecting to a Server.

Inside your event-handling code, you will want to execute different logic for different kinds of events. A typical way to do this is using a `switch` statement, based on the event identifier:

```
[Java]

switch (message.messageId()) {
        case EventAgentLogin.ID:
                OnEventAgentLogin(message);
                break;
        case EventAgentLogout.ID:
                OnEventAgentLogout(message);
                break;
}
```

Receiving Messages Synchronously

Some kinds of applications, such as batch applications, benefit from receiving messages synchronously. This means that received messages will queue up and be handled by the application on demand.

In order to receive messages this way, you simply **do not** set a protocol `MessageHandler` as described in the previous section.

Tip

For releases prior to Platform SDK 8.1.1, messages were received synchronously by default. Please note that 8.1.1 behavior is backwards-compatible, and pre-8.1.1 applications will continue to work as expected without any modification.

To receive a message synchronously, use the `Receive` method. This method blocks processing, waiting for the next message to be received before continuing. Take into account that the maximum time to wait is set by a configurable timeout value. If the timeout expires and no event is received, you will receive a `null` value.

```
[Java]

Message message = tserverProtocol.receive();
```

If you want to set your own timeout, you can use the `Receive` method overload that takes a timeout parameter. Otherwise, if you use `Receive` with no parameters, the protocol Timeout property will be used.

Sending Requests Asynchronously

This is the easiest way to send a message to a server. Suppose you have created and filled a request object, for example, a `RequestAgentLogin` message for Interaction Server:

```
[Java]

RequestAgentLogin loginRequest = RequestAgentLogin.create();
loginRequest.setTenantId(tenantId);
loginRequest.setAgentId(agentId);
loginRequest.setPlaceId(placeId);
```

You can then send it to the server using the following code:

```
[Java]

interactionServerProtocol.send(loginRequest);
```

This will result in your application receiving a response from the Interaction Server: either an `EventAck` or an `EventError` message. By using the Send method, you will ignore that response at the place where you make the request. You will get the response, like any other unsolicited event, using the techniques described in the *Receiving Messages* section.

Handling Responses

The understanding of how to send requests and receive events is all you need to communicate with Genesys servers. However, the Platform SDK also provides the ability to easily associate a response with the particular request that originated it.

Receiving a Response Synchronously

The easiest way to handle responses is with the `Request` method. This is a blocking method, as your application stops to wait for a response to come from the server. Using the same request example above:

```
[Java]

Message response = interactionServerProtocol.request(loginRequest);
if (response.messageId() == EventAck.ID) {
        EventAck eventAck = (EventAck)response;
        // continue here
}
else {
        // handle the error here
}
```

Notice that you will need to cast the message to a specific message type in order to access its attributes. If a request fails on the server side, you will typically receive an `EventError`.

Take into account that the `Request` method blocks until a message is received or a timeout occurs. If the timeout elapses and no response was received from the server, then a `null` value is received. The timeout parameter can be specified in the request method. If you do not use the timeout parameter then, then the protocol `Timeout` property is used.

The Request method will only return one message from the server. In the case that the server returns subsequent messages, apart from the first response, as a result of the requested operation, then you must process those messages separately as unsolicited events. Please make sure that your code handles all messages received from your servers.

When using the Request method, your application only receives the response to that request as a return value. The response will not be received as an unsolicited event as well. (You can change this behavior by using the CopyResponse protocol property, described below.)

Receiving a Response Asynchronously

For many applications, blocking your thread while waiting for a response to your request is not appropriate. For example GUI applications, where the GUI can appear "frozen" if the response takes too much time to be received. It can also be true for batch applications that may want to send multiple requests at the same time, while waiting for all responses concurrently. For these scenarios it is possible to receive responses asynchronously.

Receiving a Response Asynchronously Using a Callback

By using requestAsync, your thread will not block, and it will permit you to handle the response by using callback methods that will get called asynchronously.

First, you will need to implement a CompletionHandler which will contain the logic for handling the response to your request:

[Java]

```
private static final CompletionHandler loginResponseHandler = new  CompletionHandler() {

        @Override
        public void completed(Message message, Void notUsed) {
                // handle message here
        }

        @Override
        public void failed(Throwable exc, Void notUsed) {
                // handle error here
        }

};
```

> **Important**
> The CompletionHandler callback methods will be executed by the protocol invoker.

Then you can use the CompletionHandler as a parameter to the requestAsync method:

[Java]

```
interactionServerProtocol.requestAsync(loginRequest, null, loginResponseHandler);
```

Notice that in this example, the attachment parameter has not been used. If you are sharing the same CompletionHandler implementation for handling the responses to different requests then you may want to use an attachment to make it easy to differentiate among those requests.

Receiving the Response as a Future

Alternatively, you may want to handle responses using the same thread that did the request, but have the option to do something concurrently while waiting for the response. To accomplish this, use the `beginRequest` method.

As an example, you might perform two agent login requests concurrently: one for logging into the T-Server, and another for logging into Interaction Server.

```
[Java]

RequestFuture loginVoiceFuture = tserverProtocol.beginRequest(loginVoiceRequest);
RequestFuture loginMultimediaFuture =
interactionServerProtocol.beginRequest(loginMultimediaRequest);

Message loginVoiceResponse = loginVoiceFuture.get();
Message loginMultimediaResponse = loginMultimediaFuture.get();

// handle responses, both are available now
```

When using the `requestAsync` or `beginRequest` methods, you will **not** receive the response as an unsolicited event. (You can change this behavior by using the CopyResponse protocol property, described below).

CopyResponse

Previously it was stated that responses returned by request methods are not received as unsolicited events by default. This behavior can be modified by using the protocol `CopyResponse` property. The default value is `false`, but it can be set to `true` like this:

```
[Java]

tserverProtocol.setCopyResponse(true);
```

This is particularly useful for protocols which define events that can be both received unsolicited and as a response to a client request (such as `EventAgentLogin` defined by the T-Server protocol). By setting the `CopyResponse` property to `true`, you can execute your agent state change logic only when handling the message as an unsolicited event, and you do not need to include it when receiving the message as a response.

.NET

Once you have connected to a server, much of the work for your application will involves sending messages to that server and handling the events you receive from the server. This article describes how to send and receive messages from a server.

Messages: Overview of Events and Requests

Messages you send to a server are called requests, while messages you receive are called events. An event that is received from a server as the result of executing a request is called a response. In summary, messages can be classified by using the following taxonomy:

- Requests: sent to the server
- Events: received from the server
    - Responses: received as the result of a request
    - Unsolicited events: not a direct result of a request

> **Tip**
> On this page, we will use the more general term "message" instead of "event", in order to avoid confusion between protocol events and programming events.

For example, you may send a request to log in an agent or to gather statistics. You might also send a request to update a configuration object, or to shut down an application.

In each of these cases, the server will respond with an event message, as shown below.

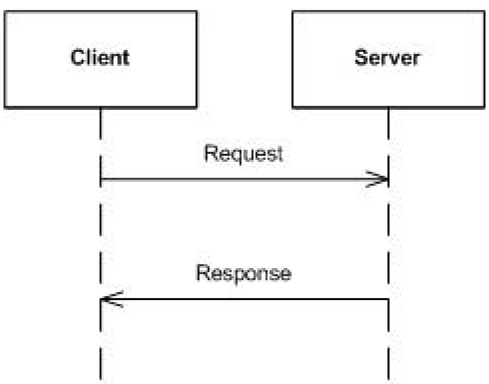

You may also get unsolicited events from the server. That means receiving events that are not a response to a specific request. For example, `EventRinging` will notify you of a call ringing on an extension that you are currently monitoring.

Receiving Messages

With the Platform SDK, you can receive messages synchronously or asynchronously. It is important that you define the way your application will work in this aspect. In general, you will probably use only one type or the other in the same application.

Interactive applications normally use asynchronous message handling, because that will prevent the UI thread from being blocked, which could make the application appear "frozen" to a user. On the other hand, non-interactive batch applications commonly use synchronous response handling, as that allows writing easy code that performs step-by-step.

Receiving Messages Asynchronously

Most Platform SDK applications need to handle unsolicited events. This is particularly true for applications that monitor the status of contact center resources, such as extensions.

You receive server messages asynchronously by subscribing to the `Received` .NET event:

```
[C#]

tserverProtocol.Received += OnTServerMessageReceived;
```

Then you can implement your event-handling logic:

```
[C#]

void OnTServerMessageReceived(object sender, EventArgs e)
{
        IMessage message = ((MessageEventArgs)e).Message;
        // your event-handling code goes here
}
```

> **Important**
> You need to know that your event-handling logic will be executed by using the protocol invoker. Please set the invoker appropriate for your application needs. For more information about the protocol invoker and how to set it, refer to Connecting to a Server.

Inside your event-handling code, you will want to execute different logic for different kinds of events. A typical way to do this is using a `switch` statement, based on the event identifier:

```
[C#]

switch (message.Id)
{
        case EventAgentLogin.MessageId:
                OnEventAgentLogin(message);
                break;
        case EventAgentLogout.MessageId:
                OnEventAgentLogout(message);
                break;
}
```

Receiving Messages Synchronously

Some kinds of applications, such as batch applications, benefit from receiving messages synchronously. This means that received messages will queue up and be handled by the application on demand.

In order to receive messages this way, you simply do not subscribe to the `Received` .NET event as described in the previous section.

> **Tip**
> For releases prior to Platform SDK 8.1.1, messages were received synchronously by default. Please note that 8.1.1 behavior is backwards-compatible, and pre-8.1.1 applications will continue to work as expected without any modification.

To receive a message synchronously, use the `Receive` method. This method blocks processing, waiting for the next message to be received before continuing. Take into account that the maximum time to wait is set by a configurable timeout value. If the timeout expires and no event is received, you will receive a `null` value.

```
[C#]

IMessage message = tserverProtocol.Receive();
```

If you want to set your own timeout, you can use the `Receive` method overload that takes a timeout parameter. Otherwise, if you use Receive with no parameters, the protocol `Timeout` property will be used.

Sending Requests Asynchronously

This is the easiest way to send a message to a server. Suppose you have created and filled a request object, for example, a `RequestAgentLogin` message for Interaction Server:

```
[C#]

var loginRequest = RequestAgentLogin.Create();
loginRequest.TenantId = tenantId;
loginRequest.AgentId = agentId;
loginRequest.PlaceId = placeId;
```

Then you can send it to the server:

```
[C#]

interactionServerProtocol.Send(loginRequest);
```

This will result in your application receiving a response from the Interaction Server: either an `EventAck` or an `EventError` message. By using the Send method, you will ignore that response at the place where you make the request. You will get the response, like any other unsolicited event, using the techniques described in the *Receiving Messages* section.

Handling Responses

The understanding of how to send requests and receive events is all you need to communicate with Genesys servers. However, the Platform SDK also provides the ability to easily associate a response with the particular request that originated it.

Receiving a Response Synchronously

The easiest way to handle responses is with the Request method. This is a blocking method, as your application stops to wait for a response to come from the server. Using the same request example above:

```
[C#]

IMessage response = interactionServerProtocol.Request(loginRequest);
if (response.Id == EventAck.MessageId)
{
        var eventAck = (EventAck)response;
        // continue here
}
else
{
        // handle the error here
}
```

Notice that you will need to cast the message to a specific message type in order to access its attributes. If a request fails on the server side, you will typically receive an EventError.

Take into account that the Request method blocks until a message is received or a timeout occurs. If the timeout elapses and no response was received from the server, then a null value is received. The timeout parameter can be specified in the request method. If you do not use the timeout parameter then the protocol Timeout property is used.

The request method will only return one message from the server. In the case that the server returns subsequent messages, apart from the first response, as a result of the requested operation, then you must process those messages separately as unsolicited events. Please make sure that your code handles all messages received from your servers.

When using the Request method, your application only receives the response to that request as a return value. The response will not be received as an unsolicited event as well. (You can change this behavior by using the CopyResponse protocol property, described below).

Receiving a Response Asynchronously

For many applications, blocking your thread while waiting for a response to your request is not appropriate. For example GUI applications, where the GUI can appear "frozen" if the response takes too much time to be received. It can also be true for batch applications that may want to send multiple requests at the same time, while waiting for all responses concurrently. For these scenarios it is possible to receive responses asynchronously.

By using BeginRequest, your thread will not block, and it will permit you to handle the response the way that best suits your application. This method complies with .NET "Asynchronous Programming Model". You can find more information about the "Asynchronous Programming Model" in the Web.

For example, your application can handle responses asynchronously by using a callback, which is a piece of logic that executes asynchronously when the response is received. Define a callback method like this:

```
[C#]

void OnLoginResponseReceived(IAsyncResult result) {
        IMessage response = interactionServerProtocol.EndRequest(result);
```

```
        if (response.Id == EventAck.MessageId)
        {
                var eventAck = (EventAck)response;
                // continue here
        }
        else
        {
                // handle the error here
        }
}
```

Then you can submit your request using the callback method.

[C#]

```
interactionServerProtocol.BeginRequest(loginRequest, OnLoginResponseReceived, null);
```

As an alternative, you may want to do something concurrently, while waiting for the response. For example, you could perform two agent login requests concurrently: one for logging the agent into the T-Server, and another for logging the agent into Interaction Server.

[C#]

```
var resultLoginVoice = tserverProtocol.BeginRequest(loginVoiceRequest, null, null);
var resultLoginMultimedia = interactionServerProtocol.BeginRequest(loginMultimediaRequest,
null, null);

var loginVoiceResponse = tserverProtocol.EndRequest(resultLoginVoice);
var loginMultimediaResponse = interactionServerProtocol.EndRequest(resultLoginMultimedia);

// handle responses, both are available now
```

When using the `BeginRequest` method, your application receives the response to your request as the return value of `EndRequest`. You will not receive the response as an unsolicited event. (You can change this behavior by using the `CopyResponse` protocol property, described below).

CopyResponse

Previously it was stated that responses returned by request methods are not received as unsolicited events by default. This behavior can be modified by using the protocol `CopyResponse` property. The default value is false, but it can be set to true like this:

[C#]

```
tserverProtocol.CopyResponse = true;
```

This is particularly useful for protocols which define events that can be both received unsolicited and as a response to a client request (such as `EventAgentLogin` defined by the T-Server protocol). By setting the `CopyResponse` property to `true`, you can execute your agent state change logic only when handling the message as an unsolicited event, and you do not need to include it when receiving the message as a response.

Setting up logging in Platform SDK

Logging for Java

Setting up log4j logging

The easiest way to set up Platform SDK logging in Java is to use the built-in integration with log4j. There are two possible ways to do this:

- Using code, by creating a `Log4JLoggerFactoryImpl` instance and setting it as the global logger factory for Platform SDK at the beginning of your program, like this:

```
com.genesyslab.platform.commons.log.Log.setLoggerFactory(new Log4JLoggerFactoryImpl());
```

Or:

- Using a Java system variable, by setting `com.genesyslab.platform.commons.log.loggerFactory` to the fully qualified name of the `ILoggerFactory` implementation class. For example, to set up log4j as the logging implementation you can start your application using the following command:

```
java
-Dcom.genesyslab.platform.commons.log.loggerFactory=com.genesyslab.platform.commons.log.Log4JLoggerFactoryImpl
<MyMainClass>
```

Providing a custom logging implementation

If log4j does not fit your needs, it is also possible to provide your own implementation of logging.

In order to do that, you will need to complete the following steps:

1. Implement the `ILogger` interface, which contains the methods that the Platform SDK uses for logging messages, by extending the `AbstractLogger` class.
2. Implement the `ILoggerFactory` interface, which should create instances of your `ILogger` implementation.
3. Finally, set up your `ILoggerFactory` implementation as the global Platform SDK `LoggerFactory`, as described above.

Setting Up Internal Logging for Platform SDK

To use internal logging in Platform SDK, you have to set a logger implementation in Log class *before* making any other call to Platform SDK. There are two ways to accomplish this:

1. Set the `com.genesyslab.platform.commons.log.loggerFactory` system property to the fully qualified

name of the factory class

2. Use the `Log.setLoggerFactory(...)` method

The only log factory available in Platform SDK itself is `com.genesyslab.platform.commons.log.Log4JLoggerFactoryImpl` which uses log4j. You will have to setup log4j according to your needs, but a simple log4j configuration file is shown below as an example.

```
log4j.logger.com.genesyslab.platform=DEBUG, A1
log4j.appender.A1=org.apache.log4j.FileAppender
log4j.appender.A1.file=psdk.log
log4j.appender.A1.layout=org.apache.log4j.PatternLayout
log4j.appender.A1.layout.ConversionPattern=%-4r [%t] %-5p %-25.25c %x - %m%n
```

The easiest way to set system property is to use `-D` switch when starting your application:

```
-Dcom.genesyslab.platform.commons.log.loggerFactory=com.genesyslab.platform.commons.log.Log4JLoggerFactoryImpl
```

Logging with AIL

In Interaction SDK (AIL) and Genesys Desktop applications, you can enable the Platform SDK logs by setting the option `log/psdk-debug = true`.

At startup, AIL calls: `Log.setLoggerFactory(new Log4JLoggerFactoryImpl());`

The default level of the logger `com.genesyslab.platform` is `WARN` (otherwise, applications would be literally overloaded with logs). The option is dynamically taken into account; it turns the logger level to `DEBUG` when set to true, and back to `WARN` when set to false.

Truncating Large Logs Using PSDK.DATA

Starting from PSDK 8.1.1, a special logger was added (in terms of log4j configuration) with the name: `PSDK.DATA`. It was initially designed for the configuration server protocol to:

1. truncate main Platform SDK logs, and
2. allow creation of logs with full protocol data dumps.

A sample of log4j configuration follows:

```
log4j.logger.com.genesyslab.platform=TRACE, A1
log4j.logger.PSDK.DATA=TRACE, A2

log4j.appender.A1=org.apache.log4j.ConsoleAppender
log4j.appender.A1.layout=org.apache.log4j.PatternLayout
log4j.appender.A1.layout.ConversionPattern=%-4r [%t] %-5p %-25.25c %x - %m%n

log4j.appender.A2=org.apache.log4j.FileAppender
log4j.appender.A2.File=d:\\psdkdata.log
log4j.appender.A2.Append=true
log4j.appender.A2.Threshold=TRACE
log4j.appender.A2.layout=org.apache.log4j.PatternLayout
log4j.appender.A2.layout.ConversionPattern=%-4r [%t] %-5p %-25.25c %x - %m%n
```

> **Tip**
>
> This feature has side effect: when log4j is configured to use "rootLogger" for all logs (including Platform SDK) then it may record protocol messages twice - once for the main logger and again for the data logger.

The goal of the extension is to resolve issues where large log files affect application performance. For example, an application may read a lot of configuration objects and require Platform SDK logging to be enabled. In this case, a configuration protocol message that arrives containing 1 MB of packed data could lead to roughly 6 MB of log data which (in most cases) is not required. These large log records can be truncated, recording enough data to ensure that configuration information is available and that data flow is ok.

In some case, a full data dump may be required in logs. In this case, the truncation enabling parameter is passed to the static context of `ToStringHelper`, which generates a string representation of abstract protocol messages with attributes. Creating an additional logger that manages separated protocol messages with this context may be useful at times, while initializing the `com.genesyslab.platform` logger for general Platform SDK logging without enabling full dumps by default is better in most cases.

Using JVM system properties, which are checked before log record generation and can enable/disable full data dumps, is an alternative way to handle this scenario. It may be a preferred solution, although usage of system properties may not work depending on how application containers are used.

Logging for .NET

Setting up logging

For .NET development, the `EnableLogging` method allows logging to be easily set up for any classes that implement the `ILogEnabled` interface. This includes:

- All protocol classes: `TServerProtocol, StatServerProtocol,` etc.
- The `WarmStandbyService` class of the Warm Standby Application Block.

For example:

```
tserverProtocol.EnableLogging(new MyLoggerImpl());
```

Providing a Custom Logging Implementation

You can provide your custom logging functionality by implementing the `ILogger` interface. Samples of how to do this are provided in the following section.

Samples

You can download some samples of classes that implement the `ILogger` interface:

- AbstractLogger: This class can make it easier to implement a custom logger, by providing a default implementation of `ILogger` methods.
- TraceSourceLogger: A logger that uses the .NET TraceSource framework. It adapts the Platform SDK logger hierarchy to the non-hierarchical TraceSource configuration.
- Log4netLogger: A logger that uses the log4net libraries.

Advanced Platform SDK Topics

Advanced Platform SDK Topics

The following articles provide details about advanced Platform SDK features you may want to take advantage of:

- Secure Connections Using TLS
    - Quick Start
    - Using the Platform SDK Commons Library
    - Using the Application Template Application Block
    - Configuring TLS Parameters in Configuration Manager
    - Using and Configuring Security Providers
    - OpenSSL Configuration File
    - Use Cases
- Lazy Parsing of Message Attributes

Secure connections using TLS

This page provides an introduction to creating and configuring Transport Layer Security (TLS) for your Platform SDK connections, as introduced in release 8.1.1.

Introduction to TLS

This page provides an overview of the TLS implementation provided in the 8.1.1 release of Platform SDK. It introduces Platform SDK users to TLS concepts and then provides links to expanded articles and examples that describe implementation details.

Before working with TLS to create secure connections, you should have a basic awareness of how public key cryptography works.

Certificates

Transport Layer Security (TLS) technology uses public key cryptography, where the key required to encrypt and decrypt information is divided into two parts: a public key and a private key. These parts are reciprocal in the sense that data encrypted using a private key can be decrypted with the public key and vice versa, but cannot be decrypted using the same key that was used for encryption.

There is an X.509 standard for public key (certificate) format, and public-key cryptography standards (PKCS) that define format for private key (PKCS#8) and related data structures.

Certificate Authority (CA)

In the context of TLS, a CA is an entity that is trusted by both sides of network connection. Each CA has a public X.509 certificate and owns a related private key that kept secret. A CA can generate and sign certificates for other parties using its private key, and then that CA certificate can be used by the parties to validate their certificates. A CA can also issue public Certificate Revocation Lists (CRLs), which are also used by parties for certificate validation.

The relation between certificates and CRL can be depicted like this:

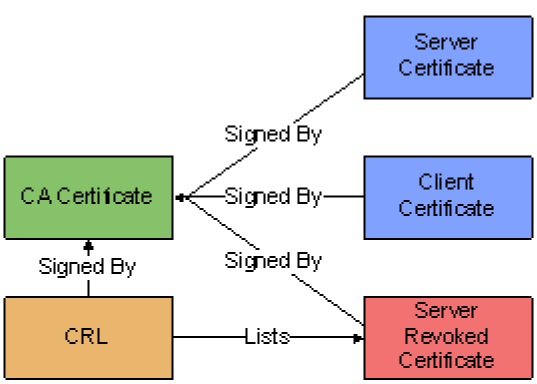

Certificate Usage

To create a secure connection, each party must have a copy of:

- a public CA certificate
- a CRL issued by the CA
- their own public certificate (with a corresponding private key)

When a network connection is established, the client initiates a TLS handshake process during which the parties exchange their public certificates, prove that they own corresponding private keys, create a shared session encryption key, and negotiate which cipher suite will be used.

Placement and exchange of certificate data is shown on the following diagram:

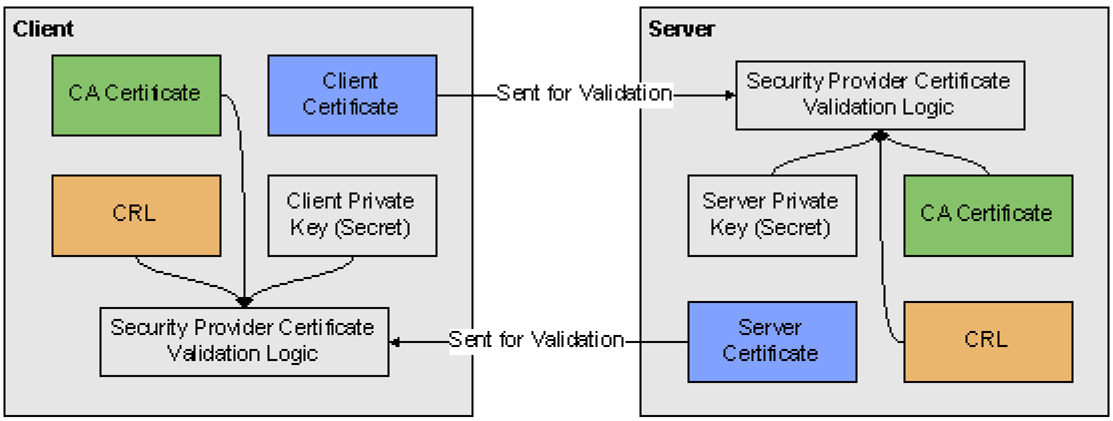

TLS only requires that servers send their certificates, but the client certificates can also be exchanged depending on server settings. Cases where the client certificates are demanded by the server are called “Mutual TLS”, as both sides send their certificates.

If all certificates pass validation and the ciphers are negotiated successfully, then a TLS connection is established and higher-level protocols may proceed.

Implementing and Configuring TLS

Genesys strongly recommends reading all TLS in Platform SDK articles in order to get understanding of how TLS works in general and how it is supported in Platform SDK. A Quick Start page is provided for reference, but the specific implementation details and expanded information provided in other pages will help you to better understand how to provide TLS support in your applications. Once you have an understanding of how TLS is implemented, you can use the Use Case guide to quickly find code snippets or relevant links for common tasks.

There are two main ways to implement TLS in your Platform SDK code:

1. Use the Platform SDK Commons Library to specify TLS settings directly when creating endpoints
2. Use the Application Template Application Block to read connection parameters inside configuration

objects retrieved from Configuration Server, then use those parameters to configure TLS settings.

Note: If using the Application Template Application Block, you will need to configure TLS Parameters in Configuration Manager before the application is tested.

Recommendations are also provided for the configuration and use of security providers. The security providers discussed on that page have been tested within the described configurations, and worked reliably.

Migrating TLS Support From Previous Versions of Platform SDK

Platform SDK for Java

Platform SDK 8.1.0 had the following connection configuration parameters for TLS:

- *Connection.TLS_KEY*
- *Connection.SSL_KEYSTORE_PATH_KEY*
- *Connection.SSL_KEYSTORE_PASS*

The *TLS_KEY* parameter is the equivalent of *enableTls* flag in the current release, while the other parameters specified the location and password for the Java keystore file containing certificates that were used by the application to authenticate itself. TLS configuration code looked like this:

```
ConnectionConfiguration connConf = new KeyValueConfiguration(new KeyValueCollection());
connConf.setOption(Connection.TLS_KEY, "1");
connConf.setOption(Connection.SSL_KEYSTORE_PATH_KEY, "c:/certificates/client-certs.keystore");
connConf.setOption(Connection.SSL_KEYSTORE_PASS, "pa$$w0rd");
```

In Platform SDK 8.1.1, this code can be translated to the following:

```
boolean tlsEnabled = true;
// By default, PSDK 8.1.0 trusted any certificate
TrustManager trustManager = TrustManagerHelper.createTrustEveryoneTrustManager();
// Keystore entries may be protected with individual password,
// but usually, these passwords are the same as keystore password
KeyManager keyManager = KeyManagerHelper.createJKSKeyManager(
    "c:/certificates/client-certs.keystore", "pa$$w0rd", "pa$$w0rd");
SSLContext sslContext = SSLContextHelper.createSSLContext(keyManager,
trustManager);
```

In most cases, certificates from other parties will need to be validated. Assuming there is a separate keystore file with a CA certificate, this can be achieved with the following code:

```
TrustManager trustManager = TrustManagerHelper.createJKSTrustManager(
    "c:/certificates/CA-cert.keystore", "pa$$w0rd", null, null);
```

Please note that different keystore files are used for the *KeyManager* and *TrustManager* objects. For more information, see Using the Platform SDK Commons Library.

Platform SDK for .Net

There were no significant changes to interfaces for the .NET version of Platform SDK 8.1.1. In this case, the same code would work for 8.1.0 and 8.1.1 releases:

```
KeyValueConfiguration config = new KeyValueConfiguration(new KeyValueCollection());
config.TLSEnabled = true;
config.TlsCertificate = "29 3f 0d d9 65 a1 a9 92 dd 1c 8c 2a e7 20 74 06 c5 ba 0f 10";
Endpoint ep = new Endpoint(AppName, Host, Port, config);
```

Known Issues

For more details about the known issues listed here, refer to Using and Configuring Security Providers.

- Java 5: MSCAPI provider is not supported.
- Java 6:
  - MSCAPI provider is only supported in 32-bit version since update 27: http://bugs.sun.com/bugdatabase/view_bug.do?bug_id=6931562.
  - MSCAPI provider is only supported in 64-bit version since update 38: http://bugs.sun.com/bugdatabase/view_bug.do?bug_id=2215540.
  - CRLs located in WCS are ignored, please use CRLs as files.
- Java 7:
  - CRL files without extension section cannot be loaded: http://bugs.sun.com/bugdatabase/view_bug.do?bug_id=7166885.
    Note: Although the bug is marked as "Will not fix", it seems to be fixed since Java 7 update 7.
  - CRLs located in WCS are ignored, please use CRLs as files.
- MSCAPI: MSCAPI does not have a documented way of programmatic setting of password to private key stored in WCS. Regardless of password returned by CallbackHandler; if private key is protected with confirmation prompt or password prompt, user will be shown OS popup dialog.

Quick Start

Understanding Port Modes

TLS is configured differently depending on target port mode:

- default - Default mode ports do not use or understand TLS protocol.
- upgrade - Upgrade mode ports allow unsecured connections to be made, switching to TLS mode only after TLS settings are retrieved from Configuration Server.
- secure - Secure mode ports require TLS to be started immediately, before sending any requests to server.

Connecting to Default Mode Ports

Default mode is supported for all protocols; no specific configuration is needed for it to work.

Example:

```
Endpoint cfgServerEndpoint = new Endpoint(appName, cfgHost, cfgPort);
ConfServerProtocol protocol = new ConfServerProtocol(cfgServerEndpoint);
protocol.setClientName(appName);
protocol.setClientApplicationType(appType);
protocol.setUserName(username);
protocol.setUserPassword(password);
protocol.open();
```

It is also OK to specify explicit null parameters for the connection configuration and TLS parameters:

```
// Explicit null ConnectionConfiguration
Endpoint cfgServerEndpoint = new Endpoint(appName, cfgHost, cfgPort, null);

// Explicit null ConnectionConfiguration and TLS parameters
Endpoint cfgServerEndpoint = new Endpoint(appName, cfgHost, cfgPort, null, false, null, null);
```

Connecting to Upgrade Mode Ports

TLS upgrade mode is supported only for Configuration Protocol, since the TLS settings for connecting clients must be retrieved from Configuration Server. No specific options are required; the TLS upgrade logic works by default.

If a user has provided custom settings, then those settings are used if the TLS parameters received from Configuration Server are empty. The only requirement that the *tlsEnabled* parameter in the Endpoint constructor is **not** to true, otherwise the client side starts TLS immediately and the connection would fail because an upgrade mode port expects the connection to be unsecured initially.

```
// Setting tlsEnabled to true would cause failure when connecting to upgrade port:
Endpoint cfgServerEndpoint = new Endpoint(appName, cfgHost, cfgPort,
```

```
    connConf, true, sslContext, sslOptions);
```

Connecting to Secure Mode Port

Secure mode is supported for all protocols. TLS configuration objects/properties must be specified before the connection is opened, and the *tlsEnabled* parameter must be set to true. Secure port mode expects the client to start TLS negotiation immediately after connecting, otherwise the connection fails.

Example:

```
boolean tlsEnabled = true;
// Here, the minimal TLS configuration is used, see the following section for details
TrustManager trustManager = TrustManagerHelper.createTrustEveryoneTrustManager();
KeyManager keyManager = KeyManagerHelper.createEmptyKeyManager();
SSLContext sslContext = SSLContextHelper.createSSLContext(keyManager, trustManager);
ConnectionConfiguration connConf = new KeyValueConfiguration(new KeyValueCollection());
Endpoint cfgServerEndpoint = new Endpoint(appName, cfgHost, cfgPort,
    connConf, tlsEnabled, sslContext, sslOptions);
ConfServerProtocol protocol = new ConfServerProtocol(cfgServerEndpoint);
protocol.setClientName(appName);
protocol.setClientApplicationType(appType);
protocol.setUserName(username);
protocol.setUserPassword(password);
protocol.open();
```

TLS Minimal Configuration

Frequently, there is a need to quickly set up code for working TLS connections, dealing with detailed TLS configuration later. The minimal configuration settings described below do exactly that.

Platform SDK for Java

The following code creates an *SSLContext* object that can be used to configure a connection to a secure port or to configure a secure server socket. This code uses *EmptyKeyManager* which indicates that the party opening connection/socket would not have any certificate to authenticate itself, and *TrustEveryoneTrustManager* which trusts any certificate presented by the other party - even expired or revoked certificates.

```
boolean tlsEnabled = true;
TrustManager trustManager = TrustManagerHelper.createTrustEveryoneTrustManager();
KeyManager keyManager = KeyManagerHelper.createEmptyKeyManager();
SSLContext sslContext = SSLContextHelper.createSSLContext(keyManager,
trustManager);
```

Note: Connections using this configuration would have a working encryption layer, but they are not secure because they can neither authenticate themselves nor validate credentials provided by the other party.

Note: If a server uses mutual TLS mode, then it requires the client to present a certificate. Minimal configuration does not have certificates, so in this case the TLS negotiation would fail.

Platform SDK for .Net

Platform SDK for .Net requires less configuration, because it always uses the MSCAPI security provider and Windows Certificate Services (WCS) by default. The following code would trust all certificates located in the WCS Trusted Root Certificates folder for the current user account.

```
KeyValueConfiguration config = new KeyValueConfiguration(new KeyValueCollection());
config.TLSEnabled = true;
Endpoint ep = new Endpoint(AppName, Host, Port, config);
```

Note: If a server uses mutual TLS mode, then it requires clients to present a certificate. Minimal configuration does not have certificates, so in this case the TLS negotiation would fail.

Using the Platform SDK Commons Library

Using the Platform SDK Commons Library to Configure TLS

Starting with Platform SDK 8.1.1, the only way to configure connections is by using `Endpoint` objects, which contain all parameters related to the endpoint connection—including TLS parameters that indicate whether TLS is enabled and provide details about the SSL context and extended options.

Note: In earlier releases, Platform SDK provided three ways to configure connections:

- using `ConnectionConfiguration` objects passed to `Protocol` constructors
- setting parameters in the protocol context
- adding a textual parameter representation to the URL query

The following diagrams show interdependencies among the Platform SDK objects used to establish network connections and support TLS.

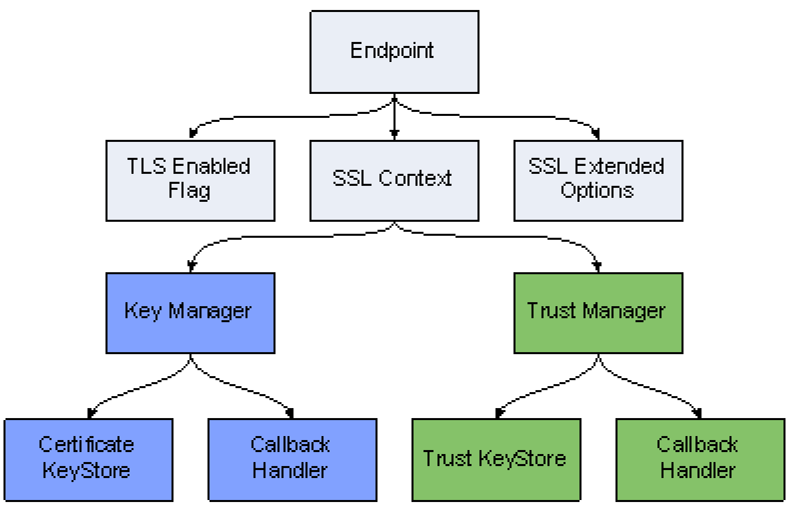

TLS Configuration Objects Containment Hierarchy

This page outlines each step required to create supporting objects for a TLS-enabled `Endpoint`.

Callback Handlers

In many cases, certificate or key storage is password-protected. This means that Platform SDK will need the password to access storage. The Java `CallbackHandler` interface offers a flexible way to pass this type of credential data:

```
package javax.security.auth.callback;
```

```
...
public interface CallbackHandler {
    void handle(Callback[] callbacks)
        throws java.io.IOException, UnsupportedCallbackException;
}
```

The `handle()` method accepts credential requests in the form of `Callback` objects that have appropriate setter methods. The most common callback implementation is `PasswordCallback`. User code may use a GUI to ask the end user to:

- enter a password
- retrieve a password from a file, pipe, network, and so on

Here is an example of a `CallbackHandler` delegating password retrieval to a GUI:

```
CallbackHandler callbackHandler = new CallbackHandler() {
    public void handle(Callback[] callbacks) throws IOException,
            UnsupportedCallbackException {
        for (Callback c : callbacks) {
            if (c instanceof PasswordCallback) {
                PasswordCallback p = (PasswordCallback) c;
                p.setPassword(gui.getKeyStorePassword());
            }
        }
    }
};
```

When No Password is Required

In some cases, certificate storage does not need a password. The API may still dictate that a `CallbackHandler` be provided however, so the Platform SDK includes a predefined class that can be used as a "dummy" `CallbackHandler` for this scenario:

```
com.genesyslab.platform.commons.connection.tls.DummyPasswordCallbackHandler
```

Here is an example of using this dummy class:

```
CallbackHandler callbackHandler = new DummyPasswordCallbackHandler();
```

Key Managers

Java provides a `KeyManager` interface. This interface defines functionality that can be used to load and contain certificates or keys, or to select appropriate certificates or keys.

Classes based on the `KeyManager` interface are used by Java TLS support to retrieve certificates that will be sent over the network to a remote party for validation. They are also used to retrieve the corresponding private keys. On the client side, `KeyManager` classes retrieve client certificates or keys; on the server side they retrieve server certificates or keys.

The Platform SDK Commons library has a helper class, `KeyManagerHelper`, which makes it easy to create key managers using several types of key stores and security providers. The built-in key manager types are:

- **PEM** — reads certificate/key pairs from X.509 PEM files.
- **MSCAPI** — uses the Microsoft CryptoAPI and Windows certificate services to retrieve certificate/key

pairs.

- **PKCS11** — delegates to an external security provider plugged in via the PKCS#11 interface, for example, Mozilla NSS.
- **JKS** — retrieves a certificate/key pair from a Java Keystore file.
- **Empty** — does not retrieve anything. This type is for use as a dummy key manager. For example, clients that do not have certificates can use it.

Here are some examples of key manager creation:

```
// From PEM file
X509ExtendedKeyManager km = KeyManagerHelper.createPEMKeyManager(
        "c:/cert/client-cert.pem", "c:/cert/client-cert-key.pem");

// From MSCAPI
CallbackHandler cbh = new DummyPasswordCallbackHandler();
// Whitespace characters are allowed anywhere inside the string
String certThumbprint =
        "4A 3F E5 08 48 3A 00 71 8E E6 C1 34 56 A4 48 34 55 49 D9 0E";
X509ExtendedKeyManager km = KeyManagerHelper.createMSCAPIKeyManager(
        cbh, certThumbprint);

// From PKCS11
// This provider does not allow customization of Key Manager
// This is required for FIPS-140 certification
// Dummy callback handler will not work, must use strong password
CallbackHandler passCallback = ...;
X509ExtendedKeyManager km = KeyManagerHelper.createPKCS11KeyManager(
        passCallback);

// From JKS
// JKS key store does not allow callback usage (bug in Java?)
// Individual entries in JKS key store can be password-protected
char[] keyStorePass = "keyStorePass".toCharArray();
char[] entryPass = "entryPass".toCharArray();
X509ExtendedKeyManager km = KeyManagerHelper.createJKSKeyManager(
        "c:/cert/client-cert.jks", keyStorePass, entryPass);

// Empty key manager
// Using KeyManagerHelper class
X509ExtendedKeyManager km1 = KeyManagerHelper.createEmptyKeyManager();
// Direct creation
X509ExtendedKeyManager km2 = new EmptyX509ExtendedKeyManager();
```

Trust Managers

A Trust Manager is an entity that decides which certificates from a remote party are to be trusted. It performs certificate validation, checks the expiration date, matches the host name, checks the certificate against a CRL list, and builds and validates the chain of trust. The chain of trust starts from a certificate trusted by both sides (for example, a CA certificate) and continues with second-level certificates signed by CA, then possibly with third-level certificates signed by second-level authorities and so on. Chain length can vary, but Platform SDK was designed to explicitly support two-level chains consisting of a CA certificate and a leaf certificate signed by CA.

Trust manager instances are created based on storage that contains trusted certificates. The number of trusted certificates can vary depending on the type of trust manager being used. With PEM files, the storage contains only a single CA certificate; other provider types can have larger sets of trusted certificates.

The Platform SDK Commons library has a helper class, `TrustManagerHelper`, which makes it easy to create trust managers that use several types of certificate stores and security providers, and which can accept additional parameters that affect certificate validation. Built-in trust manager types are:

- **PEM** — Reads a CA certificate from an X.509 PEM file.
- **MSCAPI** — Uses the Microsoft CryptoAPI and Windows certificate services to retrieve CA certificates and validate certificates.
- **PKCS11** — Delegates certificate validation to an external security provider plugged in via the PKCS#11 interface, for example, Mozilla NSS.
- **JKS** — Retrieves a CA certificate from a Java Keystore file and uses Java built-in validation logic.
- **Default** — Uses trusted certificates shipped with or configured in Java Runtime and Java built-in validation logic.
- **TrustEveryone** — Trusts any certificates. Can be used on the server side when you do not expect any certificates from clients, or during testing.

Here are some examples of trust manager creation (with generic `crlPath` and `expectedHostName` parameters defined in the first example):

```
// Generic parameters for trust manager examples
String crlPath = "c:/cert/ca-crl.pem";
String expectedHostName = "serverhost";
// From PEM file
X509TrustManager tm = TrustManagerHelper.createPEMTrustManager(
        "c:/cert/ca.pem", crlPath, expectedHostName);

// From MSCAPI
// CRL is loaded from PEM file (Platform SDK supports only file-base CRLs)
// Concrete CA is not specified, all certificates from WCS Trusted Root are used
CallbackHandler cbh = new DummyPasswordCallbackHandler();
X509TrustManager tm = TrustManagerHelper.createMSCAPITrustManager(
        cbh, crlPath, expectedHostName);

// From PKCS#11
// This provider implementation in Java does not allow custom host name check,
// but CRL can still be used
X509TrustManager tm = TrustManagerHelper.createPKCS11TrustManager(
        cbh, crlPath);

// From JKS
// JKS key store does not allow callback usage (bug in Java?)
// Certificate-only entries cannot have passwords in JKS key store
// CRL and host name check are supported
char[] keyStorePass = "keyStorePass".toCharArray();
X509ExtendedKeyManager km = KeyManagerHelper.createJKSTrustManager(
        "c:/cert/ca-cert.jks", keyStorePass, crlPath, expectedHostName);

// From Java built-in trusted certificates
// This one does not support CRL and host name check
X509ExtendedKeyManager km = KeyManagerHelper.createDefaultTrustManager();

// Trust Everyone
X509ExtendedKeyManager km =
        KeyManagerHelper.createTrustEveryoneTrustManager();
```

SSLContext and SSLExtendedOptions

An `SSLContext` instance serves as a container for all SSL and TLS parameters and objects and also as a factory for `SSLEngine` instances.

`SSLEngine` instances contain logic that deals directly with TLS handshaking, negotiation, and data encryption and decryption. `SSLEngine` instances are not reusable and must be created anew for each connection. This is a good reason for requiring users to provide an `SSLContext` instance rather than an instance of `SSLEngine`. `SSLEngine` instances are created by the Platform SDK connection layer and are not exposed to user code.

Only some of the parameters for `SSLEngine` can be pre-set in `SSLContext`. However, the `SSLExtendedOptions` class may be used to collect additional parameters.

`SSLExtendedOptions` currently contains two parameters:

- the "mutual TLS" flag
- a list of enabled cipher suites

The mutual TLS flag is used only by server applications. When the flag is turned on, the server will require connecting clients to send their certificates for validation. The connections of any clients that do not send certificates will fail.

The list of enabled cipher suites contains the names of all cipher suites that will be used as filters for `SSLEngine`. As a result, only ciphers that are supported by `SSLEngine` and that are contained in the enabled cipher suites list will be enabled for use.

Platform SDK includes the `SSLContextHelper` helper class to support one-line creation of `SSLContext` and `SSLExtendedOptions` instances.

Here are some examples:

```
// Creating SSLContext
KeyManager km = ...;
TrustManager tm = ...;
SSLContext sslContext = SSLContextHelper.createSSLContext(km, tm);

String[] cipherList = new String[] {
        "TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA",
        "TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA",
        "TLS_ECDH_RSA_WITH_3DES_EDE_CBC_SHA"};
// Can be single String with space-separated suite names
String cipherNames = "TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA " +
        "TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA " +
        "TLS_ECDH_RSA_WITH_3DES_EDE_CBC_SHA";
boolean mutualTLS = false;

// Creating SSLExtendedOptions directly
SSLExtendedOptions sslOpts1 =
        new SSLExtendedOptions(mutualTLS, cipherList);
SSLExtendedOptions sslOpts2 =
        new SSLExtendedOptions(mutualTLS, cipherNames);

// Create SSLExtendedOptions using the helper class:
SSLExtendedOptions sslOpts3 =
        SSLContextHelper.createSSLExtendedOptions(mutualTLS, cipherList);
SSLExtendedOptions sslOpts4 =
```

```
        SSLContextHelper.createSSLExtendedOptions(mutualTLS, cipherNames);
```

Endpoints

Now that supporting objects have been created and configured, you are ready to create an `Endpoint`.

The connection configuration parameters of an `Endpoint` are read-only—they cannot be changed after the `Endpoint` is created. This configuration information is then used by `Protocol` instances, the warm standby service, the connection layer and the TLS layer.

A sample `Endpoint` configuration is shown below:

```
ConnectionConfiguration connConf = ...;
SSLContext sslContext = ...;
SSLExtendedOptions sslOpts = ...;
tlsEnabled = true;
// Specifying host name and port.
Endpoint ep1 = new Endpoint("Server-1", "serverhost", 9090, connConf,
        tlsEnabled, sslContext, sslOpts);
// Specifying URI. Query part is still supported.
String uri = "tcp://Server-1@serverhost:9090/" +
        "?protocol=addp&addp-remote-timeout=5&addp-trace=remote";
Endpoint ep2 = new Endpoint("Server-1", uri, connConf,
        tlsEnabled, sslContext, sslOpts);
```

Note: Configuration parameters can be set directly in a `Protocol` instance context, but will be overwritten and lost under the following conditions:

- a new `Endpoint` is set up
- the protocol is forced to reconnect
- a warm standby switchover occurs

Configuring TLS for Client Connections

Using the information above, you are now ready to configure actual client connections.

Example:

```
// Get TLS configuration objects for connection
String clientName = "ClientApp";
String host = "serverhost";
int port = 9000;
SSLContext sslContext = ...; // Assume it is created
SSLExtendedOptions sslOptions = ...; // Assume it is created
boolean tlsEnabled = true;

ConnectionConfiguration connConf = new KeyValueConfiguration(new KeyValueCollection());
Endpoint epTSrv = new Endpoint(
        clientName, host, port, connConf, tlsEnabled, sslContext, sslOptions);

TServerProtocol tsProtocol = new TServerProtocol(epTSrv);
tsProtocol.setClientName(clientName);
tsProtocol.open();
```

Configuring TLS for Servers

Using the information above, you are now ready to configure actual server connections.

```
String serverName = "ServerApp";
String host = "serverhost";
int port = 9000;
SSLContext sslContext = ...; // Assume it is created
SSLExtendedOptions sslOptions = ...; // Assume it is created
boolean tlsEnabled = true;

ConnectionConfiguration connConf = new KeyValueConfiguration(new KeyValueCollection());
Endpoint epTSrv = new Endpoint(
        serverName, host, port, connConf, tlsEnabled, sslContext, sslOptions);

ExternalServiceProtocolListener serverChannel =
        new ExternalServiceProtocolListener(endpoint);
```

Parameter-based TLS Configuration

Platform SDK has a way to create TLS objects based on a set of parameters in a more declarative fashion rather than creating them programmatically. This feature was initially developed as a part of Application Template to configure TLS based on parameters from Configuration objects and then was generalized to use different parameter sources and moved to Commons. Currently this mechanism supports only three providers: PEM, MSCAPI and PKCS#11. Usage sequence is the following:

1. Prepare a source of TLS parameters and parse it using TLSConfigurationParser resulting in TLSConfiguration instance.
2. Customize TLSConfiguration.
    1. Add callback handlers.
    2. Clients: set expected host name.
3. Create SSLContext and SSLExtendedOptions from TLSConfiguration.

This section continues with step-by-step examples and ends with a more detailed review of helper classes.

Parsing TLS Parameters

Platform SDK Commons has a few helper classes that make it easier to extract TLS parameters from a properties files, command-line arguments, etc.: TLSConfiguration and TLSConfigurationParser. TLSConfiguration is a container for parsed TLS parameters and TLSConfigurationParser provides a general parsing method and several overloaded shortcut methods for specific cases.

Examples:

```
// Using KVList as a parameters source
KVList tlsProps = new KeyValueCollection();
tlsProps.addObject("tls", "1");
tlsProps.addObject("certificate", "client-cert.pem");
TLSConfiguration tlsConfClient =
        TLSConfigurationParser.parseClientTlsConfiguration(tlsProps);
```

```
TLSConfiguration tlsConfServer =
        TLSConfigurationParser.parseServerTlsConfiguration(tlsProps);

// Using Map as a parameters source
Map<String, String> tlsProps = new HashMap<String, String>();
tlsProps.put("tls", "1");
tlsProps.put("certificate", "client-cert.pem");
TLSConfiguration tlsConfClient =
        TLSConfigurationParser.parseClientTlsConfiguration(tlsProps);
TLSConfiguration tlsConfServer =
        TLSConfigurationParser.parseServerTlsConfiguration(tlsProps);

// Using Properties as a parameters source
Properties tlsProps = new Properties();
tlsProps.load(new FileInputStream("tls.properties"));
TLSConfiguration tlsConfClient =
        TLSConfigurationParser.parseClientTlsConfiguration(tlsProps);
TLSConfiguration tlsConfServer =
        TLSConfigurationParser.parseServerTlsConfiguration(tlsProps);

// Using String as a parameters source
// Format corresponds to Transport Parameters as they appear in Configuration Manager
String tlsProps = "tls=1;certificate=client-cert.pem"; // No spaces around ";"
TLSConfiguration tlsConfClient =
        TLSConfigurationParser.parseClientTlsConfiguration(tlsProps);
TLSConfiguration tlsConfServer =
        TLSConfigurationParser.parseServerTlsConfiguration(tlsProps);
```

Customizing TLS Configuration

When TLSConfiguration is prepared, it may still need some customization. Callback handlers for password retrieval, for example, cannot be configured in parameters and must be set explicitly. They should be set always, even if not used, because some security providers require them.

Specifying expected host name is not very straightforward and some aspects should be considered. When configuring TLS on client side, expected host names are in most cases different for primary and for backup connections. Though, on some virtualized environments, they can be the same. Users may choose to use IP addresses instead of DNS host names, or use DNS names with wildcards. Either way, expected host name must match one of names specified in server’s certificate and in extreme cases it may not relate to actual host name at all. To account for these cases, setting expected host name is not automated in Platform SDK and left for user code. Example code below shows how to set this value to actual host name of target server.

According to X.509 specification, certificate may contain not just host name or IP address, but also URI or e-mail address. Platform SDK supports only host names and IP addresses, but host name may use wildcard: a star symbol, “*”, can be used instead of any one level of domain name.

Examples:

```
TLSConfiguration tlsConfiguration = ...;

// Applicable to both clients and servers
// Passwords are not used, so set dummies:
tlsConfiguration.setKeyStoreCallbackHandler(
        new DummyPasswordCallbackHandler());
tlsConfiguration.setTrustStoreCallbackHandler(
        new DummyPasswordCallbackHandler());

// In case some real password is needed:
tlsConfiguration.setKeyStoreCallbackHandler(new CallbackHandler() {
```

```
    public void handle(Callback[] callbacks) {
        char[] password = new char[] {
                'p', 'a', 's', 's', 'w', 'o', 'r', 'd'};
            for (Callback c : callbacks) {
                if (c instanceof PasswordCallback) {
                    ((PasswordCallback) c).setPassword(password);
                }
            }
        }
    }
);

// Expected host name may contain exact host name, ...
tlsConfiguration.setExpectedHostname("someserver.ourdomain.com");
// wildcard host name, ...
tlsConfiguration.setExpectedHostname("*.ourdomain.com");
tlsConfiguration.setExpectedHostname("someserver.*.com");

// IPv4 address, ...
tlsConfiguration.setExpectedHostname("192.168.1.1");
// IPv6 address.
tlsConfiguration.setExpectedHostname("fe80::ffff:ffff:fffd");
```

Creating SSLContext

Platform SDK Commons has helper class – TLSConfigurationHelper, which creates SSLContext and SSLExtendOptions based on TLSConfiguration object. TLSConfigurationHelper has two methods:

```
public static SSLContext createSslContext(TLSConfiguration config);
```

and

```
static SSLExtendedOptions createSslExtendedOptions(TLSConfiguration config);
```

Method createSSLContext() determines security provider type if it is not set explicitly, creates necessary key store objects, key manager, trust manager, and finally wraps it all into SSLContext.

Method createSSLExtendedOptions() does not contain any logic, it just creates new SSLExtendedOptions with the exact parameters taken from TLSConfiguration.

Usage of both methods is shown in code sample below.

Example:

```
// TLS preparation section follows
KVList tlsProps = new KeyValueCollection();
tlsProps.addObject("tls", "1");
tlsProps.addObject("certificate", "client-cert.pem");
TLSConfiguration tlsConf =
        TLSConfigurationParser. parseClientTlsConfiguration(tlsProps);

boolean tlsEnabled = true;

SSLContext sslContext =
        TLSConfigurationHelper.createSslContext(tlsConfiguration);
SSLExtendedOptions sslOptions =
        TLSConfigurationHelper.createSslExtendedOptions(tlsConfiguration);

// The same as above, using shortcut methods:
sslContext = tlsConfiguration.createSslContext();
```

```
sslOptions = tlsConfiguration.createSslExtendedOptions();

Endpoint ep = new Endpoint(appName, host, port, null, tlsEnabled, sslContext, sslOptions);
```

TLSConfiguration Class

TLSConfiguration class is used as intermediate container to keep stronger-typed TLS parameters extracted from a parameter source. It contains the following:

Properties

TLSConfiguration Properties List

Name	Type	Description
tlsEnabled	boolean	Correspond to TLS parameters in Configuration; please see the list of TLS Parameters in Configuration Manager for details.
provider	String	
certificate	String	
certificateKey	String	
trustedCaCertificate	String	
mutual	boolean	
crl	String	
targetNameCheckEnabled	boolean	
cipherList	String	
fips140Enabled	boolean	
clientMode	boolean	Should be set to true for client-side of connection and false for server-side. `TLSConfigurationParser` specialized methods set it automatically.
expectedHostname	String	Host name to check against, used when `targetNameCheckEnabled` is turned on. Typically is used by client side and assigned to the host/domain part of target URL.
keyStoreCallbackHandler	CallbackHandler	Please see Callback Handlers for details.
trustStoreCallbackHandler	CallbackHandler	

Methods

TLSConfiguration Methods List

Signature	Description
SSLContext createSslContext()	A shortcut for `TLSConfigurationHelper.createSslContext`

Signature	Description
	method. Creates and configures `SSLContext` object based on the properties values.
SSLExtendedOptions createSslExtendedOptions()	A shortcut for `TLSConfigurationHelper.createSslExtendedOptions` method. Creates `SSLExtendedOptions` object based on the properties values.

Constants

The following constants define supported values for a provider property:

- String TLS_PROVIDER_PEM_FILE;
- String TLS_PROVIDER_PKCS11;
- String TLS_PROVIDER_MSCAPI;

TLSConfigurationParser Class

TLSConfigurationParser class has methods that extract TLS parameters from different sources and create TLSConfiguration instance containing the parameters. It uses interface PropertyReader and several classes implementing this interface to read TLS parameters.

Methods

TLSConfiguration Methods List

Signature	Description
public static TLSConfiguration parseTlsConfiguration(final PropertyReader prop, final boolean clientMode)	This is the main and most generic method. It reads all possible TLS parameters (parameter names and possible values are detailed in the list of TLS Parameters in Configuration Manager), converts them and assigns them to `TLSConfiguration` properties.
public static TLSConfiguration parseServerTlsConfiguration(KVList kvl)	These methods provide shortcuts to parse TLS configuration from different source types.
public static TLSConfiguration parseClientTlsConfiguration(KVList kvl)	
public static TLSConfiguration parseServerTlsConfiguration(Map<String, String> map)	
public static TLSConfiguration parseClientTlsConfiguration(Map<String, String> map)	
public static TLSConfiguration parseServerTlsConfiguration(Properties prop)	
public static TLSConfiguration parseClientTlsConfiguration(Properties prop)	

Signature	Description
public static TLSConfiguration parseServerTlsConfiguration(String transportParams)	
public static TLSConfiguration parseClientTlsConfiguration(String transportParams)	

Interface PropertyReader and Implementing Classes

Interface PropertyReader contains just one method:

```
String getProperty(String key)
```

Here, key argument contains name of parameter to extract. Implementing classes contain code that actually extract and return value corresponding to the key. Currently there are five implementations:

1. GConfigTlsPropertyReader - This class belongs to Application Template and is used to extract TLS parameters from a set of related Configuration objects. It cannot be included to Commons library since it would cause circular references between the Commons and Application Template.
2. KVListPropertyReader - Extracts String value from a KVList instance.
3. MapPropertyReader - Extracts value from a Map<String, String> instance.
4. PropertiesReader - Extracts value from a Properties instance.
5. TransportParamsPropertyReader - Parses transport parameters as they appear in Configuration Manager, for example:

```
“tls=1;certificate=c:/cert/cert.pem;mutual=1”.
```

Using the Application Template Application Block

Introduction

Instead of using the Platform SDK Commons Library to configure TLS connections with hard-coded values, you can use the Platform SDK Application Template Application Block to retrieve configuration objects from Configuration Server which contain parameters that are used to configure your TLS settings.

The steps do accomplish this are as follows:

1. Parse a configuration object.
2. Create a `TLSConfiguration` object for the configuration object.
3. Customize your `TLSConfiguration` object:
    - Add callback handlers.
    - For clients, set the expected host names for primary and backup servers.
4. Create `SSLContext` and `SSLExtendedOptions` objects based on your `TLSConfiguration` object.
5. Use your `SSLContext` and `SSLExtendedOptions` objects to create `Endpoints` and/or `WarmStandbyConfiguration` objects.
6. Use your `Endpoints` and/or `WarmStandbyConfiguration` objects to create `Protocol` instances.

The sections below describe these steps in more detail. If you plan on using this method to configure TLS settings, be sure that related application objects in Configuration Manager have been configured with TLS parameters.

Note: If you aren't familiar with TLS configuration settings then please read Using the Platform SDK Commons Library to gain a better understanding of what is required.

Parsing Configuration Objects

The Platform SDK Application Template has a helper class, `GConfigTlsPropertyReader`, which makes it easy to extract TLS parameters from Configuration Server. When used in conjunction with `TLSConfigurationParser`, `TLSConfigurationHelper`, `ClientConfigurationHelper` and `ServerConfigurationHelper` classes, all of the connection-related options found in Configuration Server are covered. They also provide other useful functionality.

`TLSConfigurationParser` has two constructors:

```
public GConfigTlsPropertyReader(
```

```
        IGApplicationConfiguration appConfig,
        IGApplicationConfiguration.IGPortInfo portConfig);
```

and

```
public GConfigTlsPropertyReader(
        IGApplicationConfiguration appConfig,
        IGApplicationConfiguration.IGAppConnConfiguration connConfig);
```

The first one is used for server-side connections while the second is for client-side connections.

For example:

```
// Client side
// Prepare configuration objects
String clientAppName = "<my-app-name>";
CfgAppType targetServerType = CfgAppType.CFGTServer;
CfgApplication cfgApplication = confService.retrieveObject(
        CfgApplication.class, new CfgApplicationQuery(clientAppName));
GCOMApplicationConfiguration appConfiguration =
        new GCOMApplicationConfiguration(cfgApplication);
IGApplicationConfiguration.IGAppConnConfiguration connConfig =
        appConfiguration.getAppServer(targetServerType);

// Parse TLS parameters
PropertyReader reader = new GConfigTlsPropertyReader(appConfiguration, connConfig);
TLSConfiguration tlsConfiguration =
        TLSConfigurationParser.parseTlsConfiguration(reader, true);
// At this point, tlsConfiguration contains TLS parameters read from
// configuration objects

// Server side
// Prepare configuration objects
String serverAppName = "<my-app-name>";
String portID = "secure";
CfgApplication cfgApplication = confService.retrieveObject(
        CfgApplication.class, new CfgApplicationQuery(serverAppName));
GCOMApplicationConfiguration appConfiguration =
        new GCOMApplicationConfiguration(cfgApplication);
IGApplicationConfiguration.IGPortInfo portConfig =
        appConfiguration.getPortInfo(portID);

// Parse TLS parameters
PropertyReader reader = new GConfigTlsPropertyReader(appConfiguration, portConfig);
TLSConfiguration tlsConfiguration =
        TLSConfigurationParser.parseTlsConfiguration(reader, false);
```

Customizing TLS Configuration

When Configuration objects are used as a source of TLS parameters, they can also provide values for expected host names.

Examples:

```
TLSConfiguration tlsConfiguration = ...;

// Client side
// Prepare configuration objects
```

```
String clientAppName = "<my-app-name>";
CfgAppType targetServerType = CfgAppType.CFGTServer;
CfgApplication cfgApplication = confService.retrieveObject(
        CfgApplication.class, new CfgApplicationQuery(clientAppName));
GCOMApplicationConfiguration appConfiguration =
        new GCOMApplicationConfiguration(cfgApplication);
IGApplicationConfiguration.IGAppConnConfiguration connConfig =
        appConfiguration.getAppServer(targetServerType);

// TLS-specific part
IGApplicationConfiguration.IGServerInfo primaryServer =
        connConfig.getTargetServerConfiguration().getServerInfo();
IGApplicationConfiguration.IGServerInfo backupServer =
        primaryServer.getBackup().getServerInfo();

tlsConfiguration.setExpectedHostname(primaryServer.getHost().getName());
// Or:
// tlsConfiguration.setExpectedHostname(backupServer.getHost().getName());
```

Creating SSLContext Objects

`SSLContext` and `SSLExtendedOptions` are created either using `TLSConfigurationHelper` or with `TLSConfiguration` shortcut methods:

Examples:

```
SSLContext sslContext =
        TLSConfigurationHelper.createSslContext(tlsConfiguration);
SSLExtendedOptions sslOptions =
        TLSConfigurationHelper.createSslExtendedOptions(tlsConfiguration);

// The same as above, using shortcut methods:
sslContext = tlsConfiguration.createSslContext();
sslOptions = tlsConfiguration.createSslExtendedOptions();
```

Configuring TLS for Client Connections

Platform SDK has a helper class, `ClientConfigurationHelper`, that makes it easier to prepare connections for client applications. This class has the following methods:

```
public static Endpoint createEndpoint(
        IGApplicationConfiguration appConfig,
        IGAppConnConfiguration connConfig,
        IGApplicationConfiguration targetServerConfig);

public static Endpoint createEndpoint(
        IGApplicationConfiguration appConfig,
        IGAppConnConfiguration connConfig,
        IGApplicationConfiguration targetServerConfig,
        boolean tlsEnabled,
        SSLContext sslContext,
        SSLExtendedOptions sslOptions);

public static WarmStandbyConfiguration createWarmStandbyConfig(
        IGApplicationConfiguration appConfig,
```

```
        IGAppConnConfiguration connConfig);

public static WarmStandbyConfiguration createWarmStandbyConfig(
        IGApplicationConfiguration appConfig,
        IGAppConnConfiguration connConfig,
        boolean primaryTLSEnabled,
        SSLContext primarySSLContext,
        SSLExtendedOptions primarySSLOptions,
        boolean backupTLSEnabled,
        SSLContext backupSSLContext,
        SSLExtendedOptions backupSSLOptions);
```

Two of these methods simply accept TLS-specific parameters and pass them through to the `Endpoint` and `WarmStandbyConfiguration` instances being created. A code sample using the `createEndpoint()` method is shown here:

```
String clientAppName = "<my-app-name>";
CfgAppType targetServerType = CfgAppType.CFGTServer;
CfgApplication cfgApplication = confService.retrieveObject(
        CfgApplication.class, new CfgApplicationQuery(clientAppName));

GCOMApplicationConfiguration appConfiguration =
        new GCOMApplicationConfiguration(cfgApplication);

IGAppConnConfiguration connConfig =
        appConfiguration.getAppServer(targetServerType);

// TLS preparation section follows
PropertyReader reader = new GConfigTlsPropertyReader(appConfiguration, connConfig);
TLSConfiguration tlsConfiguration =
        TLSConfigurationParser.parseTlsConfiguration(reader, true);

// TLS customization code goes here...
// As an example, host name verification is turned on
IGApplicationConfiguration.IGServerInfo targetServer =
        connConfig.getTargetServerConfiguration().getServerInfo();
tlsConfiguration.setExpectedHostname(targetServer.getHost().getName());

// Get TLS configuration objects for connection
SSLContext sslContext = tlsConfiguration.createSslContext();
SSLExtendedOptions sslOptions = tlsConfiguration.createSslExtendedOptions();
boolean tlsEnabled = tlsConfiguration.isTlsEnabled();
// TLS preparation section ends

Endpoint epTSrv = ClientConfigurationHelper.createEndpoint(
        appConfiguration, connConfig,
        connConfig.getTargetServerConfiguration(),
        tlsEnabled, sslContext, sslOptions);

TServerProtocol tsProtocol = new TServerProtocol(epTSrv);
tsProtocol.setClientName(clientName);
tsProtocol.open();
```

Configuring Warm Standby

In cases when the target server has a backup in warm standby mode, configuration requires a little extra effort, as shown in the following code sample.

Note: Configuring TLS for primary and backup servers in Warm Standby mode has some specifics that may not be obvious. Primary and backup servers typically share the same settings. Thus, when a server is selected as a backup for another server (the primary server), Configuration Manager copies settings from the primary server to the backup server to make them the same. This is also true of TLS settings, and the same `TLSConfiguration` object can be used to configure both the primary and backup connections. On the other hand, primary and backup servers usually reside on different hosts. This means that if a hostname check is used, each of these servers must have different `expectedHostname` parameter values. This is not hard to do, as the following code sample demonstrates, but it is not always obvious.

```
String clientAppName = "<my-app-name>";
CfgAppType targetServerType = CfgAppType.CFGStatServer;
CfgApplication cfgApplication = confService.retrieveObject(
        CfgApplication.class, new CfgApplicationQuery(appName));

GCOMApplicationConfiguration appConfiguration =
        new GCOMApplicationConfiguration(cfgApplication);

IGAppConnConfiguration connConfig =
        appConfiguration.getAppServer(targetServerType);

// TLS preparation section follows
PropertyReader reader = new GConfigTlsPropertyReader(appConfiguration, connConfig);
TLSConfiguration tlsConfiguration =
        TLSConfigurationParser.parseTlsConfiguration(reader, true);

IGApplicationConfiguration.IGServerInfo primaryServer =
        connConfig.getTargetServerConfiguration().getServerInfo();
IGApplicationConfiguration.IGServerInfo backupServer =
        primaryServer.getBackup().getServerInfo();

// Configure TLS for Primary
tlsConfiguration.setExpectedHostname(primaryServer.getHost().getName());
SSLContext primarySslContext = tlsConfiguration.createSslContext();
SSLExtendedOptions primarySslOptions = tlsConfiguration.createSslExtendedOptions();
boolean primaryTlsEnabled = tlsConfiguration.isTlsEnabled();

// Configure TLS for Backup
tlsConfiguration.setExpectedHostname(backupServer.getHost().getName());
SSLContext backupSslContext = tlsConfiguration.createSslContext();
SSLExtendedOptions backupSslOptions = tlsConfiguration.createSslExtendedOptions();
boolean backupTlsEnabled = tlsConfiguration.isTlsEnabled();
// TLS preparation section ends

WarmStandbyConfiguration wsConfig =
        ClientConfigurationHelper.createWarmStandbyConfig(
                appConfiguration, connConfig,
                primaryTlsEnabled, primarySslContext, primarySslOptions,
                backupTlsEnabled, backupSslContext, backupSslOptions);

StatServerProtocol statProtocol =
        new StatServerProtocol(wsConfig.getActiveEndpoint());
statProtocol.setClientName(clientName);

WarmStandbyService wsService = new WarmStandbyService(statProtocol);
wsService.applyConfiguration(wsConfig);
wsService.start();
statProtocol.beginOpen();
```

Configuring TLS for Servers

Platform SDK has a helper class, `ServerConfigurationHelper`, that makes it easier to prepare listening sockets for server applications. This class has the following methods:

```
public static Endpoint createListeningEndpoint(
        IGApplicationConfiguration application,
        IGApplicationConfiguration.IGPortInfo portInfo);

public static Endpoint createListeningEndpoint(
        IGApplicationConfiguration application,
        IGApplicationConfiguration.IGPortInfo portInfo,
        boolean tlsEnabled,
        SSLContext sslContext,
        SSLExtendedOptions sslOptions);
```

The overloaded version of the `createListeningEndpoint()` method accepts TLS parameters and passes them through to the `Endpoint` object that is being created. The following code sample shows how this is done:

```
String serverAppName = "<my-app-name>";
String portID = "secure";
CfgApplication cfgApplication = confService.retrieveObject(
        CfgApplication.class, new CfgApplicationQuery(appName));
GCOMApplicationConfiguration appConfig =
        new GCOMApplicationConfiguration(cfgApplication);
IGApplicationConfiguration.IGPortInfo portConfig =
        appConfig.getPortInfo(portID);

// TLS preparation section follows
PropertyReader reader = new GConfigTlsPropertyReader(appConfiguration, portConfig);
TLSConfiguration tlsConfiguration =
        TLSConfigurationParser.parseTlsConfiguration(reader, false);

// TLS customization code goes here...
// As an example, mutual TLS mode is turned on
tlsConfiguration.setMutual(true);

// Get TLS configuration objects for connection
SSLContext sslContext = tlsConfiguration.createSslContext();
SSLExtendedOptions sslOptions = tlsConfiguration.createSslExtendedOptions();
boolean tlsEnabled = tlsConfiguration.isTlsEnabled();
// TLS preparation section ends

Endpoint endpoint = ServerConfigurationHelper.createListeningEndpoint(
        appConfig, portConfig,
        tlsEnabled, sslContext, sslOptions);
ExternalServiceProtocolListener serverChannel =
        new ExternalServiceProtocolListener(endpoint);
...
```

Configuring TLS Parameters in Configuration Manager

Introduction

As described earlier, the Platform SDK Application Template Application Block allows both client and server applications to read TLS parameters from configuration objects. This page describes how to set TLS parameters correctly in those configuration objects.

Configuration objects that will be used, and their relations, are shown in the diagram below:

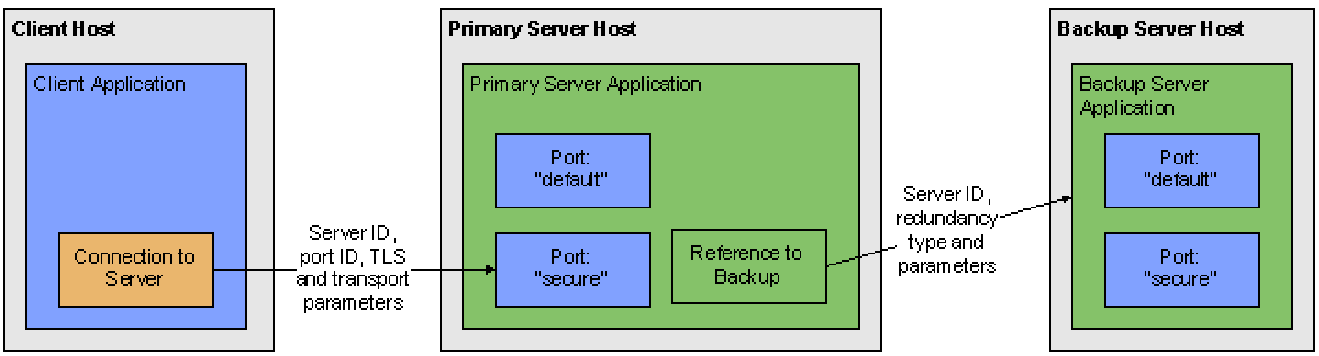

To edit TLS-related parameters for these objects, you will need to have access to the Annex tab in Configuration Manager.

Precedence of Configuration Objects

Platform SDK uses different sets of configuration objects to configure client- and server-side TLS settings. For TLS parameters, these objects are searched from the most specific object to the most general one. Parameters found in specific objects take precedence over those in more general objects.

Note: This search occurs independently for each supported TLS parameter.

Location of specific TLS parameters can differ for each object, but is detailed in the appropriate section on this page.

Configuration Object Precedence

Application type	Configuration Objects Used, in Order of Precedence
Client	1. Connection from the client application to the server.

Application type	Configuration Objects Used, in Order of Precedence
	2. Application of the client. 3. Host where client application resides. 4. Port of the target server that client connects to.[1]
Server	1. Port of the server application. 2. Application of the server. 3. Host where the server application resides.

1. If the `tls` parameter is not set to 1 in both the client Application and Connection objects, then the client application will look to the Port object for the target server to determine if TLS should be turned on. Configuration Manager does not automatically add the `tls=1` parameter to Connection Transport parameters when it is linked to a server's secure Port. This is the only case when a client application considers settings in the server's configuration objects.

Displaying the Annex Tab in Configuration Manager

By default, Configuration Manager does not show Annex tab in Object Properties windows. This tab can contain TLS parameters for Host and Application objects.

To show the Annex tab, select *View > Options...* from the main menu and ensure the *Show Annex tab in object properties* and *Show Advanced Security Information* options are selected.

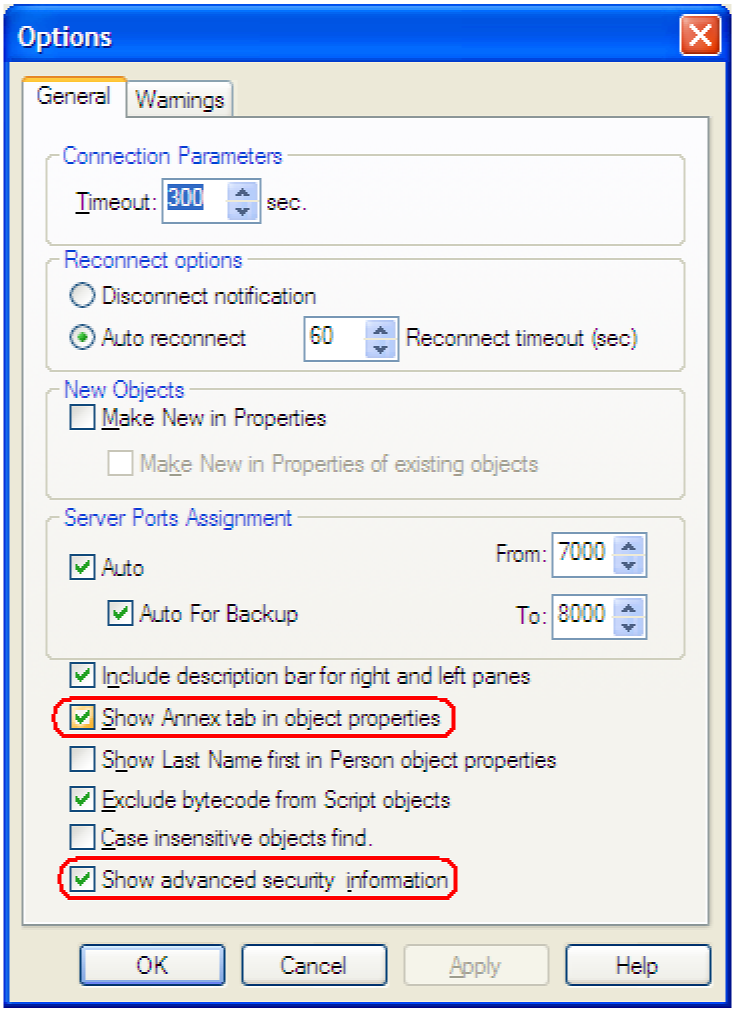

Application Objects

Host Object

The properties window for a Host object includes most common TLS parameters on the General tab:

- Certificate
- Certificate Key
- Trusted CA

These fields allow copy/paste operations, so they can be set manually by copying and pasting the "Thumbprint" field values from certificates in Windows Certificate Services (WCS) into the related field in Configuration Manager.

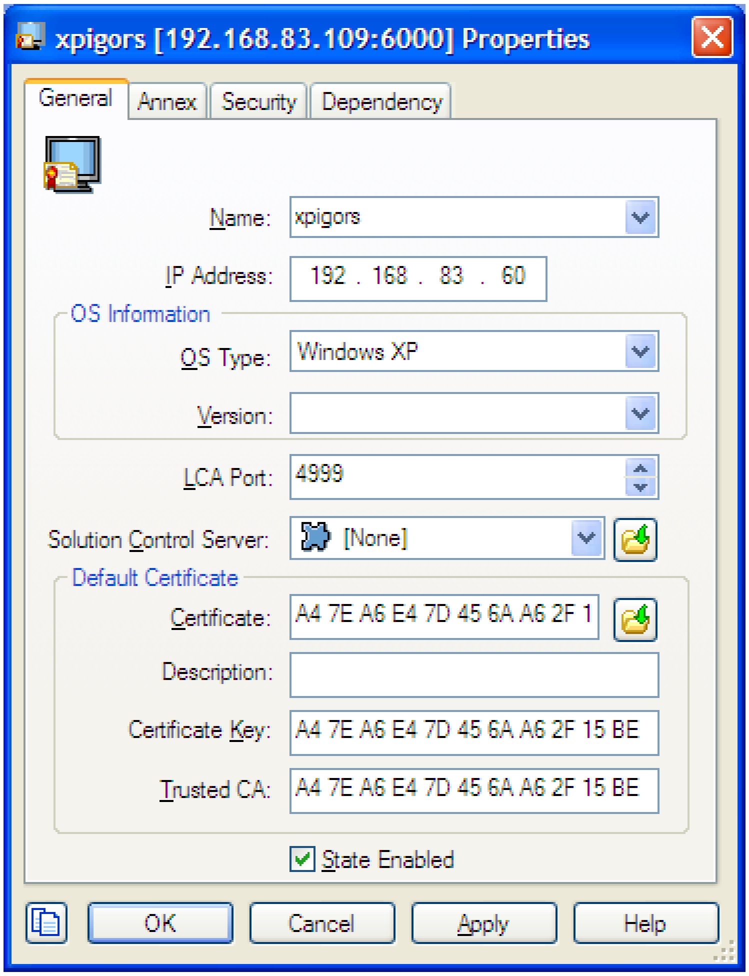

To select a certificate, use the button next to *Certificate* field. This opens the *Select certificate* window, displaying a list of certificates installed in WCS under the Local Computer account for the local machine.

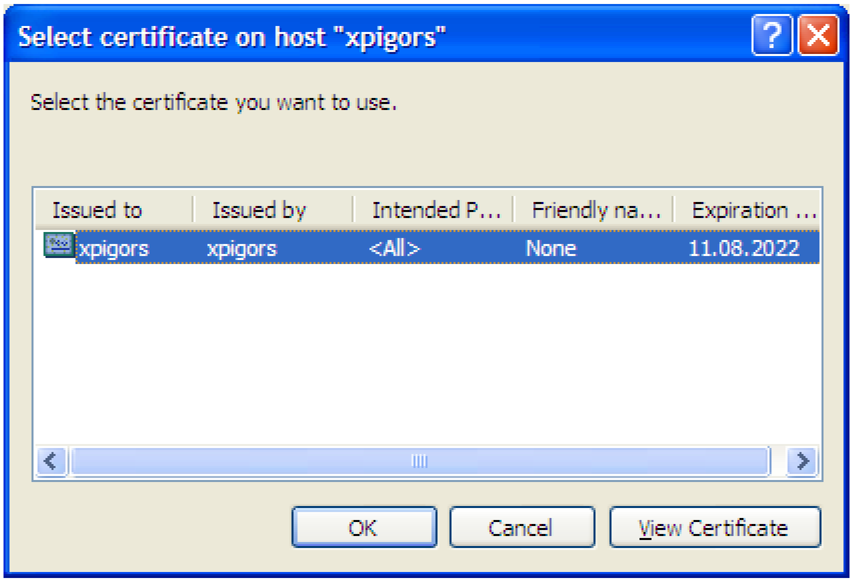

The Annex tab contains a security section that holds TLS settings for this object. Any change made to TLS-related fields on the General tab are mirrored between the Annex tab automatically. You can also specify additional TLS parameters here that aren't reflected on the General tab.

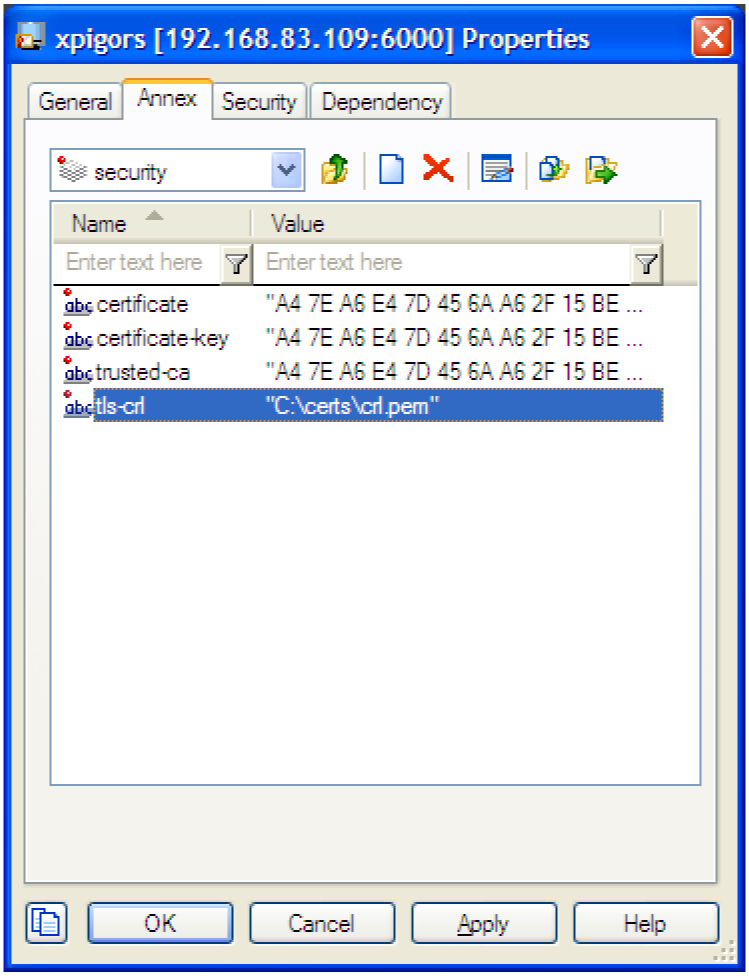

Server Application Object

For the server Application object, TLS-related fields are located on the *Server Info* tab of the properties window. Note the *Certificate View* controls group, where the server can be set to use Host TLS parameters (generally recommended for Genesys Framework) or application-specific ones.

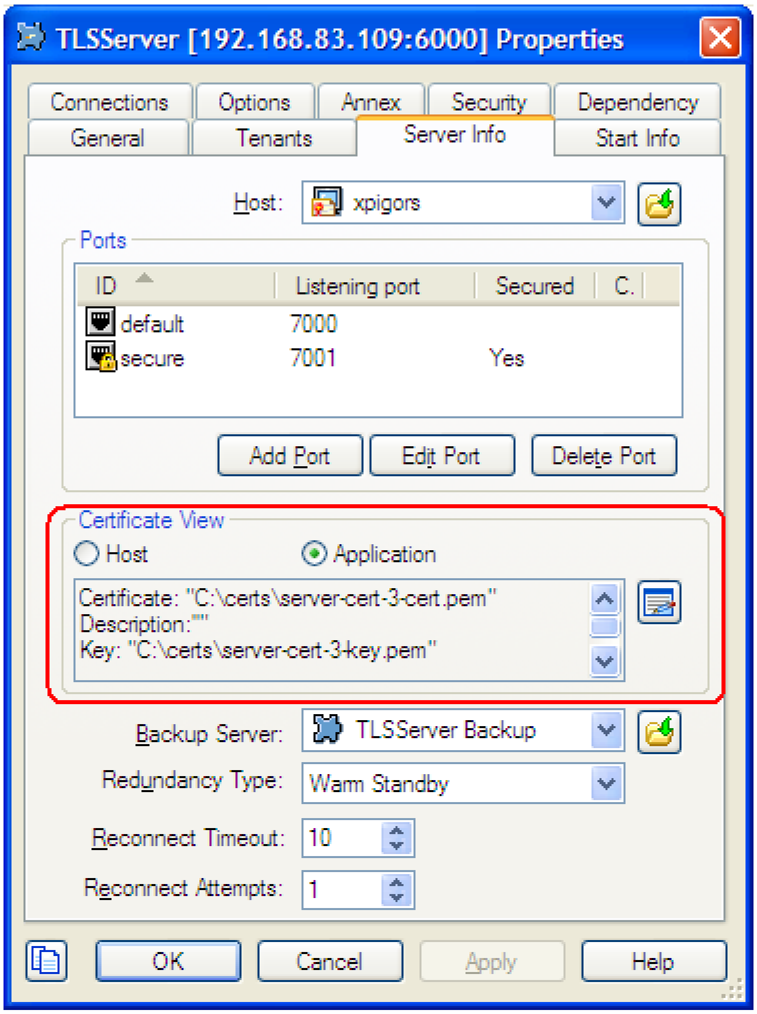

If using application-specific TLS parameters, use the button next to the certificate information field to open a certificate selection window where you can choose from a list of certificates installed for the Local Computer account or manually enter certificate information:

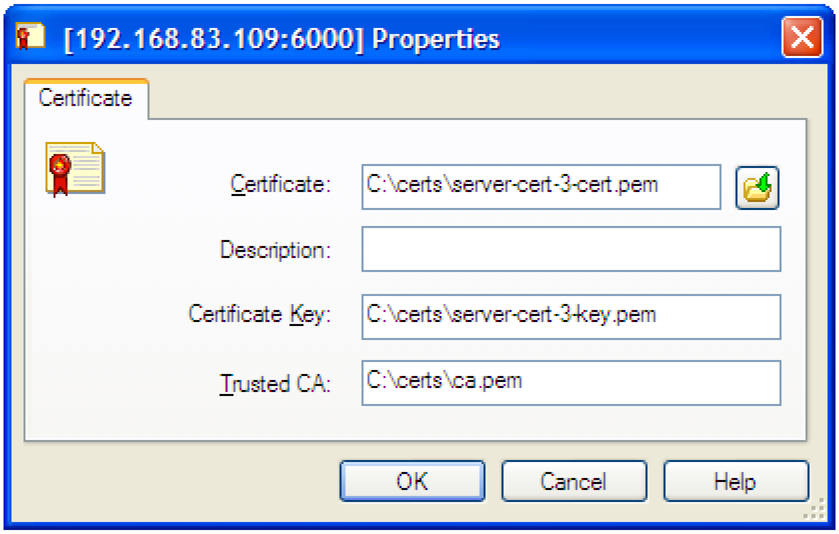

Port Object

For port objects, TLS-related fields are located on the *Server Info* tab of the properties window. You can see here whether a port is secured (TLS-enabled) or not, and have the option to edit existing ports to update TLS parameters or to add new ports.

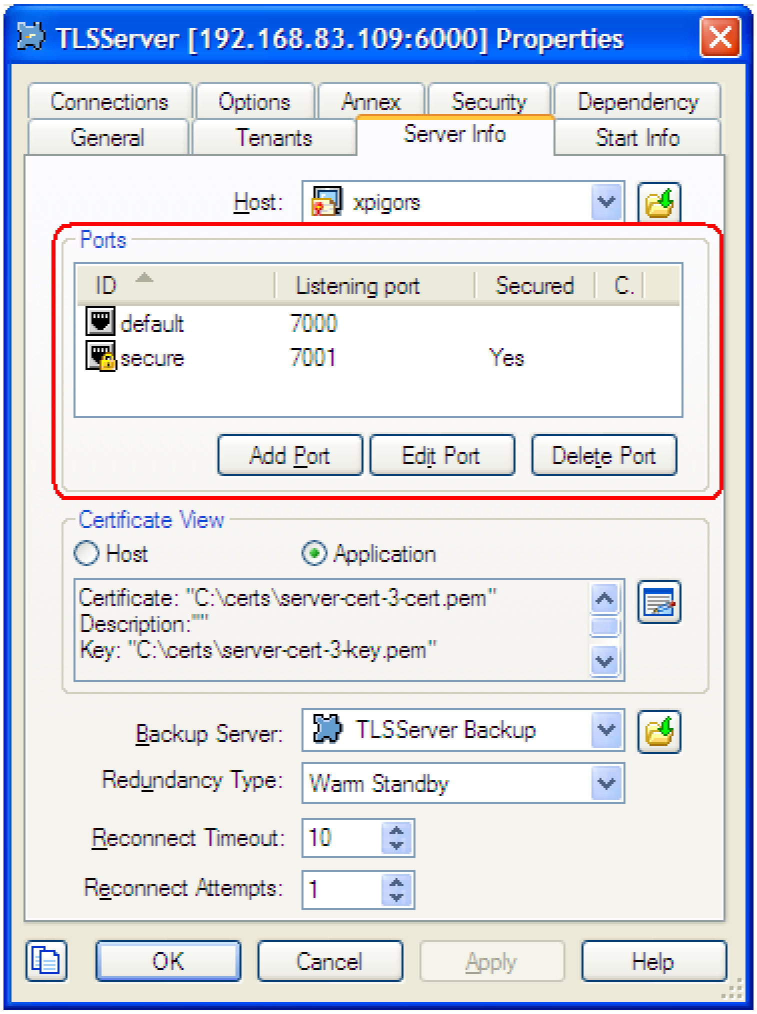

When adding or editing a port, TLS parameters are specified on the following tabs:

- *Port Info* — Turn on *Secured* listening mode for the port (the same as adding the *tls=1* string to transport parameters).
- *Certificate* — Show certificate information, open a certificate selection window, or delete the current certificate information.
- *Advanced* — Manually edit the *Transport Protocol Parameters* field. TLS parameters not reflected on the *Certificate* tab can be added here.

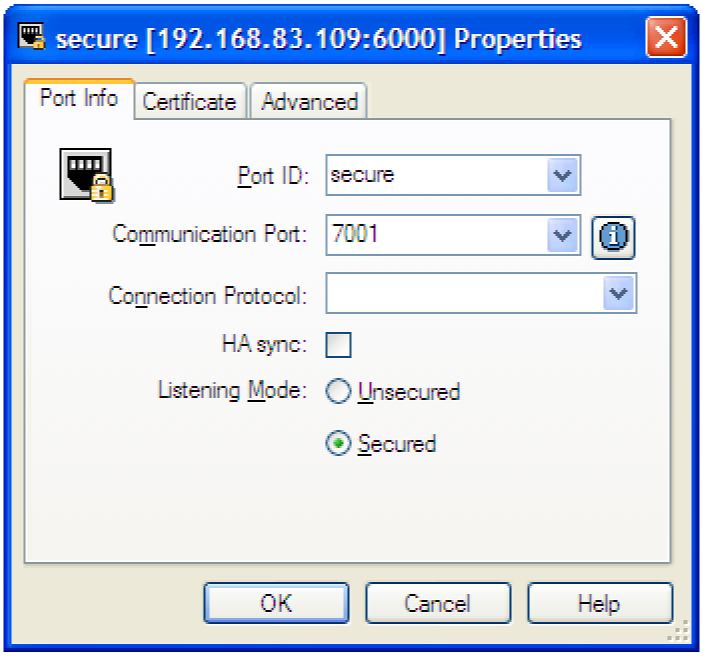
secure [192.168.83.109:6000] Properties
Port Info
Certificate
Advanced
Port ID:
secure
Communication Port:
7001
Connection Protocol:
HA sync:
Listening Mode:
Unsecured
Secured
OK
Cancel
Help

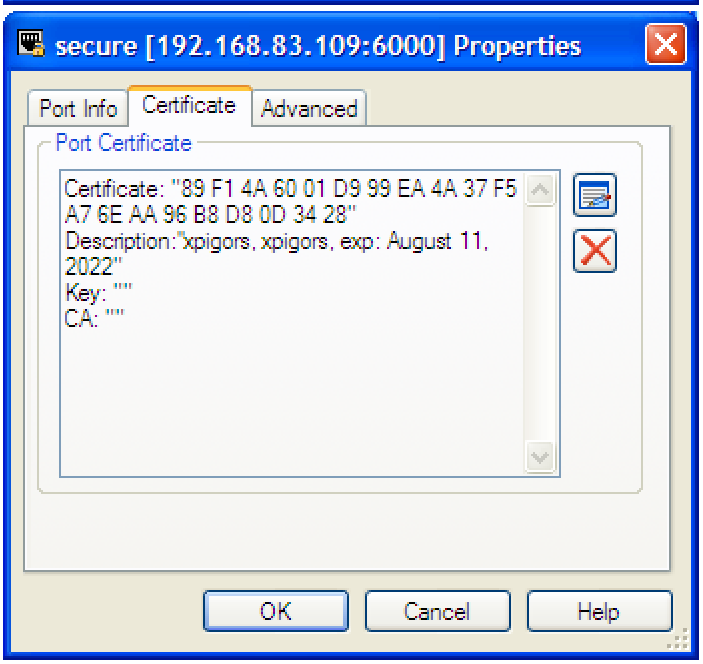
secure [192.168.83.109:6000] Properties
Port Info
Certificate
Advanced
Port Certificate
Certificate: "89 F1 4A 60 01 D9 99 EA 4A 37 F5 A7 6E AA 96 B8 D8 0D 34 28"
Description:"xpigors, xpigors, exp: August 11, 2022"
Key: ""
CA: ""
OK
Cancel
Help

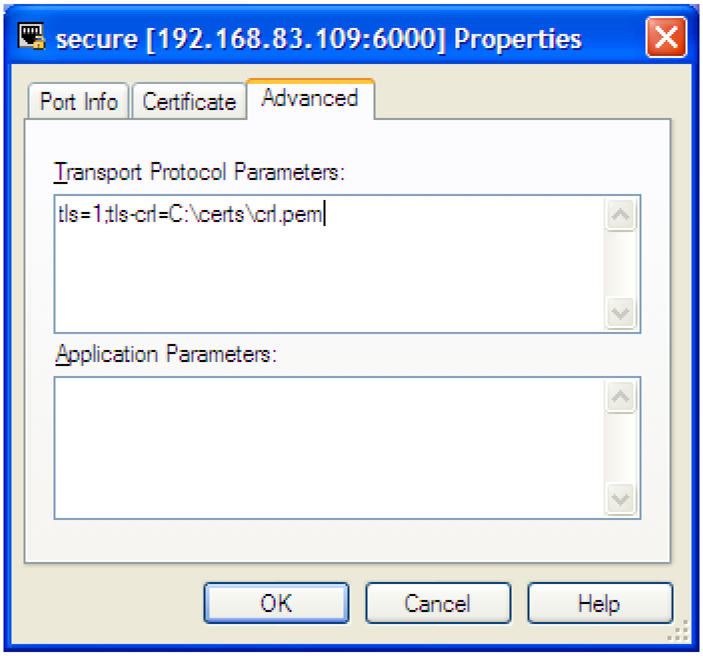

Client Application Object

For client Application objects, TLS-related fields are located under the *security* sections of both the *Options* and *Annex* tabs. There is no certificate selection window provided, but TLS parameters can be configured manually in either section.

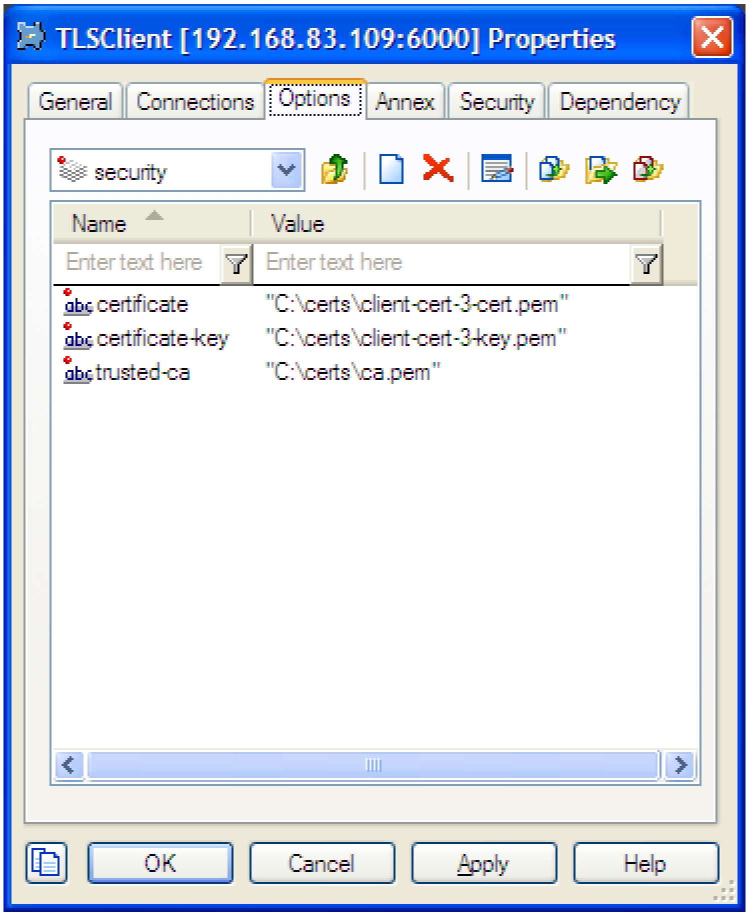
TLSClient [192.168.83.109:6000] Properties
General
Connections
Options
Annex
Security
Dependency
security
Name
Value
Enter text here
Enter text here
certificate
"C:\certs\client-cert-3-cert.pem"
certificate-key
"C:\certs\client-cert-3-key.pem"
trusted-ca
"C:\certs\ca.pem"
OK
Cancel
Apply
Help

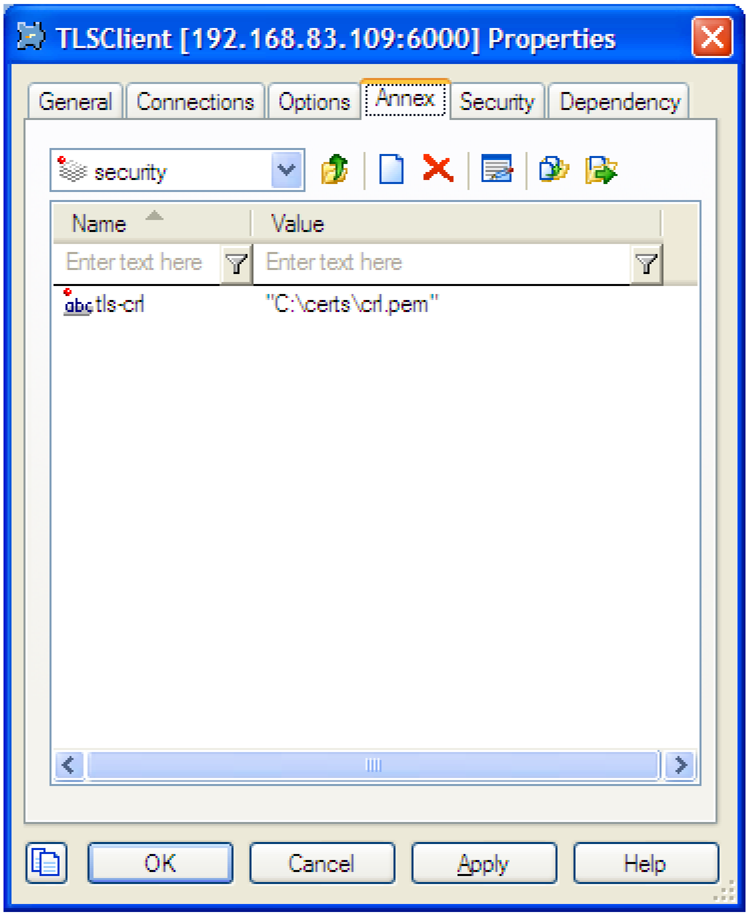

When processing a client Application object, Platform SDK looks at parameters from both sections. If any parameters are specified in both places, then the values from the *Options* tab take precedence.

Connection Object

The properties window for all Application objects includes a *Connection* tab where connections to servers can be added or edited. Each connection determines if TLS mode should be enabled based on port settings for the target server.

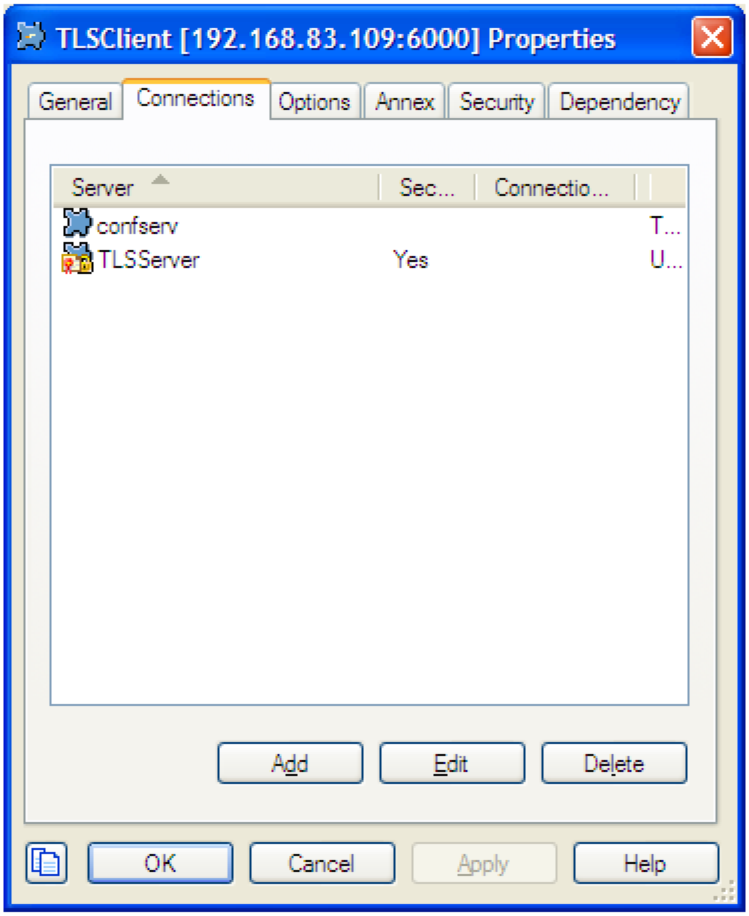
TLSClient [192.168.83.109:6000] Properties
General
Connections
Options
Annex
Security
Dependency
Server
Sec...
Connectio...
confserv
T...
TLSServer
Yes
U...
Add
Edit
Delete
OK
Cancel
Apply
Help

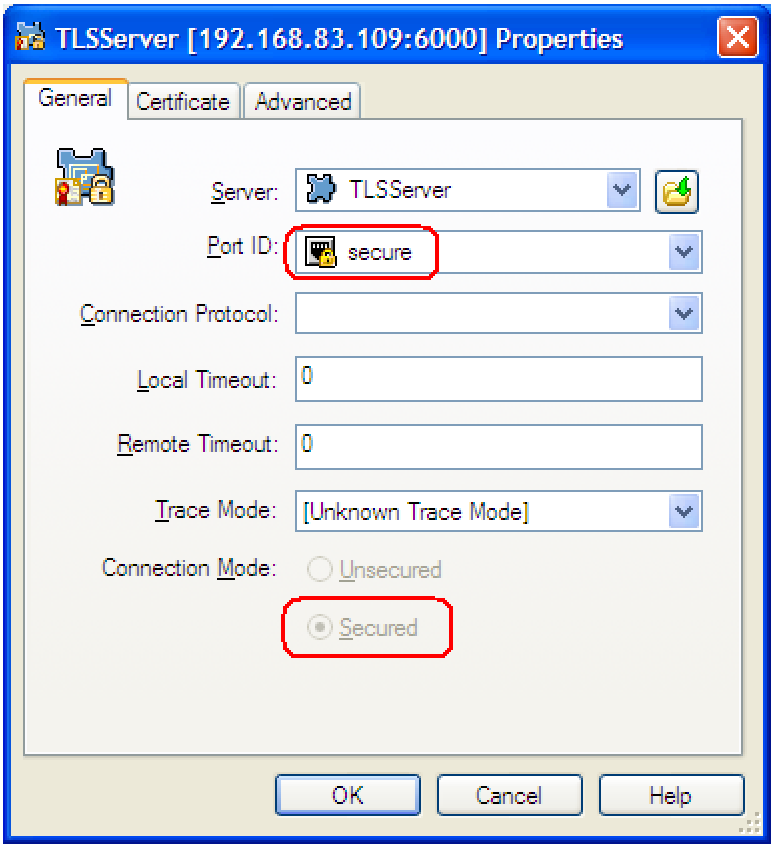

Similar to the *Port* properties window, the *Certificate* tab allows you to select from a list of certificates or manually edit certificate properties. You can also use the *Advanced* tab to edit TLS settings not included with the certificate. However, the *Transport Protocol Parameters* field behaves differently for this object — which may result in lost or incorrect settings in some cases. See the Notes and Issues section for details.

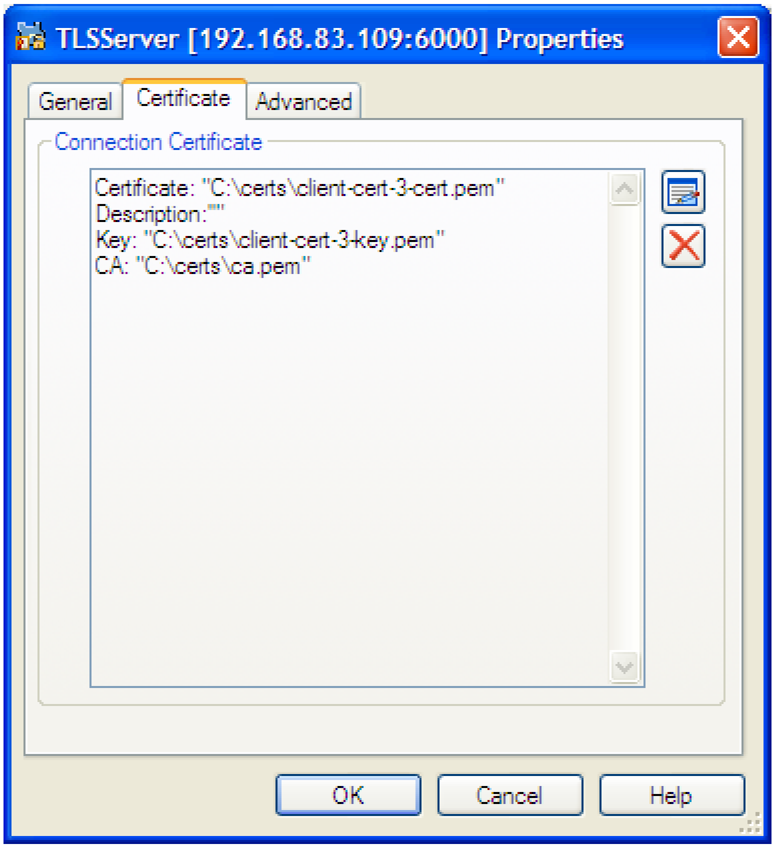
TLSServer [192.168.83.109:6000] Properties
General
Certificate
Advanced
Connection Certificate
Certificate: "C:\certs\client-cert-3-cert.pem"
Description:""
Key: "C:\certs\client-cert-3-key.pem"
CA: "C:\certs\ca.pem"
OK
Cancel
Help

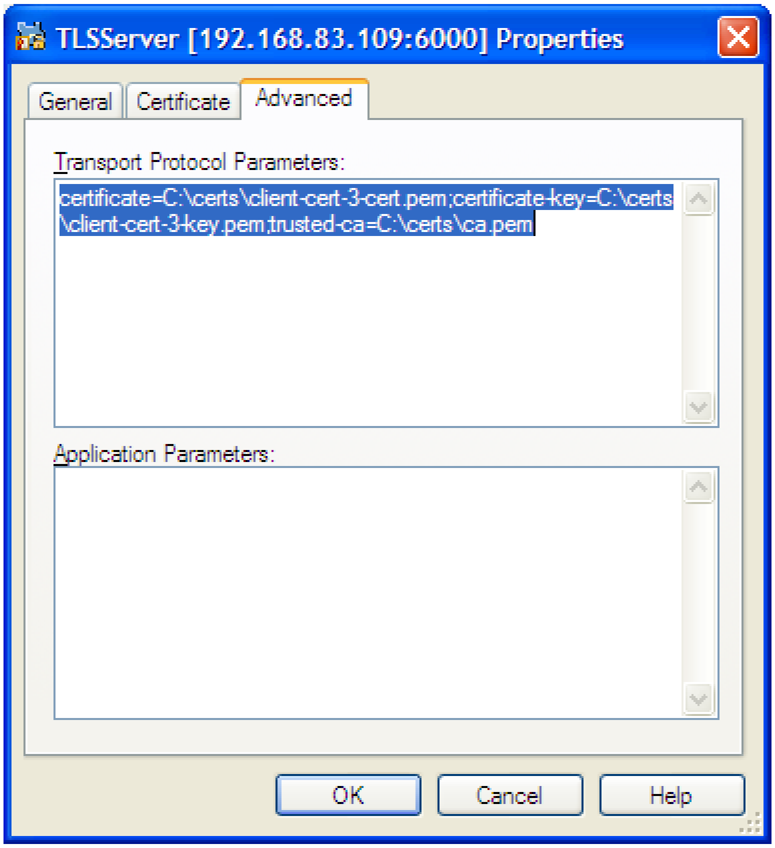

List of TLS Parameters

The following table lists all TLS parameters supported by Platform SDK, with their valid value ranges and purpose:

Parameter Name	Acceptable Values	Purpose
tls	Boolean value. Possible values are "1"/"0", "yes"/"no", "on"/"off", "true"/"false". Example: • "tls=1"	Client: 1 - perform TLS handshake immediately after connecting to server. 0 – do not turn on TLS immediately but autodetect can still work.
provider	"PEM", "MSCAPI", "PKCS11" Not case-sensitive. Example: • "provider=MSCAPI"	Explicit selection of security provider to be used. For example, MSCAPI and PKCS11 providers can contain all other parameters in their internal database. This parameter allow configuration of TLS through security provider tools.

Parameter Name	Acceptable Values	Purpose
certificate	PEM provider: path to a X.509 certificate file in PEM format. Path can use both forward and backward slash characters. MSCAPI provider: thumbprint of a certificate – string with hexadecimal SHA-1 hash code of the certificate. Whitespace characters are allowed anywhere within the string. PKCS11 provider: this parameter is ignored. Examples: • "certificate= C:\certs\client-cert-3-cert.pem" • "certificate=A4 7E A6 E4 7D 45 6A A6 2F 15 BE 89 FD 46 F0 EE 82 1A 58 B9"	Specifies location of X.509 certificate to be used by application. MSCAPI provider keeps certificates in internal database and can identify them by hash code; so called thumbprint. In Java, PKCS#11 provider does not allow selection of the certificate; it must be configured using provider tools. **Note:** When using autodetect (upgrade) TLS connection, this option MUST be specified in application configuration, otherwise Configuration Server would return empty TLS parameters even if other options are set.
certificate-key	PEM provider: path to a PKCS#8 private key file without password protection in PEM format. Path can use both forward and backward slash characters. • MSCAPI provider: this parameter is ignored; key is taken from the entry identified by "certificate" field. • PKCS11 provider: this parameter is ignored. Examples: • "certificate-key= C:\certs\client-cert-3-key.pem"	Specifies location of PKCS#8 private key to be used in pair with the certificate by application. MSCAPI provider keeps private keys paired with certificates in internal database. In Java, PKCS#11 provider does not allow selection of the private key; it must be configured using provider tools.
trusted-ca	PEM provider: path to a X.509 certificate file in PEM format. Path can use both forward and backward slash characters. MSCAPI provider: thumbprint of a certificate – string with hexadecimal SHA-1 hash code of the certificate. Whitespace characters are allowed anywhere within the string. PKCS11 provider: this parameter is ignored. Examples: • "trusted-ca= C:\certs\ca.pem" • "trusted-ca=A4 7E A6 E4 7D	Specifies location of a X.509 certificate to be used by application to validate remote party certificates. The certificate is designated as Trusted Certification Authority certificate and application will only trust remote party certificates signed with the CA certificate. MSCAPI provider keeps CA certificates in internal database and can identify them by hash code; so called thumbprint. In Java, PKCS#11 provider does not allow selection of the CA certificate; it must be configured using provider tools.

Parameter Name	Acceptable Values	Purpose
	45 6A A6 2F 15 BE 89 FD 46 F0 EE 82 1A 58 B9"	
tls-mutual	Boolean value. Possible values are "1"/"0", "yes"/"no", "on"/"off", "true"/"false". Example: • "tls-mutual=1"	Has meaning only for server application. Client applications ignore this value. When turned on, server will require connecting clients to present their certificates and validate the certificates the same way as client applications do.
tls-crl	All providers: path to a Certificate Revocation List file in PEM format. Path can use both forward and backward slash characters. Example: • "tls-crl= C:\certs\crl.pem"	Applications will use CRL during certificate validation process to check if the (seemingly valid) certificate was revoked by CA. This option is useful to stop usage of leaked certificates by unauthorized parties.
tls-target-name-check	"host" or none. Not case-sensitive. Example: • "tls-target-name-check=host"	When set to "host", enables matching of certificate’s Alternative Subject Name or Subject fields against expected host name. PSDK supports DNS names and IP addresses as expected host names.
cipher-list	String consisting of space-separated cipher suit names. Information on cipher names can be found online. Example: • "cipher-list= TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA TLS_ECDH_RSA_WITH_3DES_EDE_CBC_SHA"	Used to calculate enabled cipher suites. Only ciphers present in both the cipher suites supported by security provider and the cipher-list parameter will be valid.
fips140-enabled	Boolean value. Possible values are "1"/"0", "yes"/"no", "on"/"off", "true"/"false". Example: • "fips140-enabled=1"	PSDK Java: when set to true, effectively is the same as setting "provider=PKCS11" since only PKCS11 provider can support FIPS-140. If set to true while using other provider type, PSDK will throw exception.

Notes and Issues

- Key/value pairs in *Transport Protocol Parameters* fields should be separated only with a single semicolon character. Adding space characters to improve readability can cause applications, including those based on Platform SDK, unable to parse these parameters correctly.
- *Transport Protocol Parameters* fields in Configuration Manager are limited to 256 characters in length. Be sure to keep your parameter list as short as possible. For example: certificate thumbprints for MSCAPI provider take 40 characters without spaces and 49 characters with them, and long paths to certificate files can easily eat up all available space.
- The Connection properties window behaves differently from the Port properties window, as described below. Be sure to double-check TLS settings for Connection objects.
    - It does not save content of the *Transport Protocol Parameters* field unless a certificate was selected using UI controls on the *Certificate* tab.
    - If certificate information is deleted from the *Certificate* tab, then all transport protocol parameters are also erased (including those entered manually).
    - In some cases it does not save additional TLS parameters that were entered manually.
- Configuration Server reads its own TLS parameters from Application or from Host object only during startup. If you use an Application or Host object as a source of TLS parameters for Configuration Server, be sure to restart the server after any changes to the parameters.

Using and Configuring Security Providers

Introduction

This page deals with Security Providers — an umbrella term describing the full set of cryptographic algorithms, data formats, protocols, interfaces, and related tools for configuration and management when used together. The primary reasons for bundling together such diverse tools are: compatibility, support for specific standards, and implementation restrictions.

The security providers listed here were tested with the Platform SDK 8.1.1 implementation of TLS, and found to work reliably when used with the configuration described below.

Java Cryptography Architecture Notes

Java Cryptography Architecture (JCA) provides a general API, and a pluggable architecture for cryptography providers that supply the API implementation.

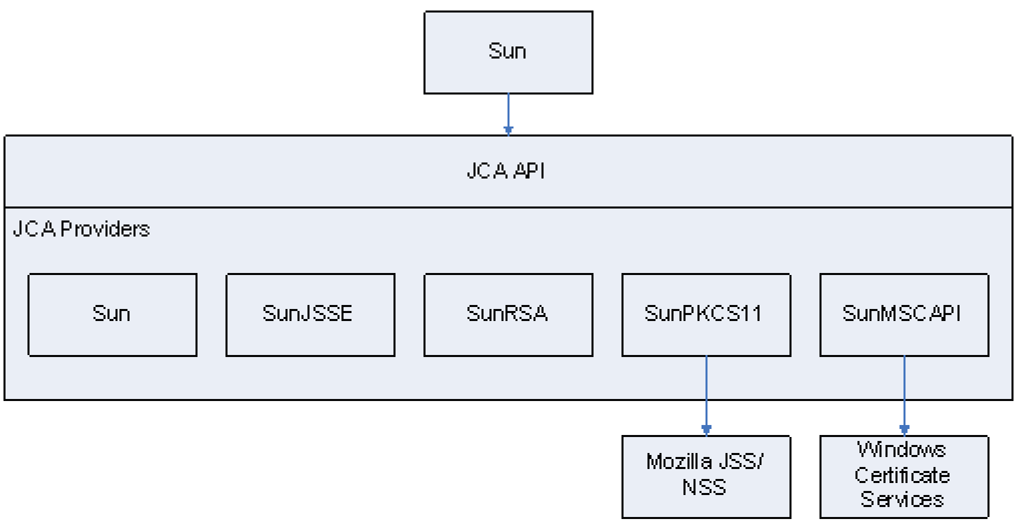

Some JCA providers (Sun, SunJSSE, SunRSA) come bundled with the Java platform and contain actual algorithm implementations, they are named PEM provider since they are used when working with certificates in PEM files. Some other (SunPKCS11, SunMSCAPI) serve as a façade for external providers. SunPKCS11 supports PKCS#11 standard for pluggable security providers, such as hardware cryptographic processors, smartcards or software tokens. Mozilla NSS/JSS is an example of pluggable software token implementation. SunMSCAPI provides access to Microsoft Cryptography API (MSCAPI), in particular, to Windows Certificate Services (WSC).

PEM Provider: OpenSSL

Note: Working with certificates and keys is also covered in the *Genesys 8.1 Security Deployment Guide*.

PEM stands for "Privacy Enhanced Mail", a 1993 IETF proposal for securing e-mail using public-key cryptography. That proposal defined the PEM file format for certificates as one containing a Base64-encoded X.509 certificate in specific binary representation with additional metadata headers. Here, the term is used to refer to Java built-in security providers that are used in conjunction with certificates and private keys loaded from X.509 PEM files.

One of the most popular free tools for creating and manipulating PEM files is OpenSSL. Instructions for installing and configuring OpenSSL are provided below.

Installing OpenSSL

OpenSSL is available two ways:

- distributed as a source code tarball: http://www.openssl.org/source/
- as a binary distribution (specific links are subject to change): http://www.openssl.org/related/binaries.html

The installation process is very easy when using a binary installer; simply follow the prompts. The only additional step required is to add the *<OpenSSL-home>\bin* folder to your *Path* system variable so that OpenSSL can run from command line directly with the `openssl` command.

Configuring OpenSSL

The OpenSSL configuration file contains settings for OpenSSL itself, and also many field values for the certificates being generated including issuer and subject names, host names and URIs, and so on. You will need to customize your OpenSSL file with your own values before using the tool. An example of a customized configuration file is available here.

The OpenSSL database consists of a set of files and folders, similar to the sample database described in the table below. To start using OpenSSL, this structure should be created manually except for files marked as "Generated by OpenSSL". Other files can be left empty as long as they exist in the expected location.

OpenSSL database file/folder structure

File or Folder	Generated by OpenSSL?	Description
openssl-ca\		
openssl-ca\openssl.cfg		OpenSSL configuration file
openssl-ca\.rnd	Yes	File filled with random data, used in key generation process.
openssl-ca\ca-password.txt		Stores the password for the CA private key. Reduces typing required, but is very insecure. Should only be used for testing and development.

File or Folder	Generated by OpenSSL?	Description
openssl-ca\export-password.txt		Stores the password used to encrypt the private keys when exporting PKCS#12 files. Reduces typing required, but is very insecure. Should only be used for testing and development.
openssl-ca\ca\		CA root folder.
openssl-ca\ca\certs\		All generated certificates are copied here. Folder contents can be safely deleted.
openssl-ca\ca\crl\		Generated CRLs stored here. Folder contents can be safely deleted.
openssl-ca\ca\newcerts\		Certificates being generated are stored here. Folder contents can be safely deleted *once generation process is finished*.
openssl-ca\ca\private\		CA private files.
openssl-ca\ca\private\cakey.pem	Yes	CA private key. Must be kept secret.
openssl-ca\ca\crlnumber		Serial number of last exported CRL.
openssl-ca\ca\serial		Serial number of last signed certificate.
openssl-ca\ca\cacert.pem	Yes	CA certificate.
openssl-ca\ca\index.txt		Textual database of all certificates.

Short Command Line Reference

- This section assumes that the OpenSSL *bin* folder was added to the local PATH environment variable, and that *openssl-ca* is the current folder for all issued commands.
- Placeholders for parameters are shown in the following form: "<param-placeholder>".
- The frequently used parameter "<request-name>" should be a unique name that identifies the certificate files.

Task	Description	Command
Create a CA Certificate/Key	This is performed in three steps: 1. Create CA Private Key 2. Create CA Certificate	1. `openssl genrsa -des3 -out ca\private\cakey.pem 1024 -passin file:ca-password.txt`

Task	Description	Command
	3. Export CA Certificate	2. `openssl req -config openssl.cfg -new -x509 -days <days-ca-cert-is-valid> -key ca\private\cakey.pem -out ca\cacert.pem -passin file:ca-password.txt` 3. `openssl x509 -in ca\cacert.pem -outform PEM -out ca.pem`

Create a Leaf Certificate/Key Pair	This is performed in three steps:
	1. Create certificate request. Certificate fields and extensions are defined during this step, and the certificate's public and private keys are created in the process.
	2. Sign the request.
	3. Export the certificate.
		1. `openssl req -new -nodes -out requests\<request name>-req.pem -keyout requests\<request name>-key.pem -days 3650 -config openssl.cfg`
		2. `openssl ca -out requests\<request-name>-signed.pem -days 3650 -config openssl.cfg -passin file:ca-password.txt -infiles requests\<request-name>-req.pem`
		3. `openssl pkcs12 -export -in requests\<request-name>-signed.pem -inkey requests\<request-name>-key.pem -certfile ca\cacert.pem -name "<entry-name-in-p12-file>" -out <request-name>.p12 -passout file:export-password.txt`
		`openssl x509 -in requests\<request-name>-signed.pem -outform PEM -out <request-name>-cert.pem`
		`openssl pkcs8 -topk8 -nocrypt -in requests\<request-name>-key.pem -out <request-name>-key.pem`

Revoke a Certificate		`openssl ca -revoke <certificate-pem-file> -config openssl.cfg -passin file:ca-password.txt`
Export the CRL		`openssl ca -gencrl -crldays`

Task	Description	Command
		`<days-crl-is-valid> -out crl.pem -config openssl.cfg -passin file:ca-password.txt`

MSCAPI Provider: Windows Certificate Services

Note: Working with Windows Certificate Services (WCS) is also covered in *Genesys 8.1 Security Deployment Guide*.

MSCAPI stands for Microsoft CryptoAPI. This provider offers the following features:

- It is available only on Windows platform.
- It implies usage of WCS to store and retrieve certificates, private keys, and CA certificates.
- Every Windows account has its own WCS storage, including the System account.
- Depends heavily on OS configuration and system security policies.
- Has its own set of supported cipher suites, different from what is provided by Java.
- When used with Java, please use the latest available version of Java to run the application. The minimum required version for correct MSCAPI support is Java 6 update 38, with additional compatibility details outlined below:
    - Java 5 and lower versions—MSCAPI is not supported.
    - Java 6 32-bit version—MSCAPI provider is only supported since update 27: http://bugs.sun.com/bugdatabase/view_bug.do?bug_id=6931562.
    - Java 6 64-bit version—MSCAPI provider is only supported since update 38: http://bugs.sun.com/bugdatabase/view_bug.do?bug_id=2215540.
    - Java 7—MSCAPI is supported in all versions.
- Java does not support CRLs located in WCS. With Java MSCAPI, CRL should be specified as a file.
- Does not accept passwords from Java code programmatically via CallbackHandler. If private key is password-protected or prompt-protected, OS popup dialog will be shown to user.
- Certificates in WCS are configured using the Certificates snap-in for Microsoft Management Console (MMC).

Note: If the version of Java being used does not support MSCAPI, a "WINDOWS-MY KeyStore not available" exception appears in the application log. If you receive such exceptions, please consider switching to a newer version of Java.

Starting Certificates Snap-in

There are two methods for accessing the Certificates Snap-in:

- Enter "certmgr.msc" at the command line. (This only gives access to Certificates for the current user account.)

- Launch the MMC console and add the Certificates Snap-in for a specific account using the following steps:
    1. Enter "mmc" at the command line.
    2. Select *File > Add/Remove Snap-in...* from the main menu.
    3. Select *Certificates* from the list of available snap-ins and click *Add*.
    4. Select the account to manage certificates for (see Account Selection for important notes) and click *Finish*.
    5. Click *OK*.

Account Selection

It is important to place certificates under the correct Windows account. Some applications are run as services under the Local Service or System account, while others are run under user accounts. The account chosen in MMC must be the same as the account used by the application that certificates are configured for, otherwise the application will not be able to access this WCS storage.

Note: Currently, most Genesys servers do not clearly report this error so WCS configuration must be checked every time there is a problem with the MSCAPI provider.

Note: Configuration Manager is also a regular application in this aspect and can access WCS only for the Local Computer (System) account on the local machine. It will not show certificates configured for different accounts or on remote machines. Please consult your system and/or security administrator for questions related to certificate configuration and usage.

Importing Certificates

There are many folders within WCS where certificates can be placed. Only two of them are used by Platform SDK:

- Personal/Certificates – Contains application certificates used by applications to identify themselves.
- Trusted Root Certification Authorities/Certificates – Contains CA certificates used to validate remote party certificates.

To import a certificate, right-click on the appropriate folder and choose *All Tasks > Import...* from the context menu. Follow the steps presented by the Certificate Import Wizard, and once finished the imported certificate will appear in the certificates list.

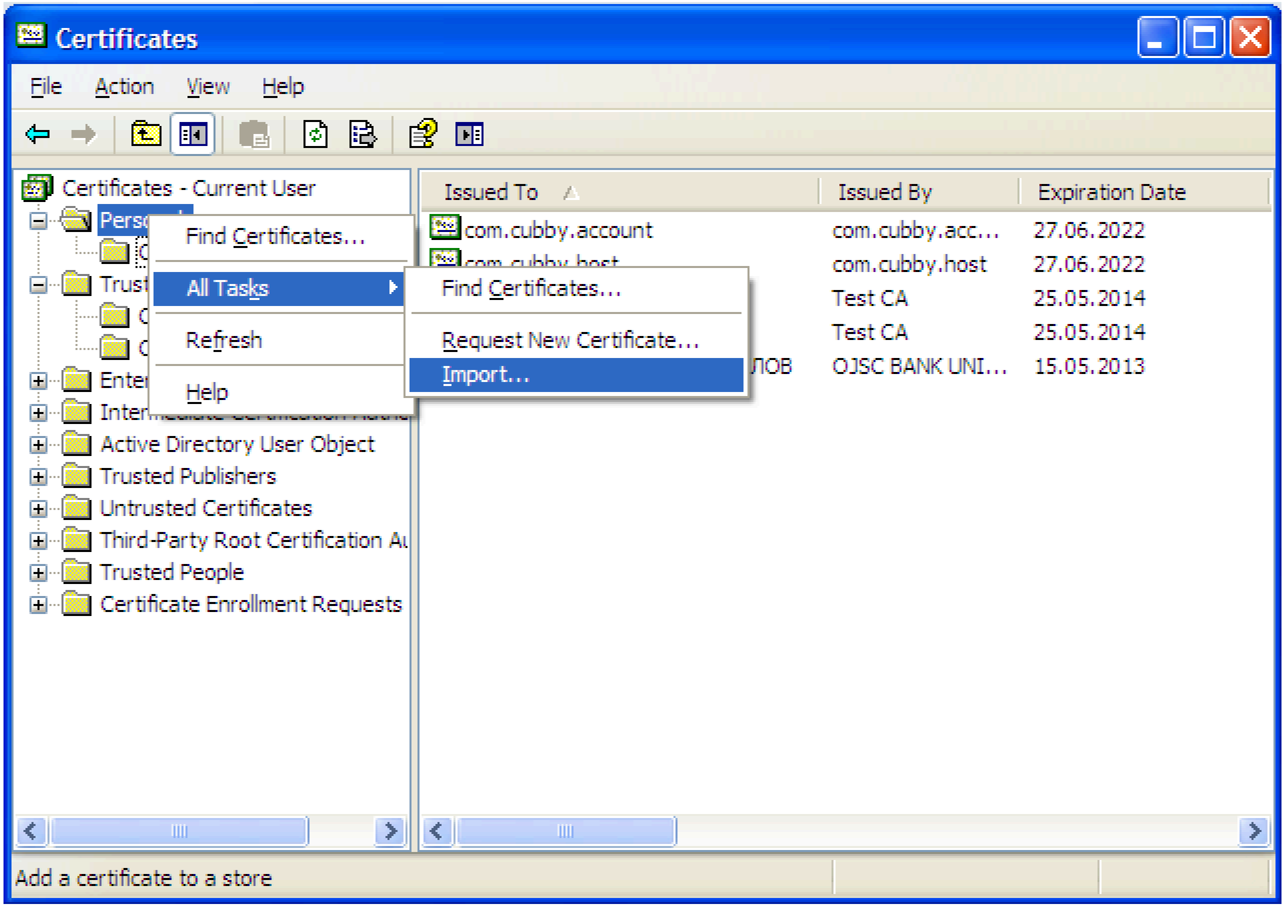

Although WCS can import X.509 PEM certificate files, these certificates cannot be used as application certificates because they do not contain a private key. It is not possible to attach a private key from a PKCS#7 PEM file to the imported certificate. To avoid this problem, import application certificates only from PKCS#12 files (*.*p12*) which contain a certificate and private key pair.

CA certificates do not have private keys attached, so it is safe to import CA certificates from X.509 PEM files.

It is possible to copy and paste certificates between folders and/or user accounts in the Management Console, but this approach is not recommended due to WCS errors which may result in the pasted certificate having an inaccessible private key. This error is not visible in Console, but applications would not be able to read the private key. A recommended and reliable workaround is to export the certificate to a file and then import from that file.

If you encounter the following error in the application log: “The credentials supplied to the package were not recognized”, the most likely cause is due to the private key being absent or inaccessible. In this case try deleting the certificate from WCS and re-importing it.

Importing CRL Files

CRL files can be imported to the following folder in WCS:

- Trusted Root Certification Authorities/Certificate Revocation List

The import procedure is the same as for importing certificate. CRL file types are automatically recognized by the import wizard.

Note: Although an MSCAPI provider may choose to use CRL while validating remote party certificates, this functionality is not guaranteed and/or supported by Platform SDK. Platform SDK implements its own CRL matching logic using CRL PEM files.

PKCS11 Provider: Mozilla NSS

PKCS11 stands for the PKCS#11 family of Public-Key Cryptography Standards (PKCS), published by RSA Laboratories. These standards define platform-independent API-to-cryptographic tokens, such as Hardware Security Modules (HSM) and smart cards, allowing you to connect to external certificate storage devices and/or cryptographic engines.

In Java, the PKCS#11 interface is a simple pass-through and all processing is done externally. When used together with a FIPS-certified security provider, such as Mozilla NSS, the whole provider chain is FIPS-compliant.

Platform SDK uses PKCS11 because it is the only way to achieve FIPS-140 compliance with Java.

Installing Mozilla NSS

Currently Platform SDK only supports FIPS when used with the Mozilla NSS security provider. (Java has FIPS certification only when working with a PKCS#11-compatible pluggable security provider, and the only provider with FIPS certification and Java support is Mozilla NSS.)

Note: In theory, BSafe can be used since it supports JCA interfaces. However, Platform SDK was not tested with RSA BSafe and such system would not be FIPS-certifiable as a while.

Generally, some security parameters and data must be configured on client host, requiring the involvement of a system/security administrator. At minimum, the client host must have a copy of the CA Certificate to be able to validate the Configuration Server certificate. The exact location of the CA certificate depends on the security provider being used. It can be present as a PEM file, Java Keystore file, a record in WCS, or as an entry in the Mozilla NSS database. Once the application is connected to Configuration Server, the Application Template Application Block can be used to extract connection parameters from Configuration Server and set up TLS.

Mozilla NSS is the most complex security provider to deploy and configure. In order to use NSS, the following steps must be completed:

1. Deploy Mozilla NSS.
2. Create Mozilla NSS database (a "soft token" in terms of NSS), and set it to FIPS mode.
3. Adjust the Java security configuration, or implement dynamic loading for the Mozilla NSS provider.
4. Import the CA certificate to the Mozilla NSS database.
5. Use the Platform SDK interface to select PKCS11 as a provider (with no specific configuration options required).

Configuring FIPS Mode in Mozilla NSS

To configure FIPS mode in Mozilla NSS, create a file named *nss-client.cfg* in Mozilla NSS deployment folder with the following values configured:

- name - Name of a software token.
- nssLibraryDirectory - Library directory, located in the Mozilla NSS deployment folder.
- nssSecmodDirectory - Folder where the Mozilla NSS database for the listed software token is located.
- nssModule - Indicates that FIPS mode should be used.

An example is provided below:

```
name = NSSfips
nssLibraryDirectory = C:/nss-3.12.4/lib
nssSecmodDirectory = C:/nss-3.12.4/client
nssModule = fips
```

More information about configuring FIPS mode is available from external sources.

Configuring FIPS Mode in Java Runtime Environment (JRE)

To configure your Java runtime to use Mozilla NSS, the *java.security* file should be located in Java deployment folder and edited as shown below:

(Changes are shown in **bold red**, insertions are shown in **bold blue**)

```
#
# List of providers and their preference orders (see above):
#
security.provider.1=sun.security.provider.Sun
security.provider.2=sun.security.rsa.SunRsaSign
security.provider.3=sun.security.ec.SunEC
#security.provider.4=com.sun.net.ssl.internal.ssl.Provider
security.provider.4=com.sun.net.ssl.internal.ssl.Provider SunPKCS11-NSSfips
security.provider.5=com.sun.crypto.provider.SunJCE
security.provider.6=sun.security.jgss.SunProvider
security.provider.7=com.sun.security.sasl.Provider
security.provider.8=org.jcp.xml.dsig.internal.dom.XMLDSigRI
security.provider.9=sun.security.smartcardio.SunPCSC
security.provider.10=sun.security.mscapi.SunMSCAPI
security.provider.11=sun.security.pkcs11.SunPKCS11 C:/nss-3.12.4/nss-client.cfg
```

After those updates are complete, the Java runtime instance works with FIPS mode, with only the PKCS#11/Mozilla NSS security provider enabled.

Short Command Line Reference

Please refer to the following references for more information:

- https://www.mozilla.org/projects/security/pki/nss/tools/certutil.html
- https://www.mozilla.org/projects/security/pki/nss/tools/crlutil.html
- https://www.mozilla.org/projects/security/pki/nss/tools/pk12util.html

Task	Command
Create CA Certificate	certutil -S -k rsa -n "<CA-cert-name>" -s "CN=Test CA, OU=Miratech, O=Genesys, L=Kyiv, C=UA" -x -t "CTu,u,u" -m 600 -v 24 -d ./client -f "<keystore-password-file>"
Import CA Certificate	certutil -A -a -n “<CA-cert-name>” -t "CTu,u,u" -i <ca-cert-file> -d ./client -f "<keystore-password-file>"
Create New Leaf Certificate	certutil -S -k rsa -n "<cert-name>" -s "CN=Test CA, OU=Miratech, O=Genesys, L=Kyiv, C=UA" -x -t "u,u,u" -m 666 -v 24 -d ./client -f "<keystore-password-file>" -z "<noise-file>"
Import Leaf Certificate	pk12util -i <cert-file.p12> -n <cert-name> -d ./client -v -h "NSS FIPS 140-2 Certificate DB" -K <keystore-password>
Create CRL	crlutil -d ./client -f "<keystore-password-file>" -G -c "<crl-script-file>" -n "<CA-cert-name>" -l SHA512
Modify CRL	crlutil -d ./client -f "<keystore-password-file>" -M -c "<crl-script-file>" -n "<CA-cert-name>" -l SHA512 -B
Show Certificate Information	certutil -d ./client -f "<keystore-password-file>" -L -n "<cert-name>"
Show CRL Information	crlutil -d ./client -f "<keystore-password-file>" -L -n "<CA-cert-name>"
List Certificates	certutil -d ./client –L
List CRLs	crlutil -L -d ./client

JKS Provider: Java Built-in

This provider is supported by the Platform SDK Commons library, but the Application Template Application Block does not support this provider due to compatibility guidelines with Genesys Framework Deployment.

This provider can only be used when TLS is configured programmatically by Platform SDK users.

Short Command Line Reference

Refer to the following reference for more information:

- http://docs.oracle.com/javase/1.5.0/docs/tooldocs/solaris/keytool.html

Task	Command
Creating and Importing - These commands allow you to generate a new Java Keytool keystore file, create a Certificate Signing Request (CSR), and import certificates. Any root or intermediate certificates	

Task	Command
will need to be imported before importing the primary certificate for your domain.	
Generate a Java keystore and key pair	`keytool -genkey -alias mydomain -keyalg RSA -keystore keystore.jks -keysize 2048`
Generate a certificate signing request (CSR) for an existing Java keystore	`keytool -certreq -alias mydomain -keystore keystore.jks -file mydomain.csr`
Import a root or intermediate CA certificate to an existing Java keystore	`keytool -import -trustcacerts -alias root -file Thawte.crt -keystore keystore.jks`
Import a signed primary certificate to an existing Java keystore	`keytool -import -trustcacerts -alias mydomain -file mydomain.crt -keystore keystore.jks`
Generate a keystore and self-signed certificate	`keytool -genkey -keyalg RSA -alias selfsigned -keystore keystore.jks -storepass password -validity 360 -keysize 2048`
Java Keytool Commands for Checking - If you need to check the information within a certificate, or Java keystore, use these commands.	
Check a stand-alone certificate	`keytool -printcert -v -file mydomain.crt`
Check which certificates are in a Java keystore	`keytool -list -v -keystore keystore.jks`
Check a particular keystore entry using an alias	`keytool -list -v -keystore keystore.jks -alias mydomain`
Other Java Keytool Commands	
Delete a certificate from a Java Keytool keystore	`keytool -delete -alias mydomain -keystore keystore.jks`
Change a Java keystore password	`keytool -storepasswd -new new_storepass -keystore keystore.jks`
Export a certificate from a keystore	`keytool -export -alias mydomain -file mydomain.crt -keystore keystore.jks`
List Trusted CA Certs	`keytool -list -v -keystore $JAVA_HOME/jre/lib/security/cacerts`
Import New CA into Trusted Certs	`keytool -import -trustcacerts -file /path/to/ca/ca.pem -alias CA_ALIAS -keystore $JAVA_HOME/jre/lib/security/cacerts`

OpenSSL Configuration File

This page provides an example of a customized OpenSSL configuration file that has been edited to work with the Platform SDK implementation of TLS. For more details about OpenSSL and how it relates to the Platform SDK implementation of TLS, refer to the Using and Configuring Security Providers page.

Sample File

Customized file content is listed below.

- Changes are marked with **bold red**.
- Added lines are marked with **bold blue**.

```
#
# OpenSSL example configuration file.
# This is mostly being used for generation of certificate requests.
#

# This definition stops the following lines choking if HOME isn't
# defined.
HOME                        = .
RANDFILE                = $ENV::HOME/.rnd

# Extra OBJECT IDENTIFIER info:
#oid_file                = $ENV::HOME/.oid
oid_section                = new_oids

# To use this configuration file with the "-extfile" option of the
# "openssl x509" utility, name here the section containing the
# X.509v3 extensions to use:
# extensions                =
# (Alternatively, use a configuration file that has only
# X.509v3 extensions in its main [= default] section.)

[ new_oids ]

# We can add new OIDs in here for use by 'ca', 'req' and 'ts'.
# Add a simple OID like this:
# testoid1=1.2.3.4
# Or use config file substitution like this:
# testoid2=${testoid1}.5.6

# Policies used by the TSA examples.
tsa_policy1 = 1.2.3.4.1
tsa_policy2 = 1.2.3.4.5.6
tsa_policy3 = 1.2.3.4.5.7

####################################################################
[ ca ]
default_ca        = CA_default                # The default ca section

####################################################################
```

```
[ CA_default ]

dir                = ./ca                        # Where everything is kept
certs                = $dir/certs                # Where the issued certs are kept
crl_dir                = $dir/crl                # Where the issued crl are kept
database        = $dir/index.txt        # database index file.
#unique_subject        = no                        # Set to 'no' to allow creation of
                                        # several ctificates with same subject.
new_certs_dir        = $dir/newcerts                # default place for new certs.

certificate        = $dir/cacert.pem         # The CA certificate
serial                = $dir/serial                 # The current serial number
crlnumber        = $dir/crlnumber        # the current crl number
                                        # must be commented out to leave a V1 CRL
crl                = $dir/crl.pem                 # The current CRL
private_key        = $dir/private/cakey.pem # The private key
RANDFILE        = $dir/private/.rand        # private random number file

x509_extensions        = usr_cert                # The extentions to add to the cert

# Comment out the following two lines for the "traditional"
# (and highly broken) format.
name_opt         = ca_default                # Subject Name options
cert_opt         = ca_default                # Certificate field options

# Extension copying option: use with caution.
# copy_extensions = copy

# Extensions to add to a CRL. Note: Netscape communicator chokes on V2 CRLs
# so this is commented out by default to leave a V1 CRL.
# crlnumber must also be commented out to leave a V1 CRL.
# crl_extensions        = crl_ext

default_days        = 365                        # how long to certify for
default_crl_days= 30                        # how long before next CRL
default_md        = default                # use public key default MD
preserve        = no                        # keep passed DN ordering

# A few difference way of specifying how similar the request should look
# For type CA, the listed attributes must be the same, and the optional
# and supplied fields are just that :-)
policy                = policy_anything

# For the CA policy
[ policy_match ]
countryName                = match
stateOrProvinceName        = match
organizationName        = match
organizationalUnitName        = optional
commonName                = supplied
emailAddress                = optional

# For the 'anything' policy
# At this point in time, you must list all acceptable 'object'
# types.
[ policy_anything ]
countryName                = optional
stateOrProvinceName        = optional
localityName                = optional
organizationName        = optional
organizationalUnitName        = optional
commonName                = supplied
emailAddress                = optional
```

```
####################################################################
[ req ]
default_bits                = 1024
default_keyfile         = privkey.pem
distinguished_name        = req_distinguished_name
attributes                = req_attributes
x509_extensions        = v3_ca        # The extentions to add to the self signed cert

# Passwords for private keys if not present they will be prompted for
# input_password = secret
# output_password = secret

# This sets a mask for permitted string types. There are several options.
# default: PrintableString, T61String, BMPString.
# pkix         : PrintableString, BMPString (PKIX recommendation before 2004)
# utf8only: only UTF8Strings (PKIX recommendation after 2004).
# nombstr : PrintableString, T61String (no BMPStrings or UTF8Strings).
# MASK:XXXX a literal mask value.
# WARNING: ancient versions of Netscape crash on BMPStrings or UTF8Strings.
string_mask = utf8only

req_extensions = v3_req # The extensions to add to a certificate request

[ req_distinguished_name ]
countryName                        = Country Name (2 letter code)
countryName_default                = UA
countryName_min                        = 2
countryName_max                        = 2

stateOrProvinceName                = State or Province Name (full name)
stateOrProvinceName_default        = None

localityName                        = Locality Name (eg, city)
localityName_default        = Kyiv

0.organizationName                = Organization Name (eg, company)
0.organizationName_default        = Genesys

# we can do this but it is not needed normally :-)
#1.organizationName                = Second Organization Name (eg, company)
#1.organizationName_default        = World Wide Web Pty Ltd

organizationalUnitName                = Organizational Unit Name (eg, section)
organizationalUnitName_default        = Engineering

commonName                        = Common Name (e.g. server FQDN or YOUR name)
commonName_default        = xpigors
commonName_max                        = 64

emailAddress                        = Email Address
emailAddress_max                = 64

# SET-ex3                        = SET extension number 3

[ req_attributes ]
challengePassword                = A challenge password
challengePassword_min                = 0
challengePassword_max                = 20

unstructuredName                = An optional company name

[ usr_cert ]
```

```
# These extensions are added when 'ca' signs a request.

# This goes against PKIX guidelines but some CAs do it and some software
# requires this to avoid interpreting an end user certificate as a CA.

basicConstraints=CA:FALSE

# Here are some examples of the usage of nsCertType. If it is omitted
# the certificate can be used for anything *except* object signing.

# This is OK for an SSL server.
# nsCertType                        = server

# For an object signing certificate this would be used.
# nsCertType = objsign

# For normal client use this is typical
# nsCertType = client, email

# and for everything including object signing:
# nsCertType = client, email, objsign

# This is typical in keyUsage for a client certificate.
# keyUsage = nonRepudiation, digitalSignature, keyEncipherment

# This will be displayed in Netscape's comment listbox.
nsComment                         = "OpenSSL Generated Certificate"

# PKIX recommendations harmless if included in all certificates.
subjectKeyIdentifier=hash
authorityKeyIdentifier=keyid,issuer

# This stuff is for subjectAltName and issuerAltname.
# Import the email address.
#subjectAltName=issue:copy
subjectAltName = @alt_names
# An alternative to produce certificates that aren't
# deprecated according to PKIX.
# subjectAltName=email:move

# Copy subject details
# issuerAltName=issuer:copy

#nsCaRevocationUrl = http://www.domain.dom/ca-crl.pem
#nsBaseUrl
#nsRevocationUrl
#nsRenewalUrl
#nsCaPolicyUrl
#nsSslServerName

# This is required for TSA certificates.
# extendedKeyUsage = critical,timeStamping

[ v3_req ]

# Extensions to add to a certificate request

basicConstraints = CA:FALSE
keyUsage = nonRepudiation, digitalSignature, keyEncipherment
subjectAltName = @alt_names

[ alt_names ]
```

```
DNS.1 = hostname.emea.int.genesyslab.com
DNS.2 = hostname
IP.1 = 192.168.1.1
IP.2 = fe80::21d:7dff:fe0d:682c
IP.3 = fe80::ffff:ffff:fffd
IP.4 = fe80::5efe:192.168.1.1
URI.1 = http://hostname/
URI.2 = https://hostname/
email.1 = UserName1@genesyslab.com
email.2 = UserName2@genesyslab.com

[ v3_ca ]


# Extensions for a typical CA


# PKIX recommendation.

subjectKeyIdentifier=hash

authorityKeyIdentifier=keyid:always,issuer

# This is what PKIX recommends but some broken software chokes on critical
# extensions.
#basicConstraints = critical,CA:true
# So we do this instead.
basicConstraints = CA:true

# Key usage: this is typical for a CA certificate. However since it will
# prevent it being used as an test self-signed certificate it is best
# left out by default.
# keyUsage = cRLSign, keyCertSign

# Some might want this also
# nsCertType = sslCA, emailCA

# Include email address in subject alt name: another PKIX recommendation
# subjectAltName=email:copy
# Copy issuer details
# issuerAltName=issuer:copy

# DER hex encoding of an extension: beware experts only!
# obj=DER:02:03
# Where 'obj' is a standard or added object
# You can even override a supported extension:
# basicConstraints= critical, DER:30:03:01:01:FF

[ crl_ext ]

# CRL extensions.
# Only issuerAltName and authorityKeyIdentifier make any sense in a CRL.

# issuerAltName=issuer:copy
authorityKeyIdentifier=keyid:always

[ proxy_cert_ext ]
# These extensions should be added when creating a proxy certificate

# This goes against PKIX guidelines but some CAs do it and some software
# requires this to avoid interpreting an end user certificate as a CA.

basicConstraints=CA:FALSE
```

```
# Here are some examples of the usage of nsCertType. If it is omitted
# the certificate can be used for anything *except* object signing.

# This is OK for an SSL server.
# nsCertType                        = server

# For an object signing certificate this would be used.
# nsCertType = objsign

# For normal client use this is typical
# nsCertType = client, email

# and for everything including object signing:
# nsCertType = client, email, objsign

# This is typical in keyUsage for a client certificate.
# keyUsage = nonRepudiation, digitalSignature, keyEncipherment

# This will be displayed in Netscape's comment listbox.
nsComment                        = "OpenSSL Generated Certificate"

# PKIX recommendations harmless if included in all certificates.
subjectKeyIdentifier=hash
authorityKeyIdentifier=keyid,issuer

# This stuff is for subjectAltName and issuerAltname.
# Import the email address.
# subjectAltName=email:copy
# An alternative to produce certificates that aren't
# deprecated according to PKIX.
# subjectAltName=email:move

# Copy subject details
# issuerAltName=issuer:copy

#nsCaRevocationUrl = http://www.domain.dom/ca-crl.pem
#nsBaseUrl
#nsRevocationUrl
#nsRenewalUrl
#nsCaPolicyUrl
#nsSslServerName

# This really needs to be in place for it to be a proxy certificate.
proxyCertInfo=critical,language:id-ppl-anyLanguage,pathlen:3,policy:foo

####################################################################
[ tsa ]

default_tsa = tsa_config1        # the default TSA section

[ tsa_config1 ]

# These are used by the TSA reply generation only.
dir                = ./demoCA                # TSA root directory
serial                = $dir/tsaserial        # The current serial number (mandatory)
crypto_device        = builtin                # OpenSSL engine to use for signing
signer_cert        = $dir/tsacert.pem         # The TSA signing certificate
                                        # (optional)
certs                = $dir/cacert.pem        # Certificate chain to include in reply
                                        # (optional)
signer_key        = $dir/private/tsakey.pem # The TSA private key (optional)
```

```
default_policy        = tsa_policy1                # Policy if request did not specify it
                                       # (optional)
other_policies        = tsa_policy2, tsa_policy3        # acceptable policies (optional)
digests                = md5, sha1                # Acceptable message digests (mandatory)
accuracy        = secs:1, millisecs:500, microsecs:100        # (optional)
clock_precision_digits  = 0        # number of digits after dot. (optional)
ordering                = yes        # Is ordering defined for timestamps?
                                # (optional, default: no)
tsa_name                = yes        # Must the TSA name be included in the reply?
                                # (optional, default: no)
ess_cert_id_chain        = no        # Must the ESS cert id chain be included?
                                # (optional, default: no)
```

Use Cases

Introduction

This page examines TLS functionality as a series of common use cases. Use cases are broken into two categories: server or application.

Examples and explanations are provided for some use cases, while others simply provide links to the related TLS documentation needed to understand the functionality.

Genesys Server Use Cases

Opening a TLS Port

Code snippets explaining how to open a basic TLS port are provided both with, and without using the Application Template Application Block:

- Opening a TLS port using the Platform SDK Commons Library
- Opening a TLS port using the Application Template Application Block

Opening a Mutual TLS Port (With Expiration, Revocation and CA Checks)

This use case is an advanced variation on opening a simple TLS port. As such, it already has a CA and expiration check, but needs additional parameters to turn on mutual mode and to enable a CRL check.

Mutual Mode

If TLS is configured programmatically, then the *mutualTLS* parameter should be set to *true* when creating an *SSLExtendedOptions* object:

```
SSLExtendedOptions sslOptions = new SSLExtendedOptions(true, (String) null);
```

If TLS is configured in Configuration Manager, then the *tls-mutual* parameter for the server port, application or host should be set to *1*. Please refer to the list of TLS parameters for details.

Revocation Check

If TLS is configured programmatically, then a valid path to the CRL file should be provided in the *crlFilePath* parameter when creating a trust manager:

```
X509TrustManager tm = TrustManagerHelper.createPEMTrustManager(
        "c:/cert/ca-cert.pem","c:/cert/crl.pem", null);
```

If TLS is configured in Configuration Manager, then the *tls-crl* parameter for the server port, application or host should contain the path to the CRL file located on server. Please refer to the list of TLS parameters for details.

Opening a FIPS-Compliant Port

FIPS mode is not a property of a port or application; it is defined mostly by the type of security provider in use and the OS/environment settings. For Java, the PKCS#11 security provider should be used to support FIPS; for .Net, FIPS is configured at the OS level (http://technet.microsoft.com/en-us/library/cc750357.aspx).

If TLS is configured programmatically, then a PKCS11 key/trust managers should be used:

```
X509TrustManager tm = TrustManagerHelper.createPKCS11TrustManager(
        new DummyPasswordCallbackHandler(), (String) null);
X509ExtendedKeyManager km = KeyManagerHelper.createPKCS11KeyManager(
        new DummyPasswordCallbackHandler());
```

If TLS is configured in Configuration Manager, then the *fips140-enabled* parameter for the server port, application or host should be set to "1". Please refer to the TLSTLSParametersinConfigManager for details.

Note: This parameter is used to detect the security provider type to use. If this setting conflicts with other TLS parameters or points to a FIPS security provider that is not installed on host, then Platform SDK will generate an exception when attempting to accept or open a connection.

Genesys Application Use Cases

Opening a TLS Connection to a TLS Autodetect Server Port

TLS autodetect ports (also called upgrade mode ports) allow you to establish an unsecured connection to the server before specifying TLS settings. For details, please refer to Connecting to Upgrade Mode Ports in the quick start instructions.

Opening a TLS Connection to a Backend Server (With Expiration, Revocation and CA Checks)

Code snippets explaining how to open a basic TLS connection to a backend server are provided both with, and without using the Application Template Application Block:

- Configuring TLS for Client Connections using the Platform SDK Commons Library
- Configuring TLS for Client Connections using the Application Template Application Block

Opening a FIPS-Compliant Connection to a FIPS-Compliant Port

In this use case, the application does not need to provide any special behavior because the server will only handshake for FIPS-compliant ciphers. Details about setting up a FIPS-compliant port are described above.

Ensuring the Certificate is Checked with CA

If TLS is configured programmatically, then a valid CA certificate data should be provided when creating the trust manager:

```
X509TrustManager tm = TrustManagerHelper.createPEMTrustManager(
        "c:/cert/ca-cert.pem","c:/cert/crl.pem", null);
```

If TLS is configured in Configuration Manager, then the *trusted-ca* parameter for the port, connection, application or host should contain valid CA certificate data. Please refer to the list of TLS parameters for details.

Note: CA certificates are configured differently for each type of security provider. Please refer to the page on using and configuring security providers for detailed information.

Ensuring the Certificate Expiration is Checked

Certificate expiration is checked by default during the certificate validation process.

Note: If a server certificate is placed in a trusted certificates store on the client host, it will be automatically trusted without any validation. A trust certificates store should not include application certificates; instead, it should contain only CA certificates.

Handling a Certificate Revocation List

If TLS is configured programmatically, then a valid path to a CRL file should be provided in the *crlFilePath* parameter when creating trust manager:

```
X509TrustManager tm = TrustManagerHelper.createPEMTrustManager(
        "c:/cert/ca-cert.pem","c:/cert/crl.pem", null);
```

If TLS is configured in Configuration Manager, then the *tls-crl* parameter for the application connection, application or host should contain the path to the CRL file located on the application's host. Please refer to the list of TLS parameters for details.

Handling a User-Specified Cipher List

If TLS is configured programmatically, then the *enabledCipherSuites* constructor parameter should contain a list of allowed ciphers when the *SSLExtendedOptions* object is being created:

```
SSLExtendedOptions sslOptions = new SSLExtendedOptions(
        true, "TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA " +
        "TLS_ECDH_RSA_WITH_3DES_EDE_CBC_SHA");
```

If TLS is configured in Configuration Manager, then the *cipher-list* parameter for the port, connection, application or host should be set to contain list of allowed ciphers. Please refer to the list of TLS parameters for details.

Lazy Parsing of Message Attributes

This page provides:

- an overview and list of requirements for the lazy parsing feature
- design details explaining how this feature works
- code examples showing how to implement lazy parsing in your applications

Introduction to Lazy Parsing

Lazy parsing allows users to specify which attributes should always be parsed immediately, and which attributes should be parsed only on demand.

Some complex attributes (such as the `ConfObject` attribute found in some Configuration Server protocol messages) are large and very complex. Unpacking these attributes can be time-consuming and, in cases when an application is not interested in that data, can affect program performance. This issue is resolved by using the "lazy parsing" feature included with the Platform SDK 8.1 release, which is described in this article.

When this feature is turned off, all message attributes are parsed immediately - which is normal behavior for previous version of the Platform SDK. When lazy parsing is enabled, any attributes that were tagged for lazy parsing are only parsed on demand. In this case, if the application does not explicitly check the value of an attribute tagged for lazy parsing then that attribute is never parsed at all.

Feature Overview

- Platform SDK includes configuration options to turn the lazy parsing functionality on or off for each individual protocol that supports this feature.
- Potentially time-consuming attributes that are candidates for lazy parsing are selected and marked by Platform SDK developers. Refer to your Platform SDK API Reference for details.
- To maintain backwards compatibility, there is no change in how user applications access attribute values.
- By default, the lazy parsing feature is turned off.

Java

System Requirements

Platform SDK for Java:

- Configuration SDK protocol release 8.1 or later<ref name="ConfObjectJava">**Note:** Currently, lazy parsing is only used with the `EventObjectsRead.ConfObject` property of the Configuration Platform SDK.</ref>
- J2SE 5.0 or Java 6 SE runtime

<references/>

Design Details

This section describes the main classes and interfaces you will need to be familiar with to implement lazy parsing in your own application.

Enabling and Disabling the Lazy Parsing Feature

At any time, a running application can enable or disable lazy parsing for a specific protocol in just a few lines of code. This is done in three easy steps:

1. Create a new `KeyValueCollection` object.
2. Set the appropriate value for the `Connection.LAZY_PARSING_ENABLED_KEY` field. A value of `True` enables the feature, while `False` disables lazy parsing.
3. Use a `KeyValueConfiguration` object to apply that setting to the desired protocol(s).

> **Tip**
> The default value of the `Connection.LAZY_PARSING_ENABLED_KEY` field is always `False`, with the lazy parsing feature disabled.

Once lazy parsing mode is enabled for a protocol, the change is applied immediately. Every new message that is received takes the lazy parsing setting into account: parsing entire messages if the feature is disabled, or leaving some attributes unparsed until their values are requested if the feature is enabled.

To enable lazy parsing for the Configuration Server protocol, an application would use the following code:

```
[Java]

KeyValueCollection kv = new KeyValueCollection();
kv.addString(Connection.LAZY_PARSING_ENABLED_KEY, "true");
KeyValueConfiguration kvcfg = new KeyValueConfiguration(kv);
ConfServerProtocol cfgChannel = new ConfServerProtocol(endpoint);
cfgChannel.configure(kvcfg); //lazy parsing is immediately active after this line
```

To disable lazy parsing for the protocol only the second line of code is changed before re-applying the configuration, as shown below:

```
[Java]

kv.addString(Connection.LAZY_PARSING_ENABLED_KEY, "false");
```

.NET

System Requirements

- Configuration SDK protocol release 8.1 or later<ref name="ConfObjectNet">**Note:** Currently, lazy parsing is only used with the `EventObjectsRead.ConfObject` property of the Configuration Platform SDK.</ref>
- .NET Framework 3.5
- Visual Studio 2008 (required for .NET project files)

<references/>

Design Details

This section describes the main classes and interfaces you will need to be familiar with to implement lazy parsing in your own application.

Enabling and Disabling the Lazy Parsing Feature

At any time, a running application can enable or disable lazy parsing for a specific protocol in just a few lines of code. This is done in three easy steps:

1. Create a new `KeyValueCollection` object.
2. Set the appropriate value for the `CommonConnection.LazyParsingEnabledKey` field. A value of `True` enables the feature, while `False` disables lazy parsing.
3. Use a `KeyValueConfiguration` object to apply that setting to the desired protocol(s).

> **Tip**
> The default value of the `CommonConnection.LazyParsingEnabledKey` field is always `False`, with the lazy parsing feature disabled.

Once lazy parsing mode is enabled for a protocol, the change is applied immediately. Every new message that is received takes the lazy parsing setting into account: parsing entire messages if the feature is disabled, or leaving some attributes unparsed until their values are requested if the feature is enabled.

To enable lazy parsing for the Configuration Server protocol, an application would use the following code:

```
[C#]

KeyValueCollection kvc = new KeyValueCollection();
kvc[CommonConnection.LazyParsingEnabledKey] = "true";
KeyValueConfiguration kvcfg = new KeyValueConfiguration(kvc);
ConfServerProtocol cfgChannel = new ConfServerProtocol(endpoint);
cfgChannel.Configure(kvcfg); //lazy parsing is immediately active after this line
```

To disable lazy parsing for the protocol only the second line of code is changed before re-applying the configuration, as shown below:

```
[C#]

kvc[CommonConnection.LazyParsingEnabledKey] = "false";
```

Accessing Attribute Values

There is no difference in how applications access attribute values from returned messages. Whether the lazy parsing feature is enabled or disabled, whether the attribute being access supports lazy parsing or not, your code remains exactly the same.

However, you should consider differences in timing when accessing attribute values.

- When lazy parsing is disabled, the entire message is parsed immediately when it is received. Accessing attribute values is very fast, as the requested information is already prepared.
- When lazy parsing is enabled, the delay to parse the message upon arrival is smaller but accessing any attributes that support lazy parsing causes a slightly delay as that information must first be parsed. Note that accessing the same attribute a second time will not result in the attribute information being parsed a second time; Platform SDK saves parsed data.

Additional Notes

- XML Serialization — The `XmlMessageSerializer` class has been updated to support lazy parsing. If a message that contains unparsed attributes is serialized, then `XmlMessageSerializer` will trigger parsing before the serialization process begins.
- `ToString` method — Use of the `ToString` method does not trigger parsing of attributes that support lazy parsing. In this case, each unparsed attribute has its name printed along with a value of: "<value is not yet parsed>".

Server-Specific Overviews

- Using the Voice Platform SDK
    - List of TLib Functions
    - List of TLib Datatypes
    - List of TLib Unstructured Data
- Using the Configuration Platform SDK
    - Connecting Using the UTF-8 Enconding
    - Change Password On Next Login
    - Getting the Last Login Info
    - Using the Configuration Object Model Application Block
    - Introduction to Configuration Layer Objects
    - List of Configuration Layer Objects
    - List of Configuration Layer Enumerations
- Using the Statistics Platform SDK
- Using the Open Media Platform SDK
- Using the Contacts Platform SDK
    - Creating an E-Mail
- Using the Web Media Platform SDK
- Using the Outbound Contact Platform SDK
- Using the Management Platform SDK
    - LCA Hang-Up Detection Support
    - Handle Application "Graceful Stop" with the LCA Protocol
- Using the Universal Routing Platform SDK

Telephony (T-Server)

Java

You can use the Voice Platform SDK to write Java or .NET applications that monitor and handle voice interactions from a traditional or IP-based telephony device. These applications can range from the simple to the advanced. This document shows how to implement the basic functions you will need to write a simple voice application. It is organized to show the kind of structure you will probably use to write your own applications.

Setting Up a TServerProtocol Object

The first thing you need to do to use the Voice Platform SDK is instantiate a `TServerProtocol` object. To do that, you must supply information about the T-Server you want to connect with. This example provides the server's name, host, and port information:

```
[Java]

TServerProtocol tServerProtocol =
        new TServerProtocol(
                new Endpoint(
                        tServerName, host, port));
```

After instantiating the `TServerProtocol` object, you need to open the connection to the T-Server:

```
[Java]

tServerProtocol.open();
```

Registering an Address

Now you need to register a DN for your agent to use. To do this, you must send a `RequestRegisterAddress` request to the server.

Here is how to create this request:

```
[Java]

RequestRegisterAddress requestRegisterAddress =
        RequestRegisterAddress.create(
                thisDn,
                RegisterMode.ModeShare,
                ControlMode.RegisterDefault,
                AddressType.DN);
```

The thisDn argument refers to the DN you want to associate with your agent, while RegisterMode.ModeShare tells the T-Server to share information about the DN with other applications. The next argument asks to use the switch's default value for deciding whether to let the switch know that you have registered this DN. And finally, you are specifying that the object you are registering is a DN.

After you create the request, you will need to send it to the T-Server:

```
[Java]

Message response =
        tServerProtocol.request(requestRegisterAddress);
```

Remember that the request() method is synchronous. If you use this method, your application will block until you hear back from the server. When you get the response, you can execute code to handle the response. In this case, you probably don't need to do anything if the request is successful:

```
[Java]

switch(response.messageId())
{
        case EventRegistered.ID:
        case EventUnregistered.ID:
                break;
        .
        .
        .
}
```

Logging in an Agent

Once you have registered a DN to your agent, you can log him or her in. To do this, you need to create a RequestAgentLogin request:

```
[Java]

RequestAgentLogin requestAgentLogin =
        RequestAgentLogin.create(
                thisDn,
                AgentWorkMode.AutoIn);
```

After you create the request, you will need to indicate the queue the agent will be using, and you may need to supply the agent's user name and password. Once you have done this, you can send the request to the server:

```
[Java]

requestAgentLogin.setThisQueue(thisQueue);
// Your switch may not need a user name and password:
requestAgentLogin.setAgentID(userName);
requestAgentLogin.setPassword(password);
Message response = tServerProtocol.request(requestAgentLogin);
```

If your request is successful, the server will respond with an EventAgentLogin event. At that point, you may need to update the state of your user interface to indicate that the agent can no longer log in, but that, for example, he or she can now log out.

Answering a Call

Now that your agent is logged in, he or she can handle calls. Let's start by answering a call.

When a call comes in, your application will receive an `EventRinging` message. When you get this message, you will probably want to enable an answer button. Here is how to do that:

```
[Java]

switch(response.messageId())
{
        .
        .
        .
        case EventRinging.ID:
                EventRinging eventRinging = (EventRinging) response;
                connId = eventRinging.getConnID();
                if (eventRinging.getThisDN() == thisDn)
                {
                        AnswerButton.enabled = true;
                }
                break;
        .
        .
        .
}
```

It is important to note that an `EventRinging` event will also be triggered when you are sending an outbound call. So this particular snippet is only enabling the answer button if the call is ringing on `thisDN`. As you can also see, when you receive an `EventRinging` you will want to store the `ConnID` of the call associated with it.

After the agent clicks the answer button, you need to send a request to answer the call, using your DN and the `ConnID` of the call:

```
[Java]

RequestAnswerCall requestAnswerCall =
        RequestAnswerCall.create(
        thisDn,
        connId);
Message response = tServerProtocol.request(requestAnswerCall);
```

If the request is successful, you will receive an `EventEstablished`.

Releasing a Call

When your agent is finished with the call, he or she will need to release it:

```
[Java]

RequestReleaseCall requestReleaseCall =
        RequestReleaseCall.create(
                thisDn,
                connId);
```

```
Message response = tServerProtocol.request(requestReleaseCall);
```

If the request is successful, you will receive an `EventReleased`.

Making a Call

Here is how to make a call:

```
[Java]

RequestMakeCall requestMakeCall =
        RequestMakeCall.create(
                thisDn,
                thatDn,
                MakeCallType.DirectAgent);
Message response = tServerProtocol.request(requestMakeCall);
```

If the request is successful, you will receive an `EventDialing` message, an `EventRinging` message, and then, when your party responds, an `EventEstablished` message.

Setting up a Conference Call

After you make or answer a call, you can add another party to the call. Here is how to perform an ordinary two-step conference call.

To start off, you need to initiate a conference call, supplying your own DN, the connection ID of the existing call, and the DN of the party you want to add to the call:

```
[Java]

RequestInitiateConference requestInitiateConference =
        RequestInitiateConference.create(
                thisDn,
                connId,
                otherDn);
Message response = tServerProtocol.request(requestInitiateConference);
```

> **Tip**
>
> In a real telephony application, the events you would receive in response to the kinds of conferencing requests shown here could also be generated by other requests. For example, you might receive an `EventDialing` or an `EventEstablished` in response to a `RequestMakeCall` or `RequestInitiateTransfer`. Because of this, a real-world application will need to keep track of the requests that initiate these events in order to interpret them correctly.

If the initiate request is successful, you will receive an `EventDialing` message and an `EventHeld` message. When your party picks up the call, you will also receive an `EventEstablished` message.

Now you need to complete the conference call.

When you received the `EventDialing` message from the `RequestInitiateConference`, you were given a new connection ID associated with the party you want to establish the conference call with. You will need that connection ID, in addition to your own DN and the original connection ID, in order to complete the conference call:

```
[Java]

RequestCompleteConference requestCompleteConference =
        RequestCompleteConference.create(
                thisDn,
                connId,
                secondConnId);
response = tServerProtocol.request(requestCompleteConference);
```

If the completion request is successful, you will receive `EventReleased`, `EventRetrieved`, `EventPartyAdded`, and `EventAttachedDataChanged` messages.

Transferring a Call

After you make or answer a call, you may also want to transfer that call. Here is how to perform an ordinary two-step transfer.

To start off, you need to initiate a transfer, supplying your own DN, the connection ID of the existing call, and the DN of the party you want to transfer the call to.

```
[Java]

RequestInitiateTransfer requestInitiateTransfer =
        RequestInitiateTransfer.create(
                thisDn,
                connId,
                otherDn);
Message response = tServerProtocol.request(requestInitiateTransfer);
```

> **Tip**
>
> In a real telephony application, the events you would receive in response to the kinds of transfer requests shown here could also be generated by other requests. For example, you might receive an `EventDialing` or an `EventEstablished` in response to a `RequestMakeCall` or `RequestInitiateConference`. Because of this, a real-world application will need to keep track of the requests that initiate these events in order to interpret them correctly.

If the initiate request is successful, you will receive an `EventDialing` message and an `EventHeld` message. When the party you want to transfer to picks up the call, you will also receive an `EventEstablished` message.

Now you need to complete the transfer.

When you received the EventDialing message from the RequestInitiateTransfer, you were given a new connection ID associated with the party you want to transfer the call to. You will need that connection ID, in addition to your own DN and the original connection ID, in order to complete the transfer:

```
[Java]

RequestCompleteTransfer requestCompleteTransfer =
        RequestCompleteTransfer.create(
                thisDn,
                connId,
                secondConnId);
response = tServerProtocol.request(requestCompleteTransfer);
```

If the completion request is successful, you will receive two EventReleased messages and you will no longer be a party to the call.

Closing the Connection

Finally, when you are finished communicating with the T-Server, you should close the connection to minimize resource utilization:

```
[Java]

tServerProtocol.close();
```

.NET

You can use the Voice Platform SDK to write Java or .NET applications that monitor and handle voice interactions from a traditional or IP-based telephony device. These applications can range from the simple to the advanced. This document shows how to implement the basic functions you will need to write a simple voice application. It is organized to show the kind of structure you will probably use to write your own applications.

Setting Up a TServerProtocol Object

The first thing you need to do to use the Voice Platform SDK is instantiate a TServerProtocol object. To do that, you must supply information about the T-Server you want to connect with. This example uses the URI of the T-Server, but you can also use name, host, and port information:

```
[C#]

TServerProtocol tServerProtocol =
        new TServerProtocol(
                new Endpoint(
                        tServerUri));
```

After instantiating the TServerProtocol object, you need to open the connection to the T-Server:

```
[C#]

tServerProtocol.Open();
```

Registering an Address

Now you need to register a DN for your agent to use. To do this, you must send a `RequestRegisterAddress` request to the server.

Here is how to create this request:

```
[C#]

RequestRegisterAddress requestRegisterAddress =
        RequestRegisterAddress.Create(
        thisDn,
        RegisterMode.ModeShare,
        ControlMode.RegisterDefault,
        AddressType.DN);
```

The `thisDn` argument refers to the DN you want to associate with your agent, while `RegisterMode.ModeShare` tells the T-Server to share information about the DN with other applications. The next argument asks to use the switch's default value for deciding whether to let the switch know that you have registered this DN. And finally, you are specifying that the object you are registering is a DN.

After you create the request, you will need to send it to the T-Server:

```
[C#]

IMessage response = tServerProtocol.Request(requestRegisterAddress);
```

Remember that the `Request()` method is synchronous. If you use this method, your application will block until you hear back from the server. When you get the response, you can execute code to handle the response. In this case, you probably don't need to do anything if the request is successful:

```
[C#]

switch(response.Id )
{
        case EventRegistered.MessageId:
        case EventUnregistered.MessageId:
                break;
        .
        .
        .
}
```

Logging in an Agent

Once you have registered a DN to your agent, you can log him or her in. To do this, you need to create a `RequestAgentLogin` request:

```
[C#]

RequestAgentLogin requestAgentLogin =
        RequestAgentLogin.Create(
        thisDn,
        AgentWorkMode.AutoIn);
```

After you create the request, you will need to indicate the queue the agent will be using, and you may need to supply the agent's user name and password. Once you have done this, you can send the request to the server:

```
[C#]

requestAgentLogin.ThisQueue = thisQueue;
// Your switch may not need a user name and password:
requestAgentLogin.AgentID = userName;
requestAgentLogin.Password = password;
IMessage response = tServerProtocol.Request(requestAgentLogin);
```

If your request is successful, the server will respond with an `EventAgentLogin` event. At that point, you may need to update the state of your user interface to indicate that the agent can no longer log in, but that, for example, he or she can now log out.

Answering a Call

Now that your agent is logged in, he or she can handle calls. Let's start by answering a call.

When a call comes in, your application will receive an `EventRinging` message. When you get this message, you will probably want to enable an answer button. Here is how to do that:

```
[C#]

switch(response.Id)
{
        .
        .
        .
        case EventRinging.MessageId:
                EventRinging eventRinging = (EventRinging) response;
                connId = eventRinging.ConnID;
                if (eventRinging.ThisDN == thisDn)
                {
                        AnswerButton.Enabled = true;
                }
                break;
        .
        .
        .
}
```

It is important to note that an `EventRinging` event will also be triggered when you are sending an outbound call. So this particular snippet is only enabling the answer button if the call is ringing on `thisDN`. As you can also see, when you receive an `EventRinging` you will want to store the `ConnID` of the call associated with it.

After the agent clicks the answer button, you need to send a request to answer the call, using your

DN and the ConnID of the call:

```
[C#]

RequestAnswerCall requestAnswerCall =
        RequestAnswerCall.Create(
        thisDn,
        connId);
IMessage response = tServerProtocol.Request(requestAnswerCall);
```

If the request is successful, you will receive an EventEstablished.

Releasing a Call

When your agent is finished with the call, he or she will need to release it:

```
[C#]

RequestReleaseCall requestReleaseCall =
        RequestReleaseCall.Create(
        thisDn,
        connId);
IMessage response = tServerProtocol.Request(requestReleaseCall);
```

If the request is successful, you will receive an EventReleased.

Making a Call

Here is how to make a call:

```
[C#]

RequestMakeCall requestMakeCall =
        RequestMakeCall.Create(
        thisDn,
        thatDn,
        MakeCallType.DirectAgent);
IMessage response = tServerProtocol.Request(requestMakeCall);
```

If the request is successful, you will receive an EventDialing message, an EventRinging message, and then, when your party responds, an EventEstablished message.

Setting up a Conference Call

After you make or answer a call, you can add another party to the call. Here is how to perform an ordinary two-step conference call.

To start off, you need to initiate a conference call, supplying your own DN, the connection ID of the existing call, and the DN of the party you want to add to the call:

```
[C#]

RequestInitiateConference requestInitiateConference =
        RequestInitiateConference.Create(
        thisDn,
        connId,
        otherDn);
IMessage response = tServerProtocol.Request(requestInitiateConference);
```

> **Tip**
>
> In a real telephony application, the events you would receive in response to the kinds of conferencing requests shown here could also be generated by other requests. For example, you might receive an `EventDialing` or an `EventEstablished` in response to a `RequestMakeCall` or `RequestInitiateTransfer`. Because of this, a real-world application will need to keep track of the requests that initiate these events in order to interpret them correctly.

If the initiate request is successful, you will receive an `EventDialing` message and an `EventHeld` message. When your party picks up the call, you will also receive an `EventEstablished` message.

Now you need to complete the conference call.

When you received the `EventDialing` message from the `RequestInitiateConference`, you were given a new connection ID associated with the party you want to establish the conference call with. You will need that connection ID, in addition to your own DN and the original connection ID, in order to complete the conference call:

```
[C#]

RequestCompleteConference requestCompleteConference =
        RequestCompleteConference.Create(
        thisDn,
        connId,
        secondConnId);
response = tServerProtocol.Request(requestCompleteConference);
```

If the completion request is successful, you will receive `EventReleased, EventRetrieved, EventPartyAdded`, and `EventAttachedDataChanged` messages.

Transferring a Call

After you make or answer a call, you may also want to transfer that call. Here is how to perform an ordinary two-step transfer.

To start off, you need to initiate a transfer, supplying your own DN, the connection ID of the existing call, and the DN of the party you want to transfer the call to.

```
[C#]

RequestInitiateTransfer requestInitiateTransfer =
```

```
        RequestInitiateTransfer.Create(
        thisDn,
        connId,
        otherDn);
IMessage response = tServerProtocol.Request(requestInitiateTransfer);
```

Tip

In a real telephony application, the events you would receive in response to the kinds of transfer requests shown here could also be generated by other requests. For example, you might receive an `EventDialing` or an `EventEstablished` in response to a `RequestMakeCall` or `RequestInitiateConference`. Because of this, a real-world application will need to keep track of the requests that initiate these events in order to interpret them correctly.

If the initiate request is successful, you will receive an `EventDialing` message and an `EventHeld` message. When the party you want to transfer to picks up the call, you will also receive an `EventEstablished` message.

Now you need to complete the transfer.

When you received the `EventDialing` message from the `RequestInitiateTransfer`, you were given a new connection ID associated with the party you want to transfer the call to. You will need that connection ID, in addition to your own DN and the original connection ID, in order to complete the transfer:

```
[C#]

RequestCompleteTransfer requestCompleteTransfer =
        RequestCompleteTransfer.Create(
        thisDn,
        connId,
        secondConnId);
response = tServerProtocol.Request(requestCompleteTransfer);
```

If the completion request is successful, you will receive two `EventReleased` messages and you will no longer be a party to the call.

Closing the Connection

Finally, when you are finished communicating with the T-Server, you should close the connection to minimize resource utilization:

```
[C#]

tServerProtocol.Close();
```

List of TLib Functions

The following table provides a convenient list of TLib Functions that are available.

TAgentLogin	TEventGetStringAttr	TNetworkMerge	TSendEvent
TAgentLogout	TFreeEvent	TNetworkPrivateService	TSendEventEx
TAgentSetIdleReason	TGetAccessNumber	TNetworkReconnect	TSendUserEvent
TAgentSetNotReady	TGetMessageTypeName	TNetworkSingleStepTransfer	TSetCallAttributes
TAgentSetReady	TGetReferenceID	TNetworkTransfer	TSetDNDOff
TAlternateCall	TGetRouteTypeNames	TOpenServer	TSetDNDOn
TAnswerCall	TGetTreatmentTypeNames	TOpenServerEx	TSetInputMask
TApplyTreatment	TGetXCaps	TOpenServerX	TSetMessageWaitingOff
TAttachUserData	TGiveMusicTreatment	TOpenVoiceFile	TSetMessageWaitingOn
TCallCancelForward	TGiveRingBackTreatment	TPlayVoice	TSetMuteOff
TCallSetForward	TGiveSilenceTreatment	TPrivateService	TSetMuteOn
TCancelMonitoring	THoldCall	TQueryAddress	TSetParamHA
TCancelReqGetAccessNumber	TInitiateConference	TQueryCall	TSetRefIDLimit
TClearCall	TInitiateTransfer	TQueryLocation	TSetReferenceID
TCloseServer	TLibSetCompatibMode	TQueryServer	TSetSocketChangeCallback
TCloseVoiceFile	TListenDisconnect	TQuerySwitch	TSingleStepConference
TCollectDigits	TListenReconnect	TReconnectCall	TSingleStepTransfer
TCompleteConference	TLoginMailBox	TRedirectCall	TSockInfoStructure
TCompleteTransfer	TLogoutMailBox	TRegisterAddress	TSyncIsSet
TCopyEvent	TMakeCall	TReleaseCall	TSyncSetSelectMask
TDeleteAllUserData	TMakePredictiveCall	TReserveAgent	TUnregisterAddress
TDeleteFromConference	TMergeCalls	TRetrieveCall	TUpdateUserData
TDeleteUserData	TMonitorNextCall	TRouteCall	TXCapsSupported
TDispatch	TMuteTransfer	TScanServer	connid_to_decimal
TEventGetConnID	TNetworkAlternate	TScanServerEx	connid_to_str
TEventGetIntAttr	TNetworkConsult	TSendDTMF	decimal_to_connid
			str_to_connid

List of TLib Datatypes

The following table provides a convenient list of TLib Datatypes that are available.

Important

The names of most of these datatypes start with a "T", but the Voice Platform SDK uses names that do not contain an initial "T". For example, the TRegisterMode datatype mentioned in this section is known to the Voice Platform SDK as `RegisterMode`.

AddressStatusInfoType	TCallState	TKVResult	TRegisterMode
AssociationInfoType	TCallType	TKVType	TReliability
MsgWaitingInfoType	TClearFlag	TLocationInfoType	TRemoteParty
TAddressInfoStatus	TConnectionID	TMakeCallType	TRouteType
TAddressInfoType	TControlMode	TMediaType	TScanServerMode
TAddressType	TDNRole	TMergeType	TServer
TAgentID	TDirectoryNumber	TMessageType	TServerRole
TAgentPassword	TEvent	TMonitorNextCallType	TSetOpType
TAgentType	TEventMask	TNetworkCallState	TSwitchInfoType
TAgentWorkMode	TFile	TNetworkDestState	TTime
TAttribute	TForwardMode	TNetworkPartyRole	TTimeStamp
TCallHistoryInfo	TInterruptFlag	TOpenMode	TTreatmentType
TCallID	TKVList	TPartyState	TXCaps
TCallInfoType	TKVPair	TPrivateMsgType	TXRouteType

List of TLib Unstructured Data

The following table provides a convenient list of TLib Unstructured Data functions that are available.

These functions deal exclusively with transaction-related user data on the client side and allow you to work with all three categories of unstructured data: *User Data*, *Extensions*, and *Reasons*. None of these functions generate any requests to T-Server. The result of the function execution is confirmed by the value that the function returns.

TKVListAddBinary	TKVListCreate	TKVListGetListValue	TKVListNextPair
TKVListAddInt	TKVListDeleteAll	TKVListGetStringValue	TKVListPrint
TKVListAddList	TKVListDeletePair	TKVListGetUnicodeValue	TKVListStringValue
TKVListAddString	TKVListDup	TKVListInitScanLoop	TKVListType
TKVListAddUnicode	TKVListFree	TKVListIntValue	TKVListUnicodeValue
TKVListBinaryLength	TKVListGetBinaryValue	TKVListKey	TVKListGetPair
TKVListBinaryValue	TKVListGetIntValue	TKVListListValue	

Configuration

You can use the Configuration Platform SDK to write Java or .NET applications that access and update information from the Genesys Configuration Layer. These applications can range from the simple to the advanced.

This article shows how to implement the basic functions you will need to write a simple Configuration Layer application.

Once you have reviewed the information in this document, you should familiarize yourself with Configuration Layer Objects. Since the Configuration Platform SDK uses these objects for nearly everything it does, you will need to understand them before you start using this SDK.

> **Tip**
> The Platform SDK includes the Configuration Object Model Application Block, which is a high-performance component you can use to query on, and to create, update, and delete, Configuration Layer objects. Genesys recommends that you use this application block for most of the work you do with Configuration Layer objects.

When you are ready to write more complicated applications, take a look at the classes and methods described in the Platform SDK API Reference.

Java

Setting Up a ConfServerProtocol Object

The first thing you need to do to use the Configuration Platform SDK is instantiate a `ConfServerProtocol` object. To do that, you must supply information about the Configuration Server you want to connect with. This example uses the URI of the Configuration Server, but you can also use the server's name, host, and port information:

```
[Java]

ConfServerProtocol confServerProtocol =
        new ConfServerProtocol(
                new Endpoint(
                        confServerUri));
```

Configuration Server needs some additional information in order to create a successful connection. This information includes the type of client you wish to create, your client's name, and your user name and password:

```
[Java]

confServerProtocol.setClientApplicationType(CfgAppType.CFGSCE.asInteger());
confServerProtocol.setClientName("default");
confServerProtocol.setUserName(userName);
confServerProtocol.setUserPassword(password);
```

After instantiating the `ConfServerProtocol` object, you need to open the connection to the Configuration Server:

```
[Java]

confServerProtocol.open();
```

Creating a Query

Now that you have opened a connection, you can create a query and send it to Configuration Server. Let's say that you want to get information about a particular agent. To do this, you will need to supply the agent's user name using a *filter key*.

The filter key tells Configuration Server to narrow your query to a specific agent, rather than retrieving information about all of the persons in your contact center:

```
[Java]

KeyValueCollection filterKey = new KeyValueCollection();
filterKey.addObject("user_name", userName);
```

You can find the names of the filter keys for Person objects by looking in the *Filter Keys* section of the CfgPerson entry.

> **Tip**
> A similar reference page is available for each Configuration Layer object.

Now you are ready to create the request.

As you may know, Configuration Server considers agents to be objects of type CfgPerson. So you will need to create a request for information about a `Person` who has the user name you specified in the filter key:

```
[Java]

CfgObjectType objectType = CfgObjectType.CFGPerson;
int intPerson = objectType.asInteger();
RequestReadObjects requestReadObjects =
        RequestReadObjects.create(
                intPerson,
                filterKey);
```

> **Important**
>
> While the Configuration Layer supports the full character set in defining object names, using certain characters can cause problems in the behavior of some Genesys applications. Avoid using spaces, dashes, periods, or special characters in object names. Consider using underscores where you might normally use spaces or dashes.

After you have created your request, you can send it to Configuration Server, as shown here:

```
[Java]

confServerProtocol.send(requestReadObjects);
```

If the request is successful, you will receive an `EventObjectsRead` message.

> **Tip**
>
> When you send a RequestReadObjects message, Configuration Server may send more than one `EventObjectsRead` messages in response, depending on whether there is too much data to be handled by a single `EventObjectsRead`. Once you have received all of the `EventObjectsRead` messages, Configuration Server will also send an `EventObjectsSent`, which confirms that it has completed your request. For more information, refer to the article on event handling.

Interpreting the Response

The information you asked for is returned by invoking the `getConfObject` method of the `EventObjectsRead` message. This method returns an `org.w3c.dom.Document` representation of the object.

Here is a sample of how you might print the XML document:

```
[Java]

EventObjectsRead objectsRead =
        (EventObjectsRead) theMessage;
System.out.println(theMessage.messageName());
System.out.println("There are "
                + objectsRead.getObjectTotalCount() + " objects of this type.");
Document resultDocument =
        (Document) objectsRead.getConfObject();

... Add code to parse and print...
```

And this is what the XML document might look like:

```
<ConfData>
  <CfgPerson>
    <DBID value="105"/>
```

```
    <tenantDBID value="101"/>
    <lastName value="agent1"/>
    <firstName value="Agent"/>
    <employeeID value="agent1"/>
    <userName value="agent1"/>
    <password value="204904E461002B28511D5880E1C36A0F"/>
    <isAgent value="2"/>
    <CfgAgentInfo>
      <placeDBID value="102"/>
      <skillLevels>
        <CfgSkillLevel>
          <skillDBID value="101"/>
          <level value="9"/>
        </CfgSkillLevel>
      </skillLevels>
      <agentLogins>
        <CfgAgentLoginInfo>
          <agentLoginDBID value="103"/>
          <wrapupTime value="0"/>
        </CfgAgentLoginInfo>
      </agentLogins>
      <capacityRuleDBID value="127"/>
    </CfgAgentInfo>
    <isAdmin value="1"/>
    <state value="1"/>
    <userProperties>
      <list_pair key="desktop-redial">
        <str_pair key="phone-number0" value="5551212"/>
        <str_pair key="phone-number1" value=""/>
        <str_pair key="phone-number2" value=""/>
        <str_pair key="phone-number3" value=""/>
        <str_pair key="phone-number4" value=""/>
        <str_pair key="phone-number5" value=""/>
        <str_pair key="phone-number6" value=""/>
        <str_pair key="phone-number7" value=""/>
        <str_pair key="phone-number8" value=""/>
        <str_pair key="phone-number9" value=""/>
      </list_pair>
      <list_pair key="multimedia">
        <str_pair key="last-media-logged"
            value="voice,email"/>
      </list_pair>
    </userProperties>
    <emailAddress value="agent1@techpubs3"/>
  </CfgPerson>
</ConfData>
```

This XML document contains information about a `Person`. To interpret the information contained in the document, look at the *Parameters* section for CfgPerson.

If you compare the elements in this XML document to the CfgPerson entry, you can see that some of them contain information that is explained in detail in another entry. For example, the `CfgAgentInfo` element contains information that is described in the CfgAgentInfo entry. Similarly, the `CfgAgentLoginInfo` element contains information described in the CfgAgentLoginInfo entry.

Updating an Object

You can update a Configuration Layer object by passing in an XML `Document` containing the appropriate information about that object:

```
[Java]

CfgObjectType objectType = CfgObjectType.CFGPerson;
int intPerson = objectType.asInteger();
RequestUpdateObject requestUpdateObject =
        RequestUpdateObject.create(
                intPerson,
                xmlDocument);
```

Creating a New Object

You can also create a new Configuration Layer object by sending an XML `Document` to Configuration Server, as shown here:

```
[Java]

CfgObjectType objectType = CfgObjectType.CFGPerson;
int intPerson = objectType.asInteger();
RequestCreateObject requestCreateObject =
        RequestCreateObject.create(
                intPerson,
                xmlDocument);
```

Closing the Connection

Finally, when you are finished communicating with the Configuration Server, you should close the connection, in order to minimize resource utilization:

```
[Java]

confServerProtocol.close();
```

Working with Delta Objects

When using the Configuration Platform SDK to change attribute values of a configuration object, it is important to understand how "delta structures" work.

A delta structure contains values for each attribute in the configuration object. When a change is requested, a delta object is created that contains values for each attribute. Delta values are initialized to either zero (for integer values) or a null string - defaults that indicate no change should be made for that attribute. To change attributes of a configuration object, you first set the delta value for that attribute and then send the request to Configuration Server to be processed. Only attribute values that are changing should be specified in the delta structure for that object.

Any attributes with a delta value set to zero are left unchanged, so there are two special cases to remember when updating integer values in a configuration object:

- leaving the integer as 0 (zero) means that attribute does not change;
- setting a delta value to the current value of the configuration object attribute will change that attribute value to zero.

For example, if an Agent skill level is currently set to 5, then the following table illustrates the effect of various delta structure values:

Initial Attribute Value	Delta Structure Value	Updated Attribute Value	Comment
5	3	3	Setting the delta structure value to a non-zero integer will change the attribute to that value.
5	0	5	Leaving the delta structure value as zero will leave the attribute unchanged.
5	5	0	Setting the delta structure value to the current attribute value will change the attribute to zero.

Requests sent by SOAP clients and formed in an XML format do not use delta structures, because these types of request do not require all attributes to be present. The COM application block (which is shipped with the Platform SDKs) also does not use delta objects, as shown in the following code snippet:

```
[Java]

//retrieve an agent that has a single skill, with skill level set to 5
CfgPersonQuery query = new CfgPersonQuery();
query.setUserName("userName");
CfgPerson person = confService.retrieveObject(CfgPerson.class, query);

//Setting the skill level to 5 again will NOT result in a change in skill level  (ie: it will
remain 5).
((List<CfgSkillLevel>)person.getAgentInfo().getSkillLevels()).get(0).setLevel(5);
person.save();

//Setting the skill level to 0 will actually change the current skill level value.
((List<CfgSkillLevel>)person.getAgentInfo().getSkillLevels()).get(0).setLevel(0);
person.save();
```

.NET

Setting Up a ConfServerProtocol Object

The first thing you need to do to use the Configuration Platform SDK is instantiate a `ConfServerProtocol` object. To do that, you must supply information about the Configuration Server you want to connect with. This example uses the URI of the Configuration Server, but you can also use the server's name, host, and port information:

```
[C#]

ConfServerProtocol confServerProtocol =
        new ConfServerProtocol(
                new Endpoint(
                        confServerUri));
```

Configuration Server needs some additional information in order to create a successful connection. This information includes the type of client you wish to create, your client's name, and your user name and password:

```
[C#]

confServerProtocol.ClientApplicationType = (int) CfgAppType.CFGSCE;
confServerProtocol.ClientName = clientName;
confServerProtocol.UserName = userName;
confServerProtocol.UserPassword = password;
```

After instantiating the `ConfServerProtocol` object, you need to open the connection to the Configuration Server:

```
[C#]

confServerProtocol.Open();
```

Creating a Query

Now that you have opened a connection, you can create a query and send it to Configuration Server. Let's say that you want to get information about a particular agent. To do this, you will need to supply the agent's user name using a filter key.

The filter key tells Configuration Server to narrow your query to a specific agent, rather than retrieving information about all of the persons in your contact center:

```
[C#]

KeyValueCollection filterKey = new KeyValueCollection();
filterKey.Add("user_name", userName);
```

You can find the names of the filter keys for Person objects by looking in the Filter Keys section of the CfgPerson entry in the Configuration Objects section of this API reference. This section has a similar reference page for each Configuration Layer object.

Now you are ready to create the request.

As you may know, Configuration Server considers agents to be objects of type CfgPerson. So you will need to create a request for information about a Person who has the user name you specified in the filter key:

```
[C#]

RequestReadObjects requestReadObjects =
        RequestReadObjects.Create(
                (int) CfgObjectType.CFGPerson,
                filterKey);
```

Important

While the Configuration Layer supports the full character set in defining object names, using certain characters can cause problems in the behavior of other Genesys applications. Avoid spaces, dashes, periods, or special characters in object names. Consider using underscores where you might normally use spaces or dashes.

After you have created your request, you can send it to Configuration Server, as shown here:

```
[C#]

confServerProtocol.Send(requestReadObjects);
```

If the request is successful, you will receive an `EventObjectsRead` message.

Tip

When you send a `RequestReadObjects` message, Configuration Server may send more than one `EventObjectsRead` messages in response, depending on whether there is too much data to be handled by a single `EventObjectsRead`. Once you have received all of the `EventObjectsRead` messages, Configuration Server will also send an `EventObjectsSent`, which confirms that it has completed your request. For more information, refer to the article on event handling at the beginning of this API Reference.

Interpreting the Response

The information you asked for is returned in the `ConfObject` property of the `EventObjectsRead` message.

Here is a sample of how you might print the XML document:

```
[C#]

EventObjectsRead objectsRead = theMessage;

StringBuilder xmlAsText = new StringBuilder();
XmlWriterSettings xmlSettings = new XmlWriterSettings();
```

```
xmlSettings.Indent = true;

using (XmlWriter xmlWriter =
        XmlWriter.Create(xmlAsText, xmlSettings))
{
        XDocument resultDocument = objectsRead.ConfObject;
        resultDocument.WriteTo(xmlWriter);
}

Console.WriteLine("This is the response:\n"
                + xmlAsText.ToString() + "\n\n");
```

And this is what the XML document might look like:


```
<ConfData>
  <CfgPerson>
    <DBID value="105"/>
    <tenantDBID value="101"/>
    <lastName value="agent1"/>
    <firstName value="Agent"/>
    <employeeID value="agent1"/>
    <userName value="agent1"/>
    <password value="204904E461002B28511D5880E1C36A0F"/>
    <isAgent value="2"/>
    <CfgAgentInfo>
      <placeDBID value="102"/>
      <skillLevels>
        <CfgSkillLevel>
          <skillDBID value="101"/>
          <level value="9"/>
        </CfgSkillLevel>
      </skillLevels>
      <agentLogins>
        <CfgAgentLoginInfo>
          <agentLoginDBID value="103"/>
          <wrapupTime value="0"/>
        </CfgAgentLoginInfo>
      </agentLogins>
      <capacityRuleDBID value="127"/>
    </CfgAgentInfo>
    <isAdmin value="1"/>
    <state value="1"/>
    <userProperties>
      <list_pair key="desktop-redial">
        <str_pair key="phone-number0" value="5551212"/>
        <str_pair key="phone-number1" value=""/>
        <str_pair key="phone-number2" value=""/>
        <str_pair key="phone-number3" value=""/>
        <str_pair key="phone-number4" value=""/>
        <str_pair key="phone-number5" value=""/>
        <str_pair key="phone-number6" value=""/>
        <str_pair key="phone-number7" value=""/>
        <str_pair key="phone-number8" value=""/>
        <str_pair key="phone-number9" value=""/>
      </list_pair>
      <list_pair key="multimedia">
        <str_pair key="last-media-logged"
            value="voice,email"/>
      </list_pair>
    </userProperties>
    <emailAddress value="agent1@techpubs3"/>
  </CfgPerson>
</ConfData>
```

This XML document contains information about a Person. To interpret the information contained in the document, look at the Parameters section of the CfgPerson entry in the list of Configuration Objects.

If you compare the elements in this XML document to the CfgPerson entry, you can see that some of them contain information that is explained in detail in another entry. For example, the `CfgAgentInfo` element contains information that is described in the CfgAgentInfo entry. Similarly, the `CfgAgentLoginInfo` element contains information described in the CfgAgentLoginInfo entry.

Updating an Object

You can update a Configuration Layer object by passing in an XML document (of type `XDocument`) containing the appropriate information about that object:

```
[C#]

RequestUpdateObject requestUpdateObject =
        RequestUpdateObject.Create(
                (int) CfgObjectType.CFGPerson,
                xDocument);
```

Creating a New Object

You can also create a new Configuration Layer object by sending an XML Document (of type `XDocument`) to Configuration Server, as shown here:

```
[C#]

RequestCreateObject requestCreateObject =
        RequestCreateObject.Create(
                (int) CfgObjectType.CFGPerson,
                xDocument);
```

Closing the Connection

Finally, when you are finished communicating with the Configuration Server, you should close the connection, in order to minimize resource utilization:

```
[C#]

confServerProtocol.Close();
```

Working with Delta Objects

When using the Configuration Platform SDK to change attribute values of a configuration object, it is important to understand how "delta structures" work.

A delta structure contains values for each attribute in the configuration object. When a change is requested, a delta object is created that contains values for each attribute. Delta values are initialized to either zero (for integer values) or a null string - defaults that indicate no change should be made for that attribute. To change attributes of a configuration object, you first set the delta value for that attribute and then send the request to Configuration Server to be processed. Only attribute values that are changing should be specified in the delta structure for that object.

Any attributes with a delta value set to zero are left unchanged, so there are two special cases to remember when updating integer values in a configuration object:

- leaving the integer as 0 (zero) means that attribute does not change;
- setting a delta value to the current value of the configuration object attribute will change that attribute value to zero.

For example, if an Agent skill level is currently set to 5, then the following table illustrates the effect of various delta structure values:

Initial Attribute Value	Delta Structure Value	Updated Attribute Value	Comment
5	3	3	Setting the delta structure value to a non-zero integer will change the attribute to that value.
5	0	5	Leaving the delta structure value as zero will leave the attribute unchanged.
5	5	0	Setting the delta structure value to the current attribute value will change the attribute to zero.

Note that requests sent by SOAP clients and formed in an XML format do not use delta structures, because these types of request do not require all attributes to be present. The COM application block (which is shipped with the Platform SDKs) also does not use delta objects, as shown in the following code snippet:

```
[C#]

//retrieve a particular agent whose last name is "Jones"
CfgPersonQuery query = new CfgPersonQuery();
query.UserName = "userName";
query.LastName = "Jones";
CfgPerson person = myConfService.RetrieveObject<CfgPerson>(query);

//Setting the last name to the same value will NOT result in a change
person.LastName = "Jones";
person.Save();

//Setting the last name to a different value will change the actual value
person.LastName = "Smith";
person.Save();
```

Connecting Using UTF-8 Character Encoding

Genesys Configuration Server 8.1.2 added the ability to be configured to support multiple languages at a same time using UTF-8 encoding. Once Configuration Server is installed, configured and started in multilingual (UTF-8) mode it cannot be switched to regular mode. If Configuration Server is installed and started in normal mode, then it cannot be switched to multilingual (UTF-8) mode later.

One known issue is that the UTF-enabled protocol breaks backward compatibility, so users must add their own code for connection reconfiguration. The following samples describe connection scenarios with Platform SDK:

Scenario 1:

Configuration Server is release 8.1.2 or later and is NOT configured as multilingual (without UTF-8 transport), or is an earlier version without support for the UTF-8 feature.

In this scenario, Platform SDK connections can be created in the usual way.

Scenario 2:

Configuration Server is release 8.1.2 or later and configured as multilingual (with UTF-8 transport), with:

A) Platform SDK release 8.1.3 in use.

Reconfiguration for encoding is automatically handled by Platform SDK as described in the section below – no user action is required.

B) Platform SDK release 8.1.1 or 8.1.2 in use.

Platform SDK provides information that Configuration Server is UTF-8, so, the connection can be reopened using new connection configuration with following user code.

```
PropertyConfiguration config = new PropertyConfiguration();
config.setUseAddp(true);
config.setAddpClientTimeout(11);
config.setAddpServerTimeout(21);

ConfServerProtocol protocol = new ConfServerProtocol(new Endpoint(name, host, port, config));
protocol.setClientName(clientName);
protocol.setClientApplicationType(clientType.ordinal());
protocol.setUserName(username);
protocol.setUserPassword(password);

protocol.open();

Integer cfgServerEncoding = protocol.getServerContext().getServerEncoding();
if (cfgServerEncoding != null && cfgServerEncoding.intValue() == 1) {
   protocol.close();
   config.setStringsEncoding("UTF-8");
   protocol.setEndpoint(new Endpoint(name, host, port, config));
```

```
    protocol.open();
  }
```

It may be more comfortable to move the flag value evaluation to a separated method where a temporary `ConfServerProtocol` instance may be created - especially in the case of `ChannelListeners` usage, messages handlers, etc.

> **Important**
>
> This is not the best solution for wide usage. The `ServerEncoding` value evaluation method may fail if non-ASCII symbols are found inside the username or password, which may lead to a handshake procedure error such as "invalid username/password". This issue may be resolved with an additional test connection retry with UTF-8 enabled, but this workaround is not a best practice solution.

C) Platform SDK release 8.0.1 through 8.1.1 in use:

Platform SDK does NOT indicate whether Configuration Server is using UTF-8 mode or not, so user application should take care to evaluate this information (or have it defined by the design or configuration of the application).

In this case we have no `protocol.getServerContext().getServerEncoding()`, but we are able to configure the connection for Unicode usage.

It may be recommended to add one more property to the application configuration/parameters (along with the existing Configuration Server host and port) such as a boolean "isCSUTF8" value.

```
PropertyConfiguration config = new PropertyConfiguration();
if (isCSUTF8) {
   config.setOption(Connection.STR_ATTR_ENCODING_NAME_KEY, "UTF-8");
}

ConfServerProtocol protocol = new ConfServerProtocol(new Endpoint(name, host, port, config));
protocol.setClientName(clientName);
protocol.setClientApplicationType(clientType.ordinal());
protocol.setUserName(username);
protocol.setUserPassword(password);

protocol.open();
```

D) Platform SDK release of 8.0.0 or earlier in use.

No support if provided for string encoding of connection configuration options. The only way is to use this feature is to upgrade your release of Platform SDK.

Automatic UTF-8 Character Encoding Set Up on Handshake

Starting in Platform SDK Release 8.1.3 (which incorporates Configuration Protocol Release 3.79), support for UTF-8 encoding can be automatically detected.

The process for this features is described here:

1. The first handshake message, `EventProtocolVersion`, now includes the extra `ServerEncoding` attribute. If this attribute is 1 then Platform SDK updates string encoding for that connection to the server as UTF-8.
2. The next message from the client requests authentication from the server. These messages (`RequestRegisterClient` or `RequestRegisterClient2`) have been expanded with the `ClientEncoding` attribute, which must have the same value as the `ServerEncoding` attribute received previously.
3. After the handshake is complete, string encoding for this channel may be different from the string encoding specified in the original configuration parameters. You can access the current value through `Endpoint.GetConfiguration()` of the `ConfServerProtocol` instance.

Change Password On Next Login

An example of the usual Configuration Server connection open scenario is provided below:

```
ConfServerProtocol protocol = new ConfServerProtocol(new Endpoint("cfgsrv", csHost, csPort));
protocol.setClientName(clientAppName);
protocol.setClientApplicationType(clientAppType.ordinal());
protocol.setUserName(userName);
protocol.setUserPassword(userPasswd);
protocol.open();
```

When the user has enabled the Change Password on Next Login feature, `protocol.open()` throws `ChangePasswordException`.

So the resulting code may look like:

```
ConfServerProtocol protocol = new ConfServerProtocol(new Endpoint("cfgsrv", csHost, csPort));
protocol.setClientName(clientAppName);
protocol.setClientApplicationType(clientAppType.ordinal());
protocol.setUserName(userName);
protocol.setUserPassword(userPasswd);
try {
  protocol.open();
} catch (ChangePasswordException e) {
  String newPasswd = ...; // obtain new user password
  protocol.useChangePasswordRegistration(newPasswd);
  protocol.open();
}
```

After a successful open procedure, the new password value will be accepted, so `protocol.getUserPassword()` will be equal to newPasswd value specified.

Getting the Last Login Info

> **Tip**
> The appropriate Configuration Server version is required to use this feature, and so is the correct security configuration. For details, refer to Chapter 11 (Last Logged In Display) in Genesys 8.0 Security Deployment Guide.

Configuration Server provides last login information during the user authentication (handshake) procedure, and the Platform SDK Configuration Protocol provides it "as-is" in the form of a `KeyValueCollection`:

```
ConfServerProtocol.getServerContext().getLastLoginInfo();
```

An example of the resulting `KeyValueCollection` could look like:

```
KVList:
   'LAST_LOGIN_PERSON' [int] = 100
   'LAST_LOGIN_TIME' [int] = 1259161588
```

> **Tip**
> This information is only available while the connection is opened.

Note that "last login" is configured on Configuration Server through the `confserv.cfg` file:

```
[confserv]
...
last-login = true
last-login-synchronization = true
```

Platform SDK obtains the information using the `EventClientRegister` message:

```
2012-08-21 10:05:49,306 [New I/O client worker #4-4] DEBUG ns.protocol.DuplexChannel null -
Handling message:  'EventClientRegistered' (19) attributes:
        IATRCFG_SESSIONNUMBER [int] = 22
        IATRCFG_CFGSERVERDBID [int] = 99
        SATRCFG_PROTOCOL [str] = "CfgProtocol 5.1.3.54"
        IATRCFG_EXTERNALAUTH [int] = 0
        SATRCFG_PARAMETERS [KvListString] = KVList:
'LAST_LOGIN_PERSON' [int] = 1227
'LAST_LOGIN_TIME' [int] = 1345532749
'LAST_LOGIN_APPLICATION' [str] = "PSDK_CFGSCI"
        IATRCFG_BACKUPCFGSERVERDBID [int] = 0
        IATRCFG_UNSOLEVENTNUM [int] = 73770
        IATRCFG_CRYPTPASSW [int] = 1
        SATRCFG_SCHEMAVERSION [str] = "8.1.100.05"
        IATRCFG_REQUESTID [int] = 6
        SATRCFG_PROTOCOLEX [str] = "CfgProtocol 5.1.3.77"
```

There are two methods available in Platform SDK for retrieving last login details:

- `protocol.getServerContext().getLastLoginInfo()`
- `protocol.getServerContext().getCfgLastLogin()` (deprecated, not recommended for use)

If these methods return null, then you need to check whether Configuration Server gave the required info by looking in the debug logs for either Platform SDK or Configuration Server.

Using the Configuration Object Model Application Block

The Configuration Object Model Application Block is a reusable production-quality component that provides developers with a consistent and intuitive object model for working with Configuration Server objects. It has been designed using industry best practices and provided with source code so it can be used “as is,” extended, or tailored if you need to. Please see the License Agreement for details.

For information on the other application blocks that ship with the Genesys SDKs, consult Introducing the Platform SDK.

Java

Architecture and Design

The Configuration Platform SDK allows you to work with objects in the Genesys Configuration Layer by using the interface provided by Configuration Server. Unfortunately, this interface can be difficult to work with. For example, in order to update or create Configuration Layer objects, you have to use special “delta” objects that are distinct from the objects used to retrieve information about Configuration Layer objects.

The Configuration Object Model Application Block provides a consistent and intuitive object model that hides many of the complexities involved in working with Configuration Layer objects. This object model is implemented by way of an event subscription/delivery model, which hides key-value details of the current protocol, and is integrated with the rest of the object model.

The architecture of the Configuration Object Model Application Block consists of three functional components:

- Configuration Objects
- Configuration Service
- Query Objects
- Cache Objects

These components are shown below.

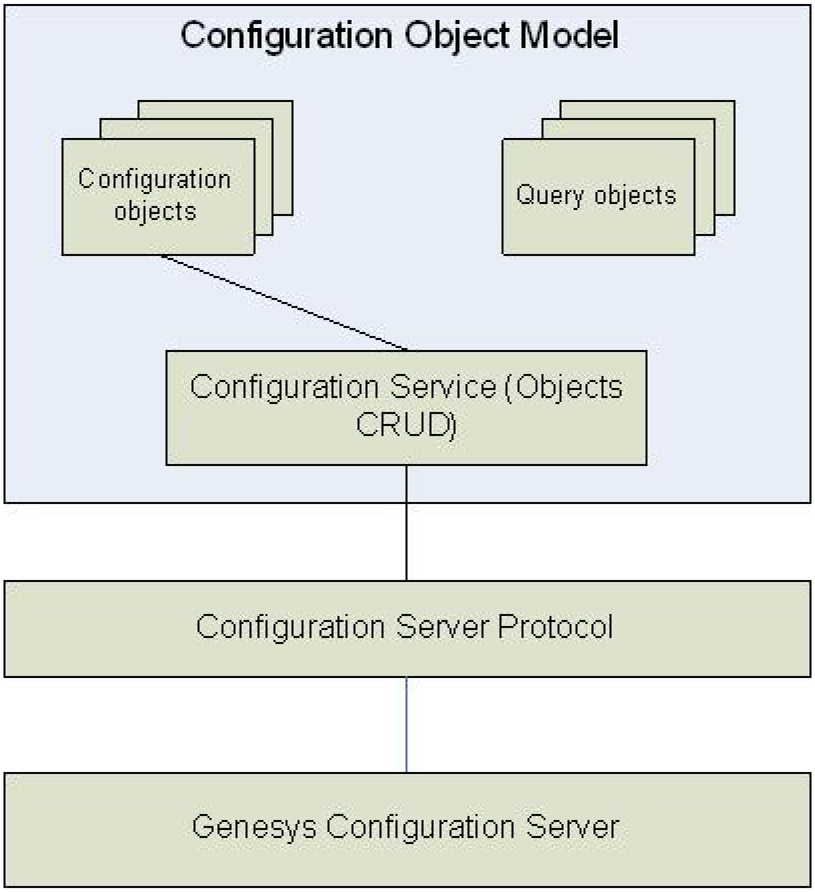

Configuration Objects

Classes and Structures

The Configuration Object Model Application Block supports two types of configuration objects:

- Classes, which can be retrieved directly from Configuration Server using queries.
- Structures, which only exist as properties of classes, and cannot be retrieved directly from Configuration Server.

Classes and structures are different in many ways, but in order to determine whether a given object is a class or a structure, all you need to do is check to see whether the object has a “DBID” property. Classes have this property, while structures do not.

Classes and structures are also different in the following ways:

- Each structure is a property of another class or structure, and therefore must have a “parent” class.
- Classes can be changed and saved to the Configuration Server and structures can only be saved through their “parent” classes.
- Clients can subscribe to events on changes in a class, but not in a structure. To retrieve events on

changes in a structure, clients have to subscribe to changes in its parent class.

Property Types

Both classes and structures have properties. Each property has its own getter and setter methods, and each property is an instance of one of the following types:

- *Simple* — A property that is represented by a value type. Configuration Server supports two types of simple properties - `string` and `integer`. For example, the `CfgPerson` object has `FirstName` and `LastName` properties, both of the `string` type.
- *KV-list* — Tree-like properties that are represented by the `KeyValueCollection` class in the Configuration Object Model. Examples of this property include `userProperties` and `CfgPerson`.
- *Structure* — A complex property that includes one or more properties. In the Configuration Object Model, structures are represented by instances of classes that are similar to configuration objects, but cannot be created directly. For example, in the `CfgPerson` class, its `AgentInfo` property contains *simple*, *kv-list* and other property types.
- *List of structures* — A property that represents more than one structure. In Configuration Object Model, lists of structures are represented by a generic type `IList<structure_type>`, so that the collection is typed, and clients can easily iterate through the collection.
- *Links to a single object* — In Configuration Server, these properties are stored as DBIDs of external objects. The Configuration Object Model automatically resolves these DBIDs into the real objects, which can be manipulated in the same way as the objects directly retrieved from Configuration Server. Links are initialized at the time of the initial request to one of its properties.

> **Tip**
>
> For each link, there are two ways to set the new value of a link. There is a setter method of the property, which uses an object reference to set a new value of a link. There is also a `Set...DBID` method, which uses an integer DBID value.

- *Links to multiple objects* — A property that contains more than one link. In the Configuration Object Model, lists of structures are represented by a generic type `IList<class_type>`, so that the collection is typed, and clients can easily iterate through the collection.

Creating Instances

One way to create an instance of an object in the Configuration Object Model is to invoke a `Retrieve...` method of a `ConfService` class. This set of methods returns instances of objects that already exist in Configuration Server.

To create a new object in Configuration Server, a client must create a new instance of a COM or "*detached*" object. The detached object does not correspond to any objects in Configuration Server until it is saved. The detached object is created using the regular Object-Oriented language object instantiation. For example, a new detached `CfgPerson` object is created using the following construction:

```
[Java]

CfgPerson person = new CfgPerson(confService);
```

An object instance can also be created by using links to external objects. The Component Object Model creates a new object instance whenever the link is called, or any of the properties of a linked object are called. For example, you can write:

```
[Java]

// person has already been retrieved from Configuration Server.
CfgTenant tenant = person.getTenant(); // this is a link to an external object. It is
initialized internally right now
CfgAddress address = tenant.getAddress();
```

Common Methods

Each configuration class contains the following methods:

- Generic `GetProperty(string propertyName)` — Retrieves the property value by its name.
- Generic `SetProperty(string propertyName)` — Sets the new value of the property by its name.
- `Save()` — Commits all changes previously made to the object to Configuration Server. If the object was created detached from Configuration Server and has never been saved, a new object is created in Configuration Server using the `RequestCreateObject` method. If the object has been saved or has been retrieved from Configuration Server, a delta-object, which contains all changes to the object, is formed and sent to Configuration Server by means of the `RequestUpdateObject` method.
- `Delete()` — Deletes the object from the Configuration Server Database.
- `Refresh()` — Retrieves the latest version of the object and refreshes the value of all its properties.

> **Tip**
> In this release, all configuration objects are “static,” which means that if the object changes in the Configuration Server, the instance of a class is not automatically changed in the Configuration Object Model. Clients must subscribe to the corresponding event and manually refresh the COM object in order for these changes to take effect.

Configuration Service

> **Tip**
> The `IConfService` interface was added to COM in release 8.0. All applications should now use this interface to work with the configuration service instead of the old `ConfService` class. This change is an example of how all COM types in the interface are now referred to by interface; for instance, if a method previously returned `CfgObject` it now returns `ICfgObject`. This is not compatible with existing code, but upgrading should not be difficult as the new interfaces support the same methods as the implementing types.

The Configuration Service (`IConfService`) interface provides services such as retrieval of objects and subscription to events from Configuration Server. Each connection to a Configuration Server (represented by a `ConfServerProtocol` class of Platform SDK) requires its own instance of the `IConfService` interface.

The protocol class should be created and initialized in the client code prior to `IConfService` initialization.

The `ConfServiceFactory` class is used to create the `IConfService`. This class uses the following syntax:

```
[Java]

IConfService service = ConfServiceFactory.CreateConfService(protocol);
```

Retrieving Objects

Objects can be retrieved from Configuration Service by using one of the following methods:

- `RetrieveObject` — Accepts a query that returns one object. If multiple objects are returned, an exception is thrown.
- `RetrieveMultipleObjects` — Accepts a query that returns one or more objects. A collection of objects is returned.

Each of the `Retrieve...` methods can be either specific (by using generic criteria entries, an object of a specified type is returned) or general (a general object is returned).

Handling Events

The following methods must be called before receiving events from Configuration Server:

1. Register - The application must register its callback by calling the `Register` method from the Configuration Service. This method supplies the client's filter, which enables the client to receive only requested events.
2. Subscribe - The application must subscribe to events from Configuration Server by calling the `Subscribe` method from the Configuration Service. This method provides a notification query object as a parameter.

   The `NotificationQuery` object determines whether the object (or set of objects) to which the client wants to subscribe has changed. The `NotificationQuery` object contains such parameters as `object type, object DBID` and `tenant DBID`.

   After calling the `Subscribe` method, Configuration Server starts sending events to the client. These events are objects, which contain information such as:

   - which object (ID and type) is affected
   - the type of event sent to the client
   - any additional information

   There are three types of events that the client might receive:

   - `ObjectCreated` — A new object has been added to Configuration Server.
   - `ObjectChanged` — Some of the object properties have been modified in Configuration Server.
   - `ObjectDeleted` — The object has been removed from Configuration Server.

Releasing a Configuration Service

Whenever a `ConfService` instance is no longer needed, the `ReleaseConfService` method can be used to remove it from the internal list.

[Java]

```
ConfServiceFactory.ReleaseConfService(service);
```

Query Objects

A query object is an instance of a class that contains information required for a successful query to a Configuration Server. This information includes an object type and its attributes (such as *name* and *tenant*), which are used in the search process.

The inheritance structure of configuration server queries is designed to allow for future expansion. The `CfgQuery` object is the base class for all query objects. Other classes extend `CfgQuery` to provide more specific functionality for different types of queries - for example, all filter-based queries use the `CfgFilterBasedQuery` class. This allows room for future query types (such as XPath) to implemented in this Application Block.

A list of currently available query types is provided below:

- CfgFilterBasedQuery — Contains mapped attribute name-value pairs, as well as the object type.

A special query class is supplied for each configuration object type, in order to facilitate the process of making queries to Configuration Server. For each searchable attribute, the query class has a property that can be set. All of these classes inherit attributes from the `CfgQuery` object, and can be supplied as parameters to the Retrieve... methods which are used to perform searches in Configuration Server.

Cache Objects

The cache functionality is intended to enhance the Configuration Object Model by allowing configuration objects to be stored locally, thereby minimizing requests to configuration server, as well as enhancing ease of use by providing automatic synchronization between locally stored objects and their server-side counterparts.

The cache functionality was designed with the following principles in mind:

- The cache functionality is designed to be extendable with custom implementations of provided interfaces and not via inheritance.
- The cache component is not designed to replicate the Configuration Server query engine or other Configuration Server functionality on the client side.
- Caching must be an optional feature. Work with Configuration Server should not be affected if caching is not used.

Use Cases

Analysis of use cases provides insight into the requirements for applications likely to require configuration cache functionality. The use cases described in the following table were selected for analysis in order to highlight different functional requirements. There are several possible actors which are referenced in the use cases. The actors are as follows:

- Application - Any application which uses the Configuration Object Model application block
- User - Human (or software) user who may perform actions upon objects in the configuration which are separate from the Application

Use Case	Description	Actor	Steps
PLACE OBJECT INTO CACHE	Place a configuration object into the configuration cache (note the object must have been saved — ie must have a DBID in order to exist in the cache).	Application	1. Application adds object to the cache
PLACE OBJECT INTO CACHE ON SAVE	Place a newly created configuration object into the configuration cache when it is saved.	Application	1. Application creates object 2. Application saves object 3. Configuration Object Model Application Block adds object to the cache
PLACE OBJECT INTO CACHE ON RETRIEVE	Allow for automatic insertion of configuration objects into the cache upon retrieval from configuration server.	Application	1. Application retrieves configuration object 2. Configuration Object Model Application Block retrieves the configuration object from the server 3. Configuration Object Model Application Block places the configuration object into the cache 4. Configuration Object Model Application Block returns the object to the application
OBJECT REMOVED IN CONFIGURATION SERVER	When configuration objects are deleted in the configuration server, the cache can delete the local representation of the object as well.	User	1. User deletes object in the Configuration Server 2. Cache removes

Use Case	Description	Actor	Steps
			corresponding local object upon receiving delete notification 3. Cache sends notification of object deletion to Application
SYNCHRONIZE OBJECT PROPERTIES WITH CONFIGURATION SERVER	When an object stored in the cache is updated in the Configuration Server the object must be updated locally as well.	User	1. User updates a configuration object 2. Cache receives notification about object update 3. Cache updates the object based on the received delta 4. Cache fires event informing any subscribers of object change
FIND OBJECT IN CACHE	The cache must support the ability to find a specific configuration object in the cache using object DBID and type as the criteria for the search.	Application	1. Application retrieves object from cache. 2. If object is in the cache, the cache returns the object. Otherwise the application is notified that the requested object is not in the cache.
ACCESS CACHED OBJECTS	The cache must provide its full object collection to the application.	Application	1. Application requests a complete list of objects from the cache. 2. The cache returns a collection of all cached objects.
RETRIEVE LINKED OBJECT FROM CACHE	If caching is turned on, object links which the Configuration Object Model currently resolves through lazy	Application	1. Application accesses a property which requires link resolution

Use Case	Description	Actor	Steps
	initialization (i.e. if a property linking to another object is accessed, we retrieve the referred-to object from configuration server) must be resolvable through cache access.		2. Configuration Object Model Application Block retrieves the linked object from configuration server and stores it in the cache before returning to the application 3. Application again accesses the property and this time the Configuration Object Model Application Block retrieves the object from the cache
PROVIDE CACHE TRANSPARENCY ON RETRIEVE	A cache search should be performed on attempt to retrieve an object from Configuration Server. If the requested object is found in the cache then the Configuration Object Model should return the cached object rather than accessing Configuration Server.		1. Application creates query to retrieve configuration object 2. Application executes query using the Configuration Object Model 3. Configuration Object Model Application Block searches the cache • If object present, return the object • If object not present, query configuration server for the object
CACHE SERIALIZATION	The cache should support serialization.	Application	1. Application provides a stream to the cache 2. The cache serializes itself into the stream in an XML format 3. Application restarts 4. Application provides

Use Case	Description	Actor	Steps
			the cache a stream of cache data in the same XML format as in step 2 5. Cache restores itself 6. Cache subscribes for updates on the restored objects

Implementation Overview

Two new interfaces for cache management have been added to the Configuration Object Model: the `IConfCache` interface and a default cache implementation (`DefaultConfCache`). Note that the `ConfCache` also implements the Subscriber interface from `MessageBroker` so that the user can subscribe to notifications from Configuration Server, as discussed in *Notification And Delta Handling*.

The `IConfCache` interface provides methods for basic functionality such as adding, updating, retrieving, and removing objects in the cache. It also includes a `Policy` property that defines cache behavior and affects method implementation. (For more details about policies, see *Cache Policy*).

The `DefaultConfCache` component provides a default implementation of the `IConfCache` interface. It serializes and deserializes cache objects using the XML format described in the *XML Format* section, below.

To enable and configure caching functionality, and to specify `ConfService` policy, there are three `CreateConfService` methods available from `ConfServiceFactory`. The original `CreateConfService` method (not shown here) creates a `ConfService` instance that uses the default policy and does not use caching.

```
[Java]

public static IConfService createConfService(Protocol protocol, boolean enableCaching)
```

This method creates an instance of a Configuration Service based on the specified protocol. If caching is enabled, the default caching policy will be used. If `enableCaching` is set to true, caching functionality will be turned on. If caching is disabled, all policy flags related to caching will be false.

```
[Java]

public static IConfService createConfService(Protocol protocol,
        IConfServicePolicy confServicePolicy, IConfCache cache)
```

This method creates a configuration service with the specified policy information. The created service will have caching enabled if a cache object (implementing the `IConfCache` interface) is passed as a parameter.

```
[Java]

public static IConfService createConfService(Protocol protocol,
        IConfServicePolicy confServicePolicy, IConfCachePolicy confCachePolicy)
```

This method creates a configuration service with the specified policy information. The created service will have caching enabled by default with the cache using the specified cache policy.

XML Format

The "Cache" node will be the root of the configuration cache XML, while "ConfData" is a child of the "Cache" node. The ConfData node contains a collection of XML representations for each configuration object in the cache. The XML format of each object is identical to that which is returned by the `ToXml` method supported by each the Configuration Object Model configuration object.

The "CacheConfiguration" element is a child of the "Cache" node. There can only be one instance of this node and it contains all cache configuration parameters, as follows:

- CONFIGURATIONSERVER NODE — There can be 1..n instances of this element. Each one will represent a configuration server for which the cache is applicable (a cache can be applicable to multiple configuration servers if they are working with the same database as in the case of a primary and backup configuration server pair). Each ConfigurationServer element will have a `URI` attribute specifying the unique URI identifying the Configuration Server, as well as a `Name` attribute specifying the name associated with the endpoint.

The example provided below shows a cache that is applicable for the configuration server at "server:2020" with some policy details specified. There are two objects in the cache for this example: a `CfgDN` and a `CfgService` object.

```
[XML]

<Cache>
  <CacheConfiguration>
    <ConfigurationServer name="serverName" uri="tcp://server:2020"/>
  </CacheConfiguration>
  <ConfData>
    <CfgDN>
      <DBID value="267" />
      <switchDBID value="111" />
      <tenantDBID value="1" />
      <type value="3" />
      <number value="1111" />
      <loginFlag value="1" />
      <registerAll value="2" />
      <groupDBID value="0" />
      <trunks value="0" />
      <routeType value="1" />
      <state value="1" />
      <name value="DNAlias" />
      <useOverride value="2" />
      <switchSpecificType value="1" />
      <siteDBID value="0" />
      <contractDBID value="0" />
      <accessNumbers />
      <userProperties />
    </CfgDN>

    <CfgService>
      <DBID value="102" />
      <name value="Solution1" />
      <type value="2" />
      <state value="1" />
      <solutionType value="1" />
      <components>
        <CfgSolutionComponent>
```

```
        <startupPriority value="3" />
        <isOptional value="2" />
        <appDBID value="153" />
      </CfgSolutionComponent>
    </components>
    <SCSDBID value="102" />
    <assignedTenantDBID value="101" />
    <version value="7.6.000.00" />
    <startupType value="2" />
    <userProperties />
    <componentDefinitions />
    <resources />
  </CfgService>
 </ConfData>
</Cache>
```

Cache Policy

The configuration cache can be assigned a policy represented by a `Policy` interface. A default implementation of the interface will be provided in the `DefaultConfCachePolicy` class.

The `IConfCache` interface interprets the policy as follows:

1. `CacheOnCreate` — When an object is created in the configuration server, the policy will be checked with the created object as the parameter. If the method returns true, the object will be added to the cache, if it is false, the object will not be added. Default implementation will always return false.
2. `RemoveOnDelete` — When an object is deleted in the configuration server, the policy will be checked with the deleted object as the parameter. If the method returns true, the object will be deleted in the cache, if it is false, the notification will be ignored. Default implementation will always return true.
3. `TrackUpdates` — When an object is updated in the configuration server, the policy will be checked with the current version of the object as the parameter. If the method returns true, the object will be updated with the received delta, if it is false, the notification will be ignored. Default implementation will always return true.
4. `ReturnCopies` — Determines whether the cache should return copies of objects when they are retrieved from the cache, or the original, cached versions. False by default.

IConfServicePolicy Interface

The `IConfServicePolicy` interface can be used to define the policy settings for the `ConfService`. Two default implementations are available:

1. `DefaultConfServicePolicy` contains the settings for a non-caching configuration service. That is, all of the cache-related policy flags will always return false.
2. `CachingConfServicePolicy` defines the default behavior for a configuration service with caching enabled. (Note that when referring to the "default" value below, we will be referring to this implementation.)

The policy interface settings are interpreted as follows:

- AttemptLinkResolutionThroughCache — Whenever a link resolution attempt is made, this policy will be checked for the type of object the link refers to. If this method returns true, the link resolution attempt will first be made through the cache. If the method returns false, or if the object has not been found in the cache, the server will queried. Default value is always true.

- `CacheOnRetrieve` — This method will be called for each object retrieved from the configuration. If the return value is "true" the object will be added to the cache. Default value is always true.
- `CacheOnSave` — This method will be called for each object that is being saved. If the return value is true, the object will be added to the cache. If the object is already in the cache, it will not be overwritten. Default value is always true.
- `ValidateBeforeSave` — This is a property from the `ConfService` which will be moved to the policy interface and is not related to caching. It is used to indicate whether property values are checked for valid values against the schema before a save attempt is made. Default value is true.
- `QueryCacheOnRetrieve` — This method will be called every time a retrieve operation is performed using a query. The `ConfService` will first check the cache for the existence of the requested configuration object. If the object exists, it will be returned and no configuration server request will be made. If there are no values returned, the `ConfService` will query the configuration server (see *Query Engine*). Default value is always false.
- `QueryCacheOnRetrieveMultiple` — This method will be called every time a retrieve multiple operation is performed. The `ConfService` will first execute the query against cache. If the returned object count is greater than 0 the found object collection will be returned and no configuration server request will be made. If there are no values returned, the `ConfService` will query the configuration server (see *Query Engine*). Default value is always false.

Note that the `RetrieveMultiple` operation is NOT implemented in the default query engine, so providing a policy where this method returns true will require a new query engine implementation.

Cache Extendability

Consistent with the design principles outlined above, the configuration cache is extendable via custom implementations of provided interfaces. The two areas of the cache which can be extended are the cache storage and the cache query engine.

Cache Storage

The storage interface defines the method by which objects are stored in the cache. When an instance of an implementing object is provided to the cache, the cache will store all cached objects in the storage component.

The default storage implementation stores cached objects using the object type and DBID as keys. Note that this means that objects in the cache are assumed to be from one configuration database. The default implementation is also thread safe using a reader/writer lock which allows for multiple concurrent readers and one writer. The storage methods are as follows:

- Add — Adds a new object to the storage. If object already exists in the storage, the default implementation thrown an exception.
- Update — Overwrites an existing object in the storage. If the object is not found in the storage, the default implementation creates a new version of the object.
- Remove — Removes an object from the storage.
- Retrieve — Retrieves an enumerable list of all objects in the storage (filtered by type), and possibly influenced by an optional helper parameter. Note that the helper parameter is not meant to provide querying logic — that should be done in the query engine. Because the query engine is to some degree dependent on the storage implementation, the helper parameter allows for some flexibility in the way stored objects are enumerated for the query engine. The default implementation can take a

`CfgObjectType` as a helper parameter.

- Clear — Removes all objects in the storage.

Query Engine

The query engine provides the ability to define the method by which objects are located in the cache.

Depending on the `IConfService` policy, `Retrieve` requests as well as link resolution can first be attempted through the cache. If the requested object is found in the cache, then that cached object is returned instead of sending a request to Configuration Server. If the object is not present in the cache, a request to Configuration Server is made.

A user-definable query engine module exists inside the cache to achieve this functionality. A query engine must implement the `IConfCacheQueryEngine` interface, which provides methods to retrieve objects (either individually, or as a list) and to test a query and determine if it can be executed.

If enabled by the policy, `IConfService` will attempt a query to its cache using the cache's query engine interface. If a result is returned, the `IConfService` will not query the Configuration Server. By following this contract, the Configuration Object Model user is then able to create a custom implementation of the `IConfCacheQueryEngine` with any extended search capabilities which may be missing from the simple default implementation.

Two implementations of the `IConfCacheQueryEngine` interface are provided in the Configuration Object Model, as described below:

- DEFAULTCONFCACHEQUERYENGINE CLASS - The `DefaultConfCacheQueryEngine` class is a default implementation of the `IConfCacheQueryEngine` interface.
- COMPOSITECONFCACHEQUERYENGINE CLASS - This class is a more advanced implementation of the query engine which allows child query engine modules to be registered in order to interpret different types of queries. It does not have a default query engine implementation, only the mechanism for working with multiple child query engines.

Notification and Delta Handling

The default configuration cache will implement the `Subscriber<ConfEvent>` interface which will allow the cache to be subscribed to receive configuration events. When a cache instance is associated with a Configuration Service, it will automatically be subscribed for configuration events from that service (note that if a custom cache implementation also implements this interface it will be subscribed for events as well). The way the cache is updated based on these notifications is determined by the cache policy.

In addition, a new filter class will be added in order to allow the subscriber to filter the cache events. The `ConfCacheFilter` will implement the MessageBroker's `Predicate` interface, allowing for the filter to be passed during registration for events via `SubscriptionService`. The `ConfCacheFilter`'s properties will specify the parameters by which the events will be filtered. Initially, the supported parameters will be object type, object DBID, and update type, allowing the user to filter events by one or a combination of these parameters assuming an AND relationship between the parameters specified.

Using the Application Block

Installing the Configuration Object Model Application Block

Before you install the Configuration Object Model Application Block, it is important to review the software requirements for using it.

Software Requirements

To work with the Configuration Object Model Application Block, you must ensure that your system meets the software requirements established in the Genesys Supported Operating Environment Reference Manual, as well as meeting the following minimum software requirements:

- JDK 1.6 or higher
- Genesys Configuration Platform SDK 8.0 or higher

Building the Configuration Object Model Application Block

To build the Configuration Object Model Application Block:

1. Open the `<Platform SDK Folder>\applicationblocks\com` folder.
2. Run either `build.bat` or `build.sh`, depending on your platform.

This will create the `comappblock.jar` file, located within the `<Platform SDK Folder>\applicationblocks\com\dist\lib` directory.

Using the QuickStart Application

The easiest way to start using the Configuration Object Model Application Block is to use the bundled QuickStart application. This application ships in the same folder as the application block.

Configuring the QuickStart Application

In order to use the QuickStart application, you will need to change some lines of code in the `quickstart.properties` file, located in the `<Platform SDK Folder>\applicationblocks\com\quickstart` directory. Change the following lines to point to your Configuration Server, and then save the updated file:

```
ConfServerUri = tcp://:

ConfServerUser =
ConfServerPassword =

ConfServerClientName = default
ConfServerClientType = CFGSCE
```

Building the QuickStart Application

1. Open the `<Platform SDK Folder>\applicationblocks\com\quickstart` folder.

2. Run either `build.bat` or `build.sh`, depending on your platform.

Running the QuickStart Application

1. Open a Command Prompt or Terminal window.
2. Navigate to the `<Platform SDK Folder>\applicationblocks\com\quickstart` directory.
3. Run either `quickstart.bat` or `quickstart.sh`, depending on your platform.

How to Properly Initialize the ConfService Instance

To work with Configuration Server, the `ConfService` instance needs `ConfServerProtocol`.

Platform SDK protocol connections allow users to manage connections, setup custom asynchronus `MessageHandler` objects, substitute message receivers, and subscribe for protocol messages and channel events. So, to maintain Platform SDK flexibility, the Configuration Object Model Application Block does not manage a `ConfServerProtocol` connection inside of the `ConfService` - this must be done by the user. Instead users may create a simple instance and initialize it with `WarmStandbyService`.

It is important to note that asycnhronous protocol events may be configured for delivery to a single destination, with only one `MessageHandler` or `MessageReceiver` for one protocol instance. Starting from Platform SDK release 8.1.1, `ConfService` may be initialized without use of legacy Message Broker Application Block. Starting from version 8.5, this is the only way to create `ConfService`.

If your application needs to receive asynchronous protocol messages from Configuration Server on the protocol instance where `ConfService` is initialized, that can be done using `ConfService.setUserMessageHandler(messageHandler)`.

Protocol Initialization

A `ConfServerProtocol` instance is required for the creation of `ConfService`. It should be initialized with an Endpoint and handshake properties, but without setting either `confServerProtocol.setMessageHandler()` or `confServerProtocol.setReceiver()`.

```
// Initialize ConfService:
PropertyConfiguration       config;
ConfServerProtocol          confServerProtocol;
IConfService                confService;

config = new PropertyConfiguration();
config.setUseAddp(true);
config.setAddpClientTimeout(15);

confServerProtocol = new ConfServerProtocol(new Endpoint("ConfigServer", csHost, csPort,
config));
confServerProtocol.setUserName(userName);
confServerProtocol.setUserPassword(password);
confServerProtocol.setClientName(clientName);
confServerProtocol.setClientApplicationType(clientType.ordinal());
```

> **Important**
>
> Do *not* open the protocol before `ConfService` is created. `ConfService` sets its own internal `MessageHandler`, and this operation can only be done on a closed channel.

ConfService Initialization

```
confService = ConfServiceFactory.createConfService(confServerProtocol);
confServerProtocol.open();
```

ConfService Shutdown

```
confServerProtocol.close();
ConfServiceFactory.releaseConfService(confService);
confService = null;
```

Application Components Usage Notes

Older releases of `ProtocolManagementService` do not support using `ConfService` without the Message Broker service - an exception raised when users try to create the `ConfService` object on a protocol instance initialized by the Protocol Manager Application Block. To migrate away from Protocol Manager Application Block usage, we recommend creating and configuring `ConfServerProtocol` without Protocol Manager Application Block usage, as shown above.

`MessageHandler` is not compatible with the deprecated `MessageReceiver`; it is only possible to use one of these components on a protocol instance. Specific to Platform SDK for Java is the limitation that one protocol instance may have only one instance of `MessageHandler`. So, if an application uses a custom `MessageHandler` on a protocol used for `ConfService`, then only one handler will be able to receive asynchronous protocol events.

If application overwrites the `ConfService` object after creation, then that service will be unable to receive Configuration Server notifications or to perform multiple objects reading operations - a timeout exception will occur. If there is a need to get those protocol messages separately from `ConfService` logic, it is possible to initialize custom `MessageHandler` with `confService.setUserMessageHandler(messageHandler)`.

Notes for Previous Releases of Platform SDK

'''[+] Platform SDK 8.1.0 Specific Notes'''

Platform SDK 8.1.0 included some improvements to the Message Broker Application Block.

There was a new `EventReceivingBrokerService` class that implements the receiver interface, which can be used as an external receiver for Platform SDK protocols. When this class is in use, protocol messages will be handled a little bit faster (compared to the older Message Broker service) with no redundant intermediate queue, and there is no additional thread sleeping/waiting.

```
EventReceivingBrokerService broker = new EventReceivingBrokerService();
broker.setInvoker(new SingleThreadInvoker("COMBrokerService-" + cfgsrvEndpointName));
```

```
ConfServerProtocol protocol = new ConfServerProtocol(endpoint);
protocol.setReceiver(broker);
protocol.setUserName(...);
protocol.set...();
protocol.open();

IConfService confService = ConfServiceFactory.createConfService(protocol, broker);
```

To shutdown the Configuration Object Model Application Block, you can use the following code:

```
protocol.close();
ConfServiceFactory.releaseConfService(confService);
```

'''[+] Platform SDK 8.0, 7.6 Specific Notes'''

In earlier releases of Platform SDK, the initialization logic could look like this:

```
ConfServerProtocol protocol = new ConfServerProtocol(endpoint);
protocol.setUserName(...);
protocol.set...();
protocol.open();

EventBrokerService broker = BrokerServiceFactory.CreateEventBroker(protocol);
IConfService confService = ConfServiceFactory.createConfService(protocol, broker);
```

If the protocol has an external receiver initialized (for example, with Protocol Manager usage), then the `EventBrokerService` should be initialized on that receiver instead of the protocol itself:

```
EventBrokerService broker =
BrokerServiceFactory.CreateEventBroker(protocolManager.getReceiver());
```

To shutdown the Configuration Object Model Application Block, you can use the following code:

```
protocol.close();
broker.dispose();
ConfServiceFactory.releaseConfService(confService);
```

> **Important**
> Legacy `EventBrokerService` objects need to be disposed on shutdown because they include an internal reading thread which should be stopped.

.NET

Architecture and Design

The Configuration Object Model Application Block provides a consistent and intuitive object model for working with Configuration Server objects, as well as a straightforward object model for queries with different filters. This Application Block hides the complexities of object creation and changing by means of "delta" objects. It also creates an event subscription/delivery model, which hides key-value details of the current protocol, and is integrated with the rest of the object model.

The architecture of the Configuration Object Model Application Block consists of three functional components:

- Configuration Objects
- Configuration Service
- Query Objects
- Cache Objects

These components are shown in the figure below.

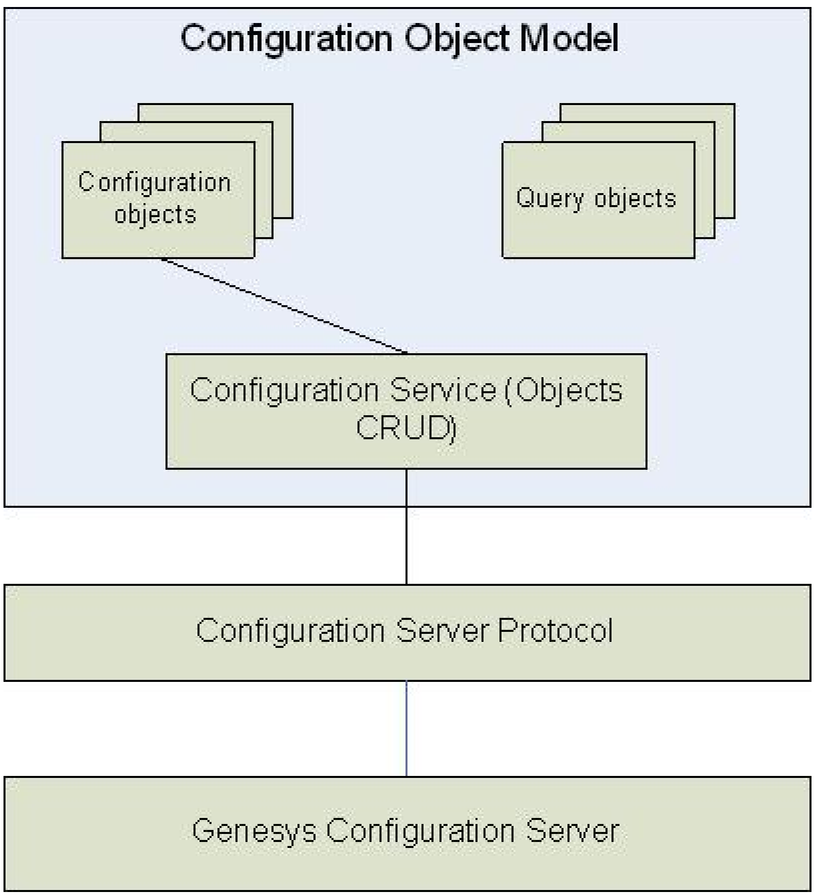

Configuration Objects

Classes and Structures

There are two types of configuration objects are supported by the Configuration Object Model

Application Block:

- Classes, which can be retrieved directly from Configuration Server using queries.
- Structures, which only exist as properties of classes, and cannot be retrieved directly from Configuration Server.

The main differences between classes and structures are as follows:

1. Each structure is a property of another class or structure, and therefore must have a "parent" class.
2. Classes can be changed and saved to the Configuration Server and structures can only be saved through their "parent" classes.
3. Clients can subscribe to events on changes in a class, but not in a structure. To retrieve events on changes in a structure, clients have to subscribe to changes in a parent class.

Property Types

Both classes and structures have properties. For each property, the object has getter and setter methods which retrieve the value of the property and set a new value correspondingly. However, some properties are read-only and therefore will only have a getter method. For each object, its properties can be one of the following types:

- *Simple* — A property that is represented by a value type. Configuration Server supports two types of simple properties - `string` and `integer`. For example, the `CfgPerson` object has `FirstName` and `LastName` properties, both of the string type.
- *KV-list* — Tree-like properties that are represented by the `KeyValueCollection` class in the Configuration Object Model. Examples of this property include `userProperties` of `CfgPerson`.
- *Structure* — A complex property that includes one or more properties. In the Configuration Object Model, structures are represented by instances of classes that are similar to configuration objects, but cannot be created directly. For example, in the `CfgPerson` class, its `AgentInfo` property contains simple, kv-list and other property types.
- *List of structures* — A property that represents more than one structure. In Configuration Object Model, lists of structures are represented by a generic type `IList<structure_type>`, so that the collection is typed, and clients can easily iterate through the collection.
- *Links to a single object* — In Configuration Server, these properties are stored as DBIDs of external objects. The Configuration Object Model automatically resolves these DBIDs into the real objects, which can be manipulated in the same way as the objects directly retrieved from Configuration Server. Links are initialized at the time of the initial request to one of its properties.

> **Tip**
>
> For each link, there are two ways to set the new value of a link. There is a setter method of the property, which uses an object reference to set a new value of a link. There is also a Set DBID method, which uses an integer DBID value.

- *Links to multiple objects* — A property that contains more than one link. In the Configuration Object Model, lists of structures are represented by a generic type `IList<class_type>`, so that the collection is typed, and clients can easily iterate through the collection.

Creating Instances

One way to create an instance of an object in the Configuration Object Model is to invoke one of the Retrieve methods of the `ConfService` class. This set of methods returns instances of objects that already exist in Configuration Server.

To create a new object in Configuration Server, a client must create a new instance of a COM or "detached" object. The detached object does not correspond to any objects in Configuration Server until it is saved. The detached object is created using the regular object-oriented language object instantiation. For example, a new detached `CfgPerson` object is created using the following construction:

```
[C#]

CfgPerson person = new CfgPerson(confService);
```

An object instance can also be created by using links to external objects. The Component Object Model creates a new object instance whenever the link is called, or any of the properties of a linked object are called. For example, you can write:

```
[C#]

// Person has already been retrieved from Configuration Server.
CfgTenant tenant = person.Tenant;
// This is a link to an external object. It is initialized internally right now...
CfgAddress address = tenant.Address;
```

Common Methods

Each configuration class contains the following methods:

- `Generic GetProperty(string propertyName)` — Retrieves the property value by its name.
- `Generic SetProperty(string propertyName)` — Sets the new value of the property by its name.
- `Save()` — Commits all changes previously made to the object to Configuration Server. If the object was created detached from Configuration Server and has never been saved, a new object is created in Configuration Server using the `RequestCreateObject` method. If the object has been saved or has been retrieved from Configuration Server, a delta-object, which contains all changes to the object, is formed and sent to Configuration Server by means of the `RequestUpdateObject` method.
- `Delete()` — Deletes the object from the Configuration Server Database.
- `Refresh()` — Retrieves the latest version of the object and refreshes the value of all its properties.

> **Tip**
> In this release, all configuration objects are "static," which means that if the object changes in the Configuration Server, the instance of a class is not automatically changed in the Configuration Object Model. Clients must subscribe to the corresponding event and manually refresh the COM object in order for these changes to take effect.

Configuration Service

> **Important**
>
> The IConfService interface was added to COM in release 8.0. All applications should now use this interface to work with the configuration service instead of the old ConfService class. This change is an example of how all COM types in the interface are now referred to by interface; for instance, if a method previously returned CfgObject it now returns ICfgObject. This is not compatible with existing code, but upgrading should not be difficult as the new interfaces support the same methods as the implementing types.

The Configuration Service (IConfService) interface provides services such as retrieval of objects and subscription to events from Configuration Server. Each connection to a Configuration Server (represented by a ConfServerProtocol class of Platform SDK) requires its own instance of the IConfService interface.

The protocol class should be created and initialized in the client code prior to IConfService initialization.

The ConfServiceFactory class is used to create the IConfService. This class uses the following syntax:

```
[C#]

IConfService service = ConfServiceFactory.CreateConfService(protocol);
```

Retrieving Objects

Objects can be retrieved from Configuration Service by using one of the following methods:

- RetrieveObject — Accepts a query that returns one object. If multiple objects are returned, an exception is thrown.
- RetrieveMultipleObjects — Accepts a query that returns one or more objects. A collection of objects is returned.

Each of the Retrieve methods can be can be strongly typed (with use of generics, an object of a specified type is returned) or general (a general object is returned).

Handling Events

The following methods must be called before receiving events from Configuration Server:

1. Register

The application must register its callback by calling the Register method from the Configuration Service. This method supplies the client's filter, which enables the client to receive only requested events.

2. Subscribe

The application must subscribe to events from Configuration Server by calling the `Subscribe` method from the Configuration Service. This method provides a notification query object as a parameter.

The `NotificationQuery` object determines whether the object (or set of objects) to which the client wants to subscribe has changed. The `NotificationQuery` object contains such parameters as object type, object DBID and tenant DBID.

After calling the `Subscribe` method, Configuration Server starts sending events to the client. These events are objects, which contain information such as:

- which object (ID and type) is affected
- the type of event sent to the client
- any additional information

There are three types of events that the client might receive:

- `ObjectCreated` — A new object has been added to Configuration Server.
- `ObjectChanged` — Some of the object properties have been modified in Configuration Server.
- `ObjectDeleted` — The object has been removed from Configuration Server.

Logging Messages

Configuration Object Model Application Block supports logging through the standard Platform SDK logging interfaces. The `IConfService` interface inherits the `EnableLogging` method that provides the ability to log messages through the provided `ILogger` interface.

Releasing a Configuration Service

Whenever a `ConfService` instance is no longer needed, the `ReleaseConfService` method can be used to remove it from the internal list.

```
[C#]

ConfServiceFactory.ReleaseConfService(service);
```

Query Objects

A query object is an instance of a class that contains information required for a successful query to a Configuration Server. This information includes an object type and its attributes (such as `name` and `tenant`), which are used in the search process.

The inheritance structure of configuration server queries is designed to allow for future expansion. The `CfgQuery` object is the base class for all query objects. Other classes extend `CfgQuery` to provide more specific functionality for different types of queries - for example, all filter-based queries use the `CfgFilterBasedQuery` class. This allows room for future query types (such as XPath) to be implemented in this Application Block.

A list of currently available query types is provided below:

- `CfgFilterBasedQuery` — Contains mapped attribute name-value pairs, as well as the object type.

A special query class is supplied for each configuration object type, in order to facilitate the process of making queries to Configuration Server. For each searchable attribute, the query class has a property that can be set. All of these classes inherit attributes from the `CfgQuery` object, and can be supplied as parameters to the `Retrieve` methods which are used to perform searches in Configuration Server.

Cache Objects

The cache functionality is intended to enhance the Configuration Object Model by allowing configuration objects to be stored locally, thereby minimizing requests to configuration server, as well as enhancing ease of use by providing automatic synchronization between locally stored objects and their server-side counterparts.

The cache functionality was designed with the following principles in mind:

- The cache functionality is designed to be extendable with custom implementations of provided interfaces and not via inheritance.
- The cache component is not designed to replicate the Configuration Server query engine or other Configuration Server functionality on the client side.
- Caching must be an optional feature. Work with Configuration Server should not be affected if caching is not used.

Use Cases

Analysis of use cases provides insight into the requirements for applications likely to require configuration cache functionality. The use cases described in the following table were selected for analysis in order to highlight different functional requirements. There are several possible actors which are referenced in the use cases. The actors are as follows:

- Application - Any application which uses the Configuration Object Model application block
- User - Human (or software) user who may perform actions upon objects in the configuration which are separate from the Application

Use Case	Description	Actor	Steps
PLACE OBJECT INTO CACHE	Place a configuration object into the configuration cache (note the object must have been saved — ie must have a DBID in order to exist in the cache).	Application	1. Application adds object to the cache
PLACE OBJECT INTO CACHE ON SAVE	Place a newly created configuration object into the configuration cache when it is saved.	Application	1. Application creates object 2. Application saves object 3. Configuration Object Model Application

Use Case	Description	Actor	Steps
			Block adds object to the cache
PLACE OBJECT INTO CACHE ON RETRIEVE	Allow for automatic insertion of configuration objects into the cache upon retrieval from configuration server.	Application	1. Application retrieves configuration object 2. Configuration Object Model Application Block retrieves the configuration object from the server 3. Configuration Object Model Application Block places the configuration object into the cache 4. Configuration Object Model Application Block returns the object to the application
OBJECT REMOVED IN CONFIGURATION SERVER	When configuration objects are deleted in the configuration server, the cache can delete the local representation of the object as well.	User	1. User deletes object in the Configuration Server 2. Cache removes corresponding local object upon receiving delete notification 3. Cache sends notification of object deletion to Application
SYNCHRONIZE OBJECT PROPERTIES WITH CONFIGURATION SERVER	When an object stored in the cache is updated in the Configuration Server the object must be updated locally as well.	User	1. User updates a configuration object 2. Cache receives notification about object update 3. Cache updates the object based on the received delta 4. Cache fires event informing any subscribers of object

Use Case	Description	Actor	Steps
			change
FIND OBJECT IN CACHE	The cache must support the ability to find a specific configuration object in the cache using object DBID and type as the criteria for the search.	Application	1. Application retrieves object from cache. 2. If object is in the cache, the cache returns the object. Otherwise the application is notified that the requested object is not in the cache.
ACCESS CACHED OBJECTS	The cache must provide its full object collection to the application.	Application	1. Application requests a complete list of objects from the cache. 2. The cache returns a collection of all cached objects.
RETRIEVE LINKED OBJECT FROM CACHE	If caching is turned on, object links which the Configuration Object Model currently resolves through lazy initialization (i.e. if a property linking to another object is accessed, we retrieve the referred-to object from configuration server) must be resolvable through cache access.	Application	1. Application accesses a property which requires link resolution 2. Configuration Object Model Application Block retrieves the linked object from configuration server and stores it in the cache before returning to the application 3. Application again accesses the property and this time the Configuration Object Model Application Block retrieves the object from the cache
PROVIDE CACHE TRANSPARENCY ON RETRIEVE	A cache search should be performed on attempt to retrieve an object from		1. Application creates query to retrieve

Use Case	Description	Actor	Steps
	Configuration Server. If the requested object is found in the cache then the Configuration Object Model should return the cached object rather than accessing Configuration Server.		configuration object 2. Application executes query using the Configuration Object Model 3. Configuration Object Model Application Block searches the cache • If object present, return the object • If object not present, query configuration server for the object
CACHE SERIALIZATION	The cache should support serialization.	Application	1. Application provides a stream to the cache 2. The cache serializes itself into the stream in an XML format 3. Application restarts 4. Application provides the cache a stream of cache data in the same XML format as in step 2 5. Cache restores itself 6. Cache subscribes for updates on the restored objects

Implementation Overview

Two new interfaces for cache management have been added to the Configuration Object Model: the `IConfCache` interface and a default cache implementation (`DefaultConfCache`). Note that the `ConfCache` also implements the `ISubscriber` interface from `MessageBroker`. The cache implements `ISubscriber` in order to allow the user to subscribe to notifications from Configuration Server, as discussed in *Notification And Delta Handling*.

The `IConfCache` interface provides methods for basic functionality such as adding, updating, retrieving, and removing objects in the cache. It also includes a `Policy` property that defines cache behavior and affects method implementation. (For more details about policies, see *Cache Policy*).

The DefaultConfCache component provides a default implementation of the IConfCache interface. It serializes and deserializes cache objects using the XML format described in the *XML Format* section, below.

To enable and configure caching functionality, and to specify ConfService policy, there are three CreateConfService methods available from ConfServiceFactory. The original CreateConfService method (not shown here) creates a ConfService instance that uses the default policy and does not use caching.

```
[C#]

public static IConfService CreateConfService(IProtocol protocol, bool enableCaching)
```

This method creates an instance of a Configuration Service based on the specified protocol. If caching is enabled, the default caching policy will be used. If enableCaching is set to true, caching functionality will be turned on. If caching is disabled, all policy flags related to caching will be false.

```
[C#]

public static IConfService CreateConfService(IProtocol protocol,
        IConfServicePolicy confServicePolicy, IConfCache cache)
```

This method creates a configuration service with the specified policy information. The created service will have caching enabled if a cache object (implementing the IConfCache interface) is passed as a parameter.

```
[C#]

public static IConfService CreateConfService(IProtocol protocol,
        IConfServicePolicy confServicePolicy, IConfCachePolicy confCachePolicy)
```

This method creates a configuration service with the specified policy information. The created service will have caching enabled by default with the cache using the specified cache policy.

XML Format

The "Cache" node will be the root of the configuration cache XML, while "ConfData" is a child of the "Cache" node. The ConfData node contains a collection of XML representations for each configuration object in the cache. The XML format of each object is identical to that which is returned by the ToXml method supported by each the Configuration Object Model configuration object.

The "CacheConfiguration" element is a child of the "Cache" node. There can only be one instance of this node and it contains all cache configuration parameters, as follows:

- CONFIGURATIONSERVER NODE – There can be 1..n instances of this element. Each one will represent a configuration server for which the cache is applicable (a cache can be applicable to multiple configuration servers if they are working with the same database as in the case of a primary and backup configuration server pair). Each ConfigurationServer element will have a URI attribute specifying the unique URI identifying the Configuration Server, as well as a Name attribute specifying the name associated with the endpoint.

The example provided below shows a cache that is applicable for the configuration server at "server:2020" with some policy details specified. There are two objects in the cache for this example: a CfgDN and a CfgService object.

[XML]

```
<Cache>
	<CacheConfiguration>
		<ConfigurationServer name="serverName" uri="tcp://server:2020"/>
	</CacheConfiguration>

	<ConfData>
		<CfgDN>
			<DBID value="267" />
			<switchDBID value="111" />
			<tenantDBID value="1" />
			<type value="3" />
			<number value="1111" />
			<loginFlag value="1" />
			<registerAll value="2" />
			<groupDBID value="0" />
			<trunks value="0" />
			<routeType value="1" />
			<state value="1" />
			<name value="DNAlias" />
			<useOverride value="2" />
			<switchSpecificType value="1" />
			<siteDBID value="0" />
			<contractDBID value="0" />
			<accessNumbers />
			<userProperties />
		</CfgDN>

		<CfgService>
			<DBID value="102" />
			<name value="Solution1" />
			<type value="2" />
			<state value="1" />
			<solutionType value="1" />
			<components>
				<CfgSolutionComponent>
					<startupPriority value="3" />
					<isOptional value="2" />
					<appDBID value="153" />
				</CfgSolutionComponent>
			</components>
			<SCSDBID value="102" />
			<assignedTenantDBID value="101" />
			<version value="7.6.000.00" />
			<startupType value="2" />
			<userProperties />
			<componentDefinitions />
			<resources />
		</CfgService>
	</ConfData>

</Cache>
```

Cache Policy

The configuration cache can be assigned a policy represented by a Policy interface. A default implementation of the interface will be provided in the `DefaultConfCachePolicy` class.

The `IConfCache` interface will interpret the policy as follows:

1. `CacheOnCreate` – When an object is created in the configuration server, the policy will be checked with the created object as the parameter. If the method returns true, the object will be added to the cache, if it is false, the object will not be added. Default implementation will always return false.
2. `RemoveOnDelete` – When an object is deleted in the configuration server, the policy will be checked with the deleted object as the parameter. If the method returns true, the object will be deleted in the cache, if it is false, the notification will be ignored. Default implementation will always return true.
3. `TrackUpdates` – When an object is updated in the configuration server, the policy will be checked with the current version of the object as the parameter. If the method returns true, the object will be updated with the received delta, if it is false, the notification will be ignored. Default implementation will always return true.
4. `ReturnCopies` – Determines whether the cache should return copies of objects when they are retrieved from the cache, or the original, cached versions. False by default.

IConfServicePolicy Interface

The `IConfServicePolicy` interface can be used to define the policy settings for the `ConfService`. Two default implementations are available:

1. `DefaultConfServicePolicy` contains the settings for a non-caching configuration service. That is, all of the cache-related policy flags will always return false.
2. `CachingConfServicePolicy` defines the default behavior for a configuration service with caching enabled. (Note that when referring to the "default" value below, we will be referring to this implementation.)

The policy interface settings are interpreted as follows:

- `AttemptLinkResolutionThroughCache` – Whenever a link resolution attempt is made, this policy will be checked for the type of object the link refers to. If this method returns true, the link resolution attempt will first be made through the cache. If the method returns false, or if the object has not been found in the cache, the server will queried. Default value is always true.
- `CacheOnRetrieve` – This method will be called for each object retrieved from the configuration. If the return value is "true" the object will be added to the cache. Default value is always true.
- `CacheOnSave` – This method will be called for each object that is being saved. If the return value is true, the object will be added to the cache. If the object is already in the cache, it will not be overwritten. Default value is always true.
- `ValidateBeforeSave` – This is a property from the `ConfService` which will be moved to the policy interface and is not related to caching. It is used to indicate whether property values are checked for valid values against the schema before a save attempt is made. Default value is true.
- `QueryCacheOnRetrieve` – This method will be called every time a retrieve operation is performed using a query. The `ConfService` will first check the cache for the existence of the requested configuration object. If the object exists, it will be returned and no configuration server request will be made. If there are no values returned, the `ConfService` will query the configuration server (see *Query Engine*). Default value is always false.
- `QueryCacheOnRetrieveMultiple` – This method will be called every time a retrieve multiple operation is performed. The ConfService will first execute the query against cache. If the returned object count is greater than 0 the found object collection will be returned and no configuration server request will be made. If there are no values returned, the `ConfService` will query the configuration server (see *Query Engine*). Default value is always false.

Note that the `RetrieveMultiple` operation is NOT implemented in the default query engine, so providing a policy where this method returns true will require a new query engine implementation.

Cache Extendability

Consistent with the design principles outlined above, the configuration cache is extendable via custom implementations of provided interfaces. The two areas of the cache which can be extended are the cache storage and the cache query engine.

Cache Storage

The storage interface defines the method by which objects are stored in the cache. When an instance of an implementing object is provided to the cache, the cache will store all cached objects in the storage component.

The default storage implementation stores cached objects using the object type and DBID as keys. Note that this means that objects in the cache are assumed to be from one configuration database. The default implementation is also thread safe using a reader/writer lock which allows for multiple concurrent readers and one writer. The storage methods are as follows:

- Add – Adds a new object to the storage. If object already exists in the storage, the default implementation thrown an exception.
- Update – Overwrites an existing object in the storage. If the object is not found in the storage, the default implementation creates a new version of the object.
- Remove – Removes an object from the storage.
- Retrieve – Retrieves an enumerable list of all objects in the storage (filtered by type), and possibly influenced by an optional helper parameter. Note that the helper parameter is not meant to provide querying logic – that should be done in the query engine. Because the query engine is to some degree dependent on the storage implementation, the helper parameter allows for some flexibility in the way stored objects are enumerated for the query engine. The default implementation can take a `CfgObjectType` as a helper parameter.
- Clear–Removes all objects in the storage.

Query Engine

The query engine provides the ability to define the method by which objects are located in the cache.

Depending on the `IConfService` policy, `Retrieve` requests as well as link resolution can first be attempted through the cache. If the requested object is found in the cache, then that cached object is returned instead of sending a request to Configuration Server. If the object is not present in the cache, a request to Configuration Server is made.

A user-definable query engine module exists inside the cache to achieve this functionality. A query engine must implement the `IConfCacheQueryEngine` interface, which provides methods to retrieve objects (either individually, or as a list) and to test a query and determine if it can be executed.

If enabled by the policy, `IConfService` will attempt a query to its cache using the cache's query engine interface. If a result is returned, the `IConfService` will not query the Configuration Server. By following this contract, the Configuration Object Model user is then able to create a custom implementation of the `IConfCacheQueryEngine` with any extended search capabilities which may be missing from the simple default implementation.

Two implementations of the `IConfCacheQueryEngine` interface are provided in the Configuration

Object Model, as described below:

- DEFAULTCONFCACHEQUERYENGINE CLASS - The `DefaultConfCacheQueryEngine` class is a default implementation of the `IConfCacheQueryEngine` interface.
- COMPOSITECONFCACHEQUERYENGINE CLASS - This class is a more advanced implementation of the query engine which allows child query engine modules to be registered in order to interpret different types of queries. It does not have a default query engine implementation, only the mechanism for working with multiple child query engines.

Notification and Delta Handling

The default configuration cache will implement the `ISubscriber<ConfEvent>` interface which will allow the cache to be subscribed to receive configuration events. When a cache instance is associated with a Configuration Service, it will automatically be subscribed for configuration events from that service (note that if a custom cache implementation also implements this interface it will be subscribed for events as well). The way the cache is updated based on these notifications is determined by the cache policy.

In addition, a new filter class will be added in order to allow the subscriber to filter the cache events. The `ConfCacheFilter` will implement the `MessageBroker`'s `IPredicate` interface, allowing for the filter to be passed during registration for events via `ISubscriptionService`. The `ConfCacheFilter`'s properties will specify the parameters by which the events will be filtered. Initially, the supported parameters will be object type, object DBID, and update type, allowing the user to filter events by one or a combination of these parameters assuming an AND relationship between the parameters specified.

The Configuration Object Model Application Block Interface

The following figures show the relationships among many of the classes that make up this application block.

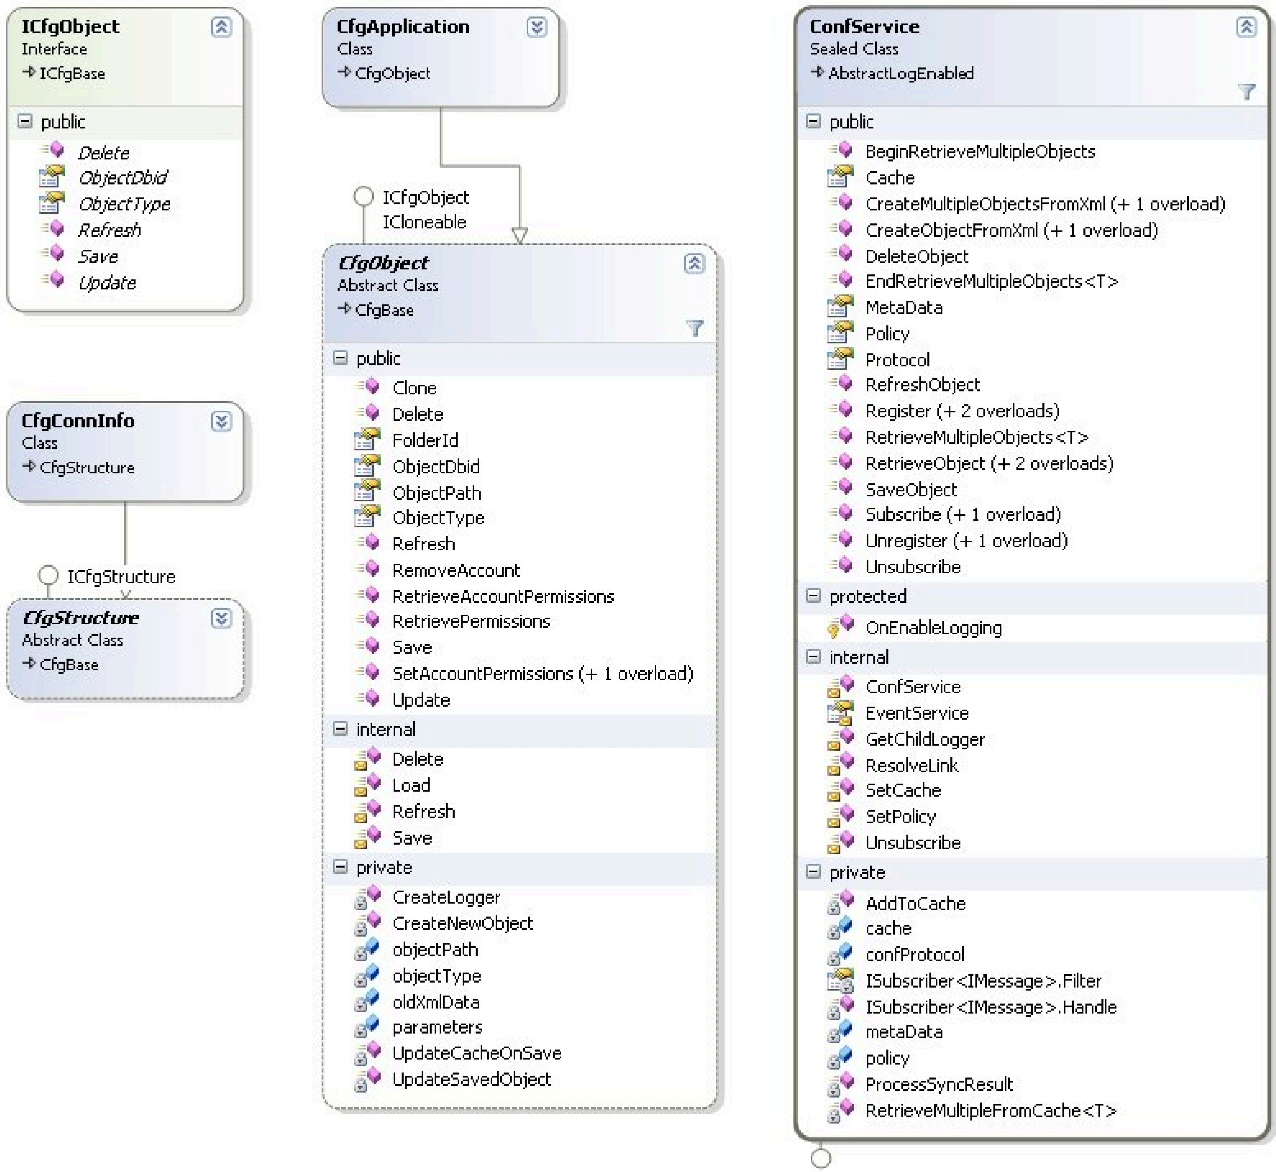
ICfgObject
Interface
ICfgBase
public
Delete
ObjectDbid
ObjectType
Refresh
Save
Update
CfgApplication
Class
CfgObject
ICfgObject
ICloneable
CfgObject
Abstract Class
CfgBase
public
Clone
Delete
FolderId
ObjectDbid
ObjectPath
ObjectType
Refresh
RemoveAccount
RetrieveAccountPermissions
RetrievePermissions
Save
SetAccountPermissions (+ 1 overload)
Update
internal
Delete
Load
Refresh
Save
private
CreateLogger
CreateNewObject
objectPath
objectType
oldXmlData
parameters
UpdateCacheOnSave
UpdateSavedObject
CfgConnInfo
Class
CfgStructure
ICfgStructure
CfgStructure
Abstract Class
CfgBase
ConfService
Sealed Class
AbstractLogEnabled
public
BeginRetrieveMultipleObjects
Cache
CreateMultipleObjectsFromXml (+ 1 overload)
CreateObjectFromXml (+ 1 overload)
DeleteObject
EndRetrieveMultipleObjects<T>
MetaData
Policy
Protocol
RefreshObject
Register (+ 2 overloads)
RetrieveMultipleObjects<T>
RetrieveObject (+ 2 overloads)
SaveObject
Subscribe (+ 1 overload)
Unregister (+ 1 overload)
Unsubscribe
protected
OnEnableLogging
internal
ConfService
EventService
GetChildLogger
ResolveLink
SetCache
SetPolicy
Unsubscribe
private
AddToCache
cache
confProtocol
ISubscriber<IMessage>.Filter
ISubscriber<IMessage>.Handle
metaData
policy
ProcessSyncResult
RetrieveMultipleFromCache<T>

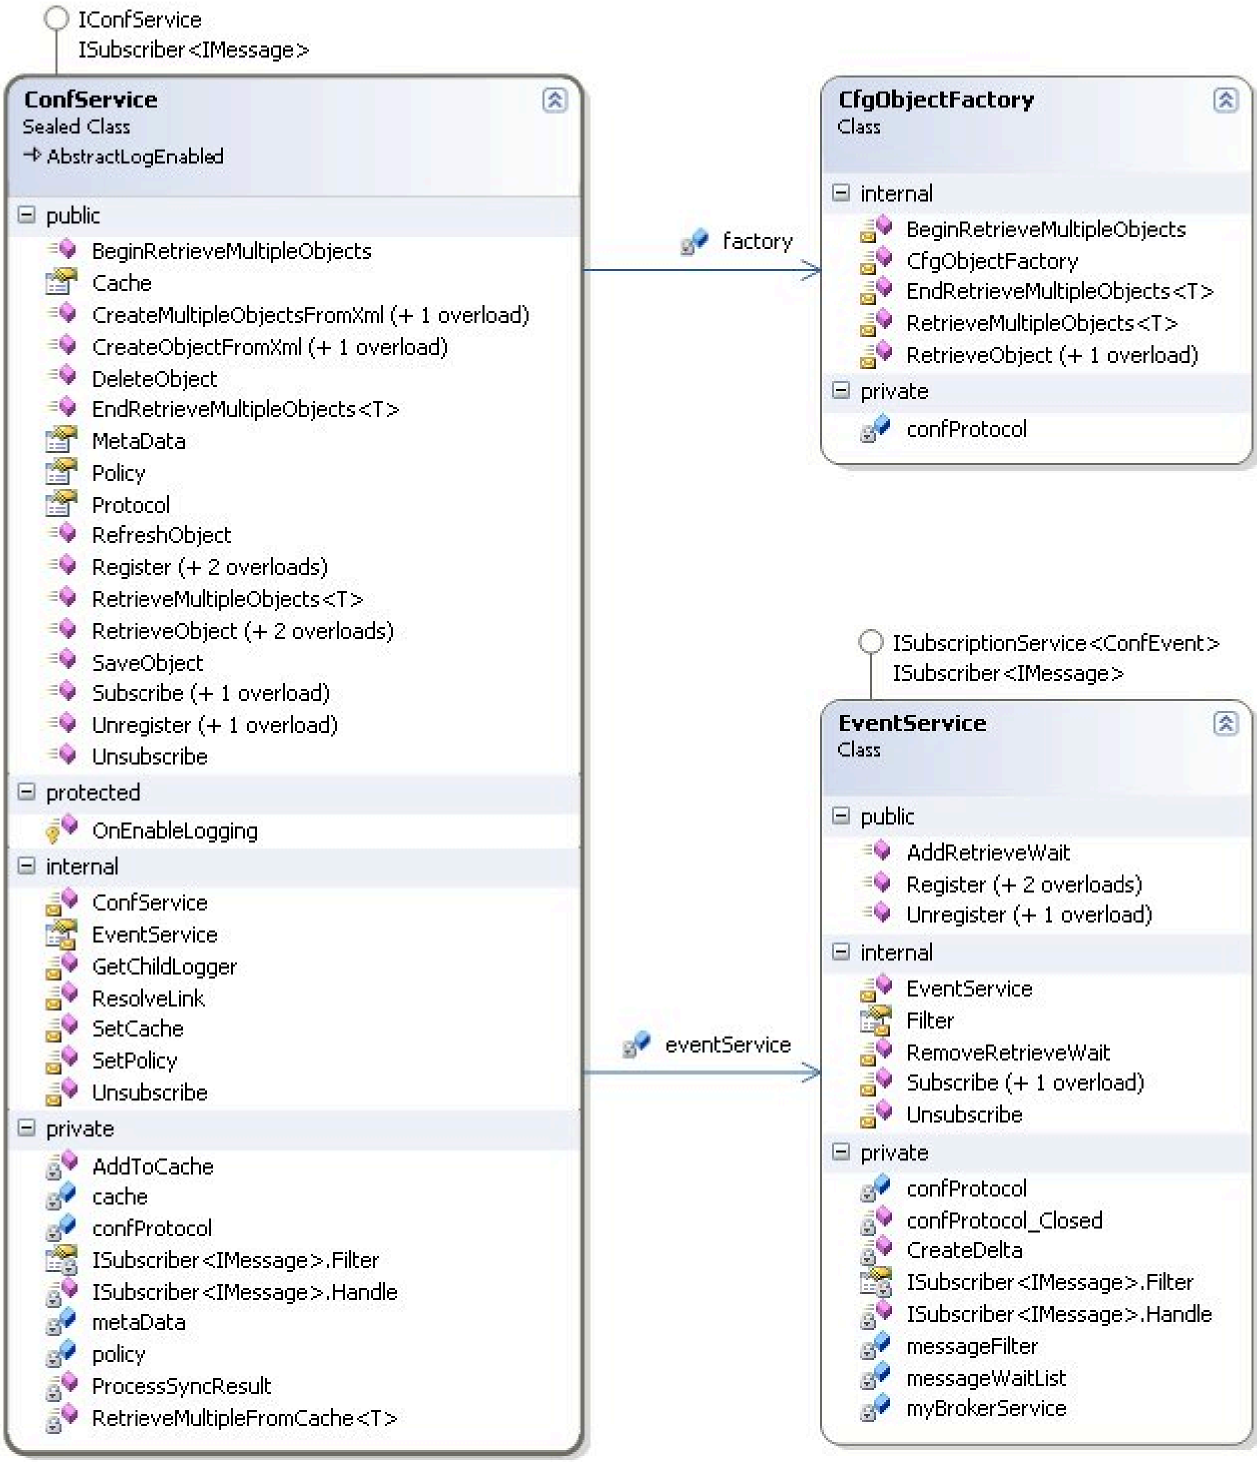
IConfService
ISubscriber<IMessage>
ConfService
Sealed Class
AbstractLogEnabled
public
BeginRetrieveMultipleObjects
Cache
CreateMultipleObjectsFromXml (+ 1 overload)
CreateObjectFromXml (+ 1 overload)
DeleteObject
EndRetrieveMultipleObjects<T>
MetaData
Policy
Protocol
RefreshObject
Register (+ 2 overloads)
RetrieveMultipleObjects<T>
RetrieveObject (+ 2 overloads)
SaveObject
Subscribe (+ 1 overload)
Unregister (+ 1 overload)
Unsubscribe
protected
OnEnableLogging
internal
ConfService
EventService
GetChildLogger
ResolveLink
SetCache
SetPolicy
Unsubscribe
private
AddToCache
cache
confProtocol
ISubscriber<IMessage>.Filter
ISubscriber<IMessage>.Handle
metaData
policy
ProcessSyncResult
RetrieveMultipleFromCache<T>
factory
CfgObjectFactory
Class
internal
BeginRetrieveMultipleObjects
CfgObjectFactory
EndRetrieveMultipleObjects<T>
RetrieveMultipleObjects<T>
RetrieveObject (+ 1 overload)
private
confProtocol
eventService
ISubscriptionService<ConfEvent>
ISubscriber<IMessage>
EventService
Class
public
AddRetrieveWait
Register (+ 2 overloads)
Unregister (+ 1 overload)
internal
EventService
Filter
RemoveRetrieveWait
Subscribe (+ 1 overload)
Unsubscribe
private
confProtocol
confProtocol_Closed
CreateDelta
ISubscriber<IMessage>.Filter
ISubscriber<IMessage>.Handle
messageFilter
messageWaitList
myBrokerService

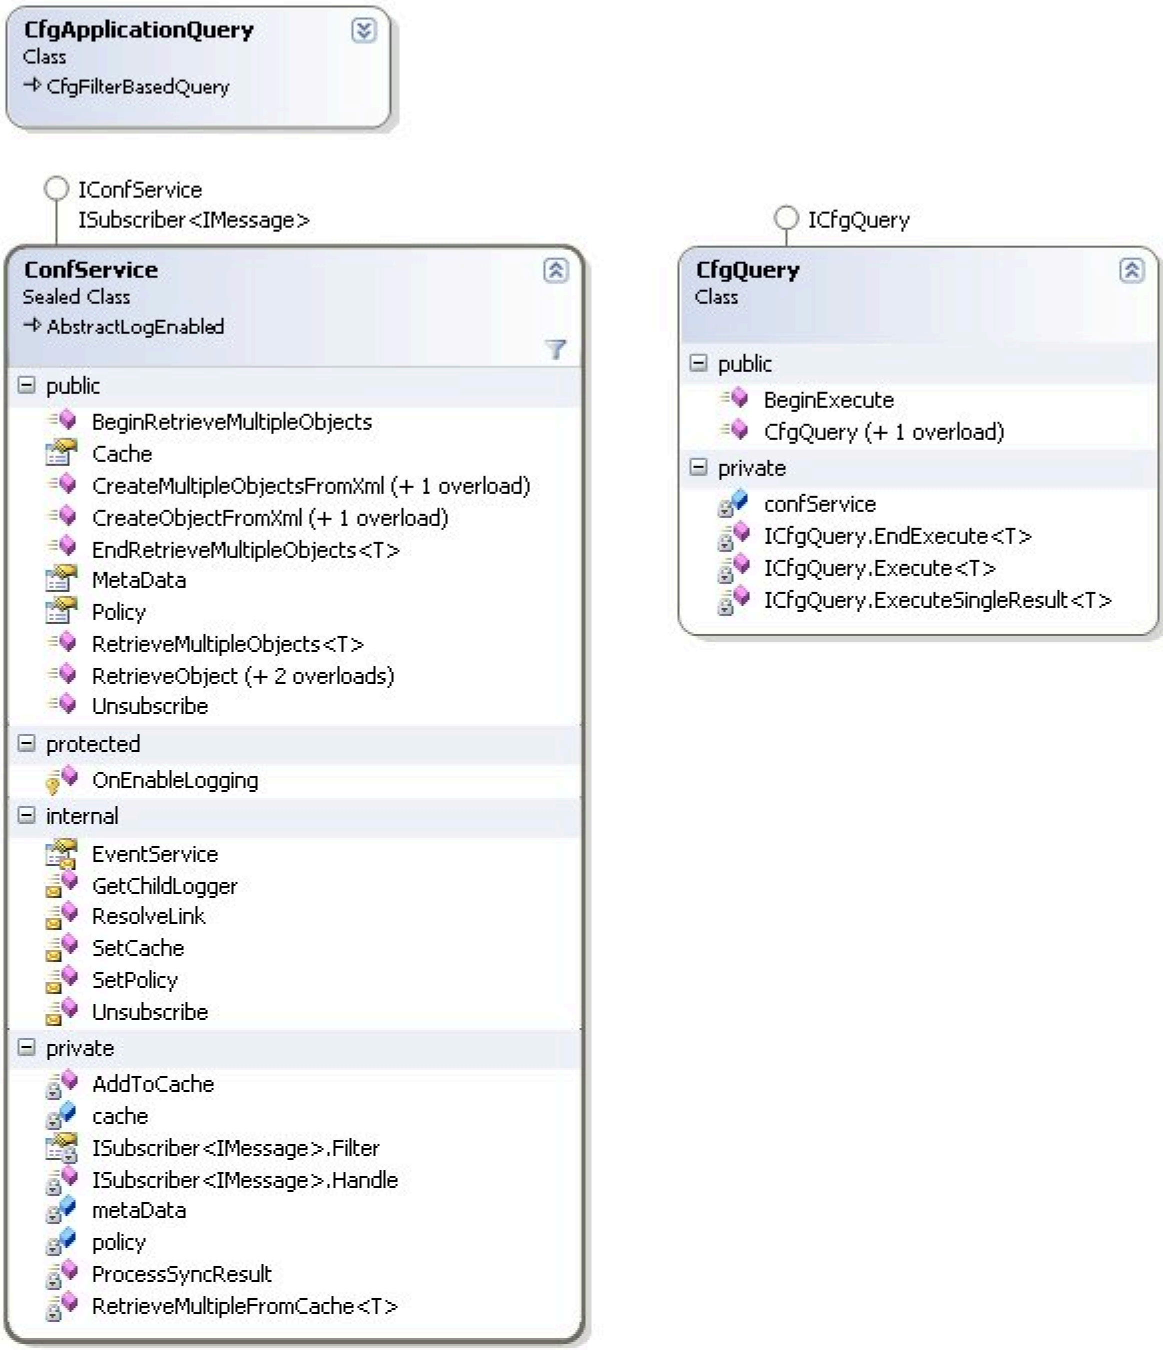
CfgApplicationQuery
Class
CfgFilterBasedQuery
IConfService
ISubscriber<IMessage>
ConfService
Sealed Class
AbstractLogEnabled
public
BeginRetrieveMultipleObjects
Cache
CreateMultipleObjectsFromXml (+ 1 overload)
CreateObjectFromXml (+ 1 overload)
EndRetrieveMultipleObjects<T>
MetaData
Policy
RetrieveMultipleObjects<T>
RetrieveObject (+ 2 overloads)
Unsubscribe
protected
OnEnableLogging
internal
EventService
GetChildLogger
ResolveLink
SetCache
SetPolicy
Unsubscribe
private
AddToCache
cache
ISubscriber<IMessage>.Filter
ISubscriber<IMessage>.Handle
metaData
policy
ProcessSyncResult
RetrieveMultipleFromCache<T>
ICfgQuery
CfgQuery
Class
public
BeginExecute
CfgQuery (+ 1 overload)
private
confService
ICfgQuery.EndExecute<T>
ICfgQuery.Execute<T>
ICfgQuery.ExecuteSingleResult<T>

Using the Application Block

Installing the Configuration Object Model Application Block

Before you install the Configuration Object Model Application Block, it is important to review the software requirements and the structure of the software distribution.

Software Requirements

To work with the Configuration Object Model Application Block, you must ensure that your system meets the software requirements established in the Genesys Supported Operating Environment Reference Manual, as well as meeting the following minimum software requirements:

- Genesys Configuration Platform SDK 8.0 or higher

Configuring the Configuration Object Model Application Block

In order to use the QuickStart application, you will need to set up the XML configuration file that comes with the application block. This file is located at `Quickstart\app.config`. This is what the contents look like:

```
<?xml version="1.0" encoding="utf-8" ?>
<configuration>
  <appSettings>
    <add key="Uri" value="tcp://yourhost:yourport"/>

    <add key="ClientName" value="StarterApp"/>

    <add key="ClientType" value="CFGAgentDesktop"/>

    <add key="UserName" value="default"/>

    <add key="Password" value="password"/>

  </appSettings>
</configuration>
```

Follow the instructions in the comments and save the file.

Building the Configuration Object Model Application Block

The Platform SDK distribution includes a `Genesyslab.Platform.ApplicationBlocks.ConfigurationObjectModel.dll` file that you can use as is. This file is located in the `bin` directory at the root level of the Platform SDK directory. To build your own copy of this application block, follow the instructions below:

To build the Configuration Object Model Application Block:

1. Open the <Platform SDK Folder>\ApplicationBlocks\Com folder.
2. Double-click Com.sln.
3. Build the solution.

Using the QuickStart Application

The easiest way to start using the Configuration Object Model Application Block is to use the bundled QuickStart application. This application ships in the same folder as the application block.

To run the QuickStart application:

1. Open the `<Platform SDK Folder>\ApplicationBlocks\Com` folder.
2. Double-click `ComQuickStart.sln`.
3. Build the solution.
4. Find the executable for the QuickStart application, which will be at `<Platform SDK Folder>\ApplicationBlocks\Com\QuickStart\bin\Debug\ComQuickStart.exe`.
5. Double-click `ComQuickStart.exe`.

Introduction to the Configuration Layer Objects

The Genesys Configuration Layer is a database containing information about the objects in your contact center environment. You may need to get information about these objects. You may also want to add, update, or delete them. The Configuration Platform SDK gives you the means to do that.

This article contains information that is common to all of these Configuration Layer objects.

Once you have reviewed the information in this section, you can look at the detailed descriptions of these objects.

General Parameters

The following parameters are common to objects of all types. They will not be described again in the listings for individual objects.

- `DBID` — An identifier of this object in the Configuration Database. Generated by Configuration Server, it is unique within an object type. Identifiers of deleted objects are not used again. Read-only.
- `state` — Current object state. Mandatory. Refer to `CfgObjectState` in section Variable Types.

> **Tip**
> Change in the state of a parent object will cause the states of all its child objects to change accordingly. Configuration Server will provide a notification for each elementary change. Changing the state of a parent object will not be allowed unless the client application has privileges to change all of the child objects of this parent object.

- `userProperties` — In objects, a pointer to the list of user-defined properties. In delta objects, a pointer to a list of user-defined properties added to the existing list. Parameter `userProperties` has the following structure: Each key-value pair of the primary list (`TKVList *userProperties`) uses the key for the name of a user-defined section, and the value for a secondary list, that also has the `TKVList` structure and specifies the properties defined within that section. Each key-value pair of the secondary list uses the key for the name of a user-defined property, and the value for its current setting. User properties can be defined as variables of integer, character, or binary type. Names of sections must be unique within the primary list. Names of properties must be unique within the secondary list.

> **Tip**
> Configuration Server is not concerned with logical meanings of user-defined sections,

> properties, or their values.

- `deletedUserProperties` — A pointer to the list of deleted user-defined properties. Has the same structure as parameter `userProperties` above. A user-defined property is deleted by specifying the name of the section that this property belongs to, and the name of the property itself with any value. A whole section is deleted by specifying the name of that section and an empty secondary list.

- `changedUserProperties` — A pointer to the list of user-defined properties whose values have been changed. Has the same structure as parameter `userProperties` above. A value of a user-defined property is changed by specifying the name of the section that this property belongs to, the name of the property itself, and a new value of that property.

- `flexibleProperties` — In objects, a pointer to the list of additional properties. In delta objects, a pointer to a list of user-defined properties added to the existing list. This parameter has the following structure: Each key-value pair of the primary list (`TKVList * flexibleProperties`) uses the key for the name of the section, and the value for a secondary list, that also has the `TKVList` structure and specifies either properties defined within that section or another section name. Each key-value pair of the secondary list uses the key for the name of a property, and the value for its current setting. Properties can be defined as variables of integer, character, or binary type or as the name of another list of properties. Names of sections must be unique within the primary list. Names of properties must be unique within the list. The data structure within the `flexibleProperties` property is object-type specific and hard-coded within Configuration Server. Each key-value in the `TKVList * flexibleProperties` is controlled and processed by Configuration Server only in the same manner as any other property in contrast with user-properties the contents of which are not Configuration Server concerned. If the structure of the property's `Extension` is not specified, the value is `NULL`. For more information, see the detailed object descriptions in this document.

Configuration Object Association

Configuration Objects can be associated with each other in a number of different ways that can be generally classified as follows:

- Parent-child relationship, where a child object cannot be created without a parent and will be deleted automatically if its parent object is deleted. Most of the object types will have an explicit reference to their parents which is marked with an asterisk in the specification below. For the object types that do not have such a reference, it is implied that their parent is the Service Provider (that is, the imaginary tenant with `DBID = 1`).
- Exclusive association, where an object cannot be associated in the same manner with more than one other object.
- Non-exclusive association, where an object can be associated in the same manner with more than one other object. Unless expressly noted otherwise, a reference to the DBID of another object without an asterisk indicates a non-exclusive assignment.

The parameters of all object-related structures are optional unless otherwise noted. However, all variables of character type must be initialized at the time an object is created. The variables of character type that are not mandatory may be initialized with an empty string (the recommended default value unless otherwise noted). The variables of character type that are mandatory may not be initialized with an empty string. Variables of character type may accept values of up to 255 symbols in length unless otherwise noted. The recommended default value for optional parameters of other types is `zero` or `NULL`, unless otherwise noted.

Filters

Filters are used to specify more precisely the kind of information that the client application is interested in. Filters reduce both volumes of data communicated by Configuration Server and data-processing efforts on the client side. Filters are structured as key-value pairs where the value of each key defines a certain condition of data selection. Filter keys are defined as variables of integer type unless otherwise noted.

Here is a list of common filter types:

- `folder_dbid` — A unique identifier of a folder. If specified, Configuration Server will return information only about objects of specific type located under specified folder. See also the description of the `ConfGetObjectInfo` function.
- delegate_dbid — A unique identifier of an account on behalf of which current query is to be executed. Produced result set will be calculated using a superposition of the registered account permissions and that passed in `delegate_dbid` filter. Must be used in conjunction with delegate_type filter in order to specify account type (`CFGPerson` or `CFGAccessGroup`).
- `delegate_type` — Object type of the account (`CFGPerson` or `CFGAccessGroup`) on behalf of which the current query is to be executed. Must be used in conjunction with `delegate_dbid`.
- `object_path` — A flag that causes Configuration Server to return a full path of the object in the folder hierarchy for every object in the result set. The path string will be returned in the `cfgDescription` field of the `CFGObjectInfo` event.
- `cmp_insensitive` — A flag that causes Configuration Server to perform case-insensitive comparison of string values in the filter. Supported from Configuration Server 7.2.000.00.
- `read_folder_dbid` — A flag that causes Configuration Server to return a Folder DBID for every object in the result set. The folder will be returned in the `cfgExtraInfo3` field of the `CFGObjectInfo` event. Supported from Configuration Server 7.2.000.00.

List of Configuration Layer Objects

The following table provides a convenient list of Configuration Layer Objects that are available. For more information, refer to Introduction to the Configuration Layer Objects.

CfgAccessGroup	CfgDeltaCampaign	CfgDeltaTableAccess	
CfgActionCode	CfgDeltaCampaignGroup	CfgDeltaTenant	
CfgAgentGroup	CfgDeltaDN	CfgDeltaTimeZone	
CfgAgentInfo	CfgDeltaDNGroup	CfgDeltaTransaction	
CfgAgentLogin	CfgDeltaEnumerator	CfgDeltaTreatment	CfgRole
CfgAgentLoginInfo	CfgDeltaEnumeratorValue	CfgDeltaVoicePrompt	CfgRoleMember
CfgAlarmCondition	CfgDeltaField	CfgEnumerator	CfgScheduledTask
CfgAlarmEvent	CfgDeltaFilter	CfgEnumeratorValue	CfgScript
CfgAppPrototype	CfgDeltaFolder	CfgField	CfgServer
CfgAppRank	CfgDeltaFormat	CfgFilter	CfgService
CfgAppServicePermission	CfgDeltaGVPCustomer	CfgFolder	CfgServiceInfo
CfgApplication	CfgDeltaGVPIVRProfile	CfgFormat	CfgSkill
CfgCallingList	CfgDeltaGVPReseller	CfgGVPCustomer	CfgSkillLevel
CfgCallingListInfo	CfgDeltaGroup	CfgGVPIVRProfile	CfgSolutionComponent
CfgCampaign	CfgDeltaHost	CfgGVPReseller	CfgSolutionComponentDefinition
CfgCampaignGroup	CfgDeltaIVR	CfgGroup	CfgStatDay
CfgConnInfo	CfgDeltaIVRPort	CfgHost	CfgStatInterval
CfgDN	CfgDeltaObjectiveTable	CfgID	CfgStatTable
CfgDNAccessNumber	CfgDeltaPerson	CfgIVR	CfgSubcode
CfgDNGroup	CfgDeltaPhysicalSwitch	CfgIVRPort	CfgSwitch
CfgDNInfo	CfgDeltaPlace	CfgOS	CfgSwitchAccessCode
CfgDeltaAccessGroup	CfgDeltaPlaceGroup	CfgObjectResource	CfgTableAccess
CfgDeltaActionCode	CfgDeltaRole	CfgObjectiveTable	CfgTenant
CfgDeltaAgentGroup	CfgDeltaScheduledTask	CfgObjectiveTableRecord	CfgTimeZone
CfgDeltaAgentInfo	CfgDeltaScript	CfgPerson	CfgTransaction
CfgDeltaAgentLogin	CfgDeltaService	CfgPhones	CfgTreatment
CfgDeltaAlarmCondition	CfgDeltaSkill	CfgPhysicalSwitch	CfgVoicePrompt
CfgDeltaAppPrototype	CfgDeltaStatDay	CfgPlace	
CfgDeltaApplication	CfgDeltaStatTable	CfgPlaceGroup	
CfgDeltaCallingList	CfgDeltaSwitch	CfgPortInfo	

List of Configuration Layer Enumerations

The following table provides a convenient list of Configuration Layer Enumerations that are available. For more information, refer to Introduction to the Configuration Layer Objects.

CfgAccessGroupType	CfgEnumeratorType	CfgOSType	CfgSolutionType
CfgActionCodeType	CfgErrorType	CfgObjectState	CfgStartupType
CfgAlarmCategory	CfgEventType	CfgObjectType	CfgStatDayType
CfgAppComponentType	CfgFieldType	CfgObjectiveTableType	CfgStatTableType
CfgAppType	CfgFilterType	CfgOperationMode	CfgSwitchType
CfgCallActionCode	CfgFlag	CfgOperationalMode	CfgTableType
CfgChargeType	CfgFolderClass	CfgOptimizationMethod	CfgTargetType
CfgCTILinkType	CfgGroupType	CfgPermissions	CfgTaskType
CfgDIDGroupType	CfgHAType	CfgPersonType	CfgTraceMode
CfgDNGroupType	CfgHostType	CfgRank	CfgTransactionType
CfgDNRegisterFlag	CfgIVRProfileType	CfgRecActionCode	GctiCallState
CfgDNType	CfgIVRType	CfgResourceType	GctiContactType
CfgDataType	CfgLanguage	CfgRouteType	GctiRecordStatus
CfgDialMode	CfgLinkType	CfgScriptType	GctiRecordType
CfgEnumeratorObjectType	CfgMediaType	CfgSelectionMode	

Stat Server

Stat Server tracks information about customer interaction networks (contact center, enterprise-wide, or multi-enterprise telephony and computer networks). It also converts the data accumulated for directory numbers (DNs), agents, agent groups, and non-telephony-specific object types, such as e-mail and chat sessions, into statistically useful information, and passes these calculations to other software applications that request data. For example, Stat Server sends data to Universal Routing Server (URS), because Stat Server reports on agent availability. You can also use Stat Server's numerical statistical values as routing criteria.

Stat Server provides contact center managers with a wide range of information, allowing organizations to maximize the efficiency and flexibility of customer interaction networks. For more information about Stat Server, consult the Reporting Technical Reference 8.0 Overview and the Framework 8.1 Stat Server User's Guide.

You can use the Platform SDK to write Java or .NET applications that gather statistical information from Stat Server. These applications may be fairly simple or quite advanced. This article shows how to implement the basic functions you will need to write a simple Statistics application.

A Typical Statistics Application

There are many ways in which you might need to use data from Stat Server, but in most cases, you will use three types of requests:

- `RequestOpenStatistic` and `RequestOpenStatisticEx` are used to ask Stat Server to start sending statistical information to your application. `RequestOpenStatistic` allows you to request information about a statistic that has already been defined in the Genesys Configuration Layer, while you can use `RequestOpenStatisticEx` to define your own statistics dynamically.
- You can use `RequestPeekStatistic` to get the value of a statistic that has already been opened using either `RequestOpenStatistic` or `RequestOpenStatisticEx`. Since it can take a while for certain types of statistical information to be sent to your application, this can be useful if you are writing an application—such as a wallboard application, for instance—for which you would like statistical values to be displayed immediately.
- Use `RequestCloseStatistic` to tell Stat Server that you no longer need information about a particular statistic.

> **Tip**
>
> When you use `RequestOpenStatistic` and `RequestOpenStatisticEx`, you have to specify a `ReferenceId`, which is a unique integer that allows Stat Server and your application to distinguish between different sets of statistical information. You must also enter this integer in the `StatisticId` field for any request that refers to the statistics generated on the basis of your Open request. For example, if you sent a request for "TotalNumberInboundCalls" for agent 001, you might give the `RequestOpenStatistic` a `ReferenceId` of `333001`. A similar request for agent 002

> might have a ReferenceId of 333002. When you want to peek at the value of "TotalNumberInboundCalls" for agent 001, or close the statistic (or suspend or resume reporting on the statistic), you need to specify a StatisticId of 333001 for each of these requests.

Java

Connecting to Stat Server

As mentioned in the article on the architecture, the Platform SDKs uses a message-based architecture to connect to Genesys servers. This section describes how to connect to Stat Server, based on the material in the article on Connecting to a Server.

After you have set up your import statements, the first thing you need to do is create a `StatServerProtocol` object:

```
[Java]

StatServerProtocol statServerProtocol =
  new StatServerProtocol(
    new Endpoint(
      statServerEPName,
      host,
      port));
statServerProtocol.setClientName(clientName);
```

You can also configure your ADDP and warm standby settings at this point, following the example shown in the Connecting to a Server article.

Once your configuration is complete, open the connection to Stat Server:

```
[Java]

try {
        statServerProtocol.open();
} catch (InterruptedException e) {
        e.printStackTrace();
} catch (ProtocolException e) {
        e.printStackTrace();
}
```

Working with Statistics

The Stat Server application object in the Genesys Configuration Layer comes with many predefined statistics. You can also define your own statistics using the options tab of this application object. The Platform SDK allows you to get information about any of these statistics by using

`RequestOpenStatistic`. There may be times, however, when you want your application to be able to create new types of statistics dynamically. The Platform SDK also supports this, with the use of `RequestOpenStatisticEx`.

This section will show you how to use `RequestOpenStatistic` to get information on a predefined statistic. After that, we will give an example of how to use `RequestOpenStatisticEx`.

The first thing you need to do to use `RequestOpenStatistic` is to create the request:

```
[Java]

RequestOpenStatistic requestOpenStatistic
        = RequestOpenStatistic.create();
```

Now you need to describe the *statistics object*, that is, the object you are monitoring. This description consists of the object's Configuration Layer ID and object type, and the tenant ID and password:

```
[Java]

StatisticObject object = StatisticObject.create();
object.setObjectId("Analyst001");
object.setObjectType(StatisticObjectType.Agent);
object.setTenantName("Resources");
object.setTenantPassword("");
```

Next, you will specify the `StatisticType` property, which must correspond to the name of the statistic definition that appears in the options tab. In this case, we are asking for the total login time for an agent identified as "Analyst001":

```
[Java]

StatisticMetric metric = StatisticMetric.create();
metric.setStatisticType("TotalLoginTime");
```

Now you can specify the desired `Notification` settings. The Statistics Platform SDK supports four ways of gathering statistics:

1. `NoNotification` allows you to retrieve statistics when you want them.
2. `Periodical` means Stat Server reports on statistics based on the time period you request.
3. `Immediate` means Stat Server reports on statistics whenever a statistical value changes. For time-related statistics, `Immediate` means that Stat Server will report the current value whenever a statistical value changes, but it will also report that value periodically, using the specified notification frequency.
4. `Reset` means Stat Server reports the current value of a statistic right before setting the statistical value to zero (0).

In this case, we are interested in receiving statistics on a regular basis, so we have asked for a notification mode of `Periodical`, with updates every 5 seconds, using a `GrowingWindow` statistic interval. For more information on notification modes, see the section on Notification Modes in the Framework 8.1 Stat Server User's Guide. For more information on statistic intervals, see the section on TimeProfiles in the same guide.

```
[Java]

Notification notification = Notification.create();
notification.setMode(NotificationMode.Periodical);
notification.setFrequency(5);
```

At this point, you can add the information about the statistic object and your notification settings to the request:

```
[Java]

requestOpenStatistic.setStatisticObject(object);
requestOpenStatistic.setStatisticMetric(metric);
requestOpenStatistic.setNotification(notification);
```

Before sending this request, you have to assign it an integer that uniquely identifies it, so that Stat Server and your application can easily distinguish it from other sets of statistical information. Note that you will also need to enter this integer in the `StatisticId` field for any subsequent requests that refer to the statistics generated on the basis of the Open request.

```
[Java]

requestOpenStatistic.setReferenceId(2);
```

Now you can send the request:

```
[Java]

System.out.println("Sending:\n" + requestOpenStatistic);
statServerProtocol.send(requestOpenStatistic);
```

After Stat Server sends the `EventStatisticOpened` in response to this request, it will start sending `EventInfo` messages every 5 seconds. You need to set up an event handler to receive these messages, as discussed in the the Event Handling article.

This is what one such message might look like:

```
'EventInfo' ('2')
message attributes:
REQ_ID [int]    = 4
USER_REQ_ID [int] = -1
TM_SERVER [int] = 1244412448
TM_LENGTH [int] = 0
LONG_VALUE [int] = 0
VOID_VALUE [object] = AgentStatus {
        AgentId = Analyst001
        AgentStatus = 23
        Time = 1240840034
        PlaceStatus =   PlaceStatus = 23
        Time = 1240840034
        LoginId = LoggedOut
}
```

Creating Dynamic Statistics

As mentioned above, there may be times when you want to get statistical information that has not already been defined in the Configuration Layer. In cases like that, you can use `RequestOpenStatisticEx`. Before you do, however, you should make sure you understand several topics covered in the Reporting Technical Reference 8.0 Overview and the Framework 8.1 Stat Server User's Guide, including the use of masks.

The first things you need to do in order to use `RequestOpenStatisticEx` are similar to what we did in the previous section. You will start by creating the request and specifying the statistic object and notification mode, which you will add to the request:

```
[Java]

RequestOpenStatisticEx request =
        RequestOpenStatisticEx.create();

StatisticObject object = StatisticObject.create();
object.setObjectId("Analyst001");
object.setObjectType(StatisticObjectType.Agent);
object.setTenantName("Resources");
object.setTenantPassword("");

Notification notification = Notification.create();
notification.setMode(NotificationMode.Immediate);

request.setNotification(notification);
request.setStatisticObject(object);
```

Now, instead of requesting a pre-defined statistic type, you need to set up your own masks, as described in the section on "Metrics: Their Composition and Definition" in the Reporting Technical Reference 8.0 Overview. The following mask and statistic metric settings give the Current State for the agent mentioned above:

```
[Java]

DnActionMask mainMask = ActionsMask.createDNActionsMask();
mainMask.setBit(DnActions.WaitForNextCall);
mainMask.setBit(DnActions.CallDialing);
mainMask.setBit(DnActions.CallRinging);
mainMask.setBit(DnActions.NotReadyForNextCall);
mainMask.setBit(DnActions.CallOnHold);
mainMask.setBit(DnActions.CallUnknown);
mainMask.setBit(DnActions.CallConsult);
mainMask.setBit(DnActions.CallInternal);
mainMask.setBit(DnActions.CallOutbound);
mainMask.setBit(DnActions.CallInbound);
mainMask.setBit(DnActions.LoggedOut);

DnActionMask relMask = ActionsMask.createDNActionsMask();

StatisticMetricEx metric = StatisticMetricEx.create();
metric.setCategory(StatisticCategory.CurrentState);
metric.setMainMask(mainMask);
metric.setRelativeMask(relMask);
metric.setSubject(StatisticSubject.DNStatus);

request.setStatisticMetricEx(metric);
```

Once you have set up the masks and the statistic metric, you can create a `ReferenceId` and send the request:

```
[Java]

request.setReferenceId(anIntThatYouSpecify);

System.out.println("Sending:\n" + request);
Message response = statServerProtocol.request(request);
System.out.println("Received:\n" + response);
```

Current Target State Events

You can use RequestGetStatisticEx and RequestOpenStatisticEx to set up the same type of current target state definitions that Universal Routing Server (URS) uses. (You can also set these up using Configuration Manager.) When this type of request has been sent, Stat Server sends some additional event types:

- EventCurrentTargetStateSnapshot
- EventCurrentTargetStateTargetUpdated
- EventCurrentTargetStateTargetAdded
- EventCurrentTargetStateTargetRemoved

The Snapshot event is returned in response to the open, while the Updated event is sent as state changes occur. In a situation where you open a CurrentTargetState-based statistic against an agent group, the Added and Removed messages occur when an agent is added to or removed from an agent group — it would behave in a similar fashion for place groups.

Here is the output from a typical request:

```
'EventCurrentTargetStateSnapshot' (17) attributes:
      TM_LENGTH [int] = 0
      USER_REQ_ID [int] = -1
      LONG_VALUE [int] = 0
      CURRENT_TARGET_STATE_INFO [CurrentTargetState] = CurrentTargetStateSnapshot (size=1) [
    [0] CurrentTargetStateInfo {
      AgentId = Analyst001
      AgentDbId = 101
      LoginId = null
      PlaceId = null
      PlaceDbId = 0
      Extensions = KVList:
'VOICE_MEDIA_STATUS' [int] = 0
'AGENT_VOICE_MEDIA_STATUS' [int] = 0
    }
  ]

      REQ_ID [int] = 5
      TM_SERVER [int] = 1245182089
```

Peeking at a Statistic

There may be times when you need to get immediate information on a statistic you have opened. For example, you may want to initialize a wallboard display. In that case, you can use RequestPeekStatistic. Note that Stat Server does not send a handshake event when you use this request, so you should use the send method rather than the request method when you use it. Note also that you need to use the StatisticId property to provide the ReferenceId of the RequestOpenStatistic or RequestOpenStatisticEx associated with the statistic you want information on:

Tip

If you use the request method on a `RequestPeekStatistic`, your request will time out and receive null, rather than retrieving the desired information from Stat Server.

```
[Java]

RequestPeekStatistic req = RequestPeekStatistic.create();
req.setStatisticId(2);

System.out.println("Sending:\n" + req);
statServerProtocol.send(req);
```

Suspending Notification

Because there are times when you do not need to collect information on a statistic for a while, the Platform SDK has requests that allow you to suspend and resume notification. These requests are like the peek request in that Stat Server does not send a handshake event when you use them, so you should use the send method rather than the request method when you use these requests. Note also that you need to use the `StatisticId` property of these requests to provide the `ReferenceId` of the `RequestOpenStatistic` or `RequestOpenStatisticEx` associated with the statistic you want information on. Here is how to suspend notification:

```
[Java]

RequestSuspendNotification req = RequestSuspendNotification.create();
req.setStatisticId(2);

System.out.println("Sending:\n" + req);
statServerProtocol.send(req);
```

Use code like this to resume notification:

```
[Java]

RequestResumeNotification req = RequestResumeNotification.create();
req.setStatisticId(2);

System.out.println("Sending:\n" + req);
statServerProtocol.send(req);
```

Closing the Statistic and the Connection

When you are finished communicating with Stat Server, you should close the statistics that you have opened and close the connection, in order to minimize resource utilization:

```
[Java]

RequestCloseStatistic req = RequestCloseStatistic.create();
req.setStatisticId(2);
```

```
System.out.println("Sending:\n" + req);
statServerProtocol.send(req);

...

statServerProtocol.beginClose();
```

.NET

Connecting to Stat Server

As mentioned in the article on the architecture, the Platform SDKs uses a message-based architecture to connect to Genesys servers. This section describes how to connect to Stat Server, based on the material in the article on Connecting to a Server.

After you have set up `using` statements, the first thing you need to do is create a `StatServerProtocol` object:

```
[C#]

StatServerProtocol statServerProtocol =
       new StatServerProtocol(new Endpoint(statServerUri));
statServerProtocol.ClientId = clientID;
statServerProtocol.ClientName = clientName;
```

You can also configure your ADDP and warm standby settings at this point, as described in the Connecting to a Server article.

Once you have finished configuring your protocol object, open the connection to Stat Server:

```
[C#]

statServerProtocol.Open();
```

Working with Statistics

The Stat Server application object in the Genesys Configuration Layer comes with many predefined statistics. You can also define your own statistics using the options tab of this application object. The Platform SDK allows you to get information about any of these statistics by using `RequestOpenStatistic`. There may be times, however, when you want your application to be able to create new types of statistics dynamically. The Platform SDK also supports this, with the use of `RequestOpenStatisticEx`.

This section will show you how to use `RequestOpenStatistic` to get information on a predefined statistic. After that, we will give an example of how to use `RequestOpenStatisticEx`.

The first thing you need to do to use `RequestOpenStatistic` is to create the request:

```
[C#]

var requestOpenStatistic = RequestOpenStatistic.Create();
```

Now you need to describe the *statistics object*, that is, the object you are monitoring. This description consists of the object's Configuration Layer ID and object type, and the tenant ID and password:

```
[C#]

requestOpenStatistic.StatisticObject = StatisticObject.Create();
requestOpenStatistic.StatisticObject.ObjectId = "Analyst001";
requestOpenStatistic.StatisticObject.ObjectType = StatisticObjectType.Agent;
requestOpenStatistic.StatisticObject.TenantName = "Environment";
requestOpenStatistic.StatisticObject.TenantPassword = "";
```

Next, you will specify the `StatisticMetric` property for this statistic. A `StatisticMetric` contains information including the `StatisticType` (which must correspond to the name of the statistic definition that appears in the options tab), along with the required `TimeRangeLeft` and `TimeRangeRight` parameters.

In this case, we are asking for the total login time for an agent identified as "Analyst001":

```
[C#]

requestOpenStatistic.StatisticMetric = StatisticMetric.Create();
requestOpenStatistic.StatisticMetric.StatisticType = "TotalLoginTime";
requestOpenStatistic.StatisticMetric.TimeProfile = "Default";
// Note: if no time profile is provided, then the default is used automatically
```

Finally, specify the desired Notification settings. The Statistics Platform SDK supports four ways of gathering statistics:

- NoNotification allows you to retrieve statistics when you want them.
- Periodical means Stat Server reports on statistics based on the time period you request.
- Immediate means Stat Server reports on statistics whenever a statistical value changes. For time-related statistics, Immediate means that Stat Server will report the current value whenever a statistical value changes, but it will also report that value periodically, using the specified notification frequency.
- Reset means Stat Server reports the current value of a statistic right before setting the statistical value to zero (0).

In this case, we are interested in receiving statistics on a regular basis, so we have asked for a notification mode of Periodical, with updates every 5 seconds. For more information on notification modes, see the section on Notification Modes in Framework 8.1 Stat Server User's Guide.

```
[C#]

requestOpenStatistic.Notification = Notification.Create();
requestOpenStatistic.Notification.Mode = NotificationMode.Periodical;

requestOpenStatistic.Notification.Frequency = 5; // seconds
```

Before sending this request, you have to assign it an integer that uniquely identifies it, so that Stat Server and your application can easily distinguish it from other sets of statistical information. Note that you will also need to enter this integer in the `StatisticId` field for any subsequent requests that refer to the statistics generated on the basis of the Open request.

```
[C#]

requestOpenStatistic.ReferenceId = 3; // Must be unique and is included as StatisticId in
// Peek/Close for the stat
```

Now you can send the request:

```
[C#]

Console.WriteLine("Sending:\n{0}", requestOpenStatistic);
var response =
        statServerProtocol.Request(requestOpenStatistic);
Console.WriteLine("Received:\n{0}", response);

if (response == null || response.Id != EventStatisticOpened.MessageId)
{
        // Open failed, proper error handling goes here
        throw new Exception("RequestOpenStatistic failed.");
}

var @event = response as EventStatisticOpened;
```

After Stat Server sends the `EventStatisticOpened` in response to this request, it will start sending EventInfo messages every 5 seconds. You need to set up an event handler to receive these messages, as discussed in the the Event Handling article.

This is what one such message might look like:

```
[C#]

'EventInfo' ('2')
message attributes:
REQ_ID [int]    = 4
USER_REQ_ID [int] = -1
TM_SERVER [int] = 1244412448
TM_LENGTH [int] = 0
LONG_VALUE [int] = 0
VOID_VALUE [object] = AgentStatus {
        AgentId = Analyst001
        AgentStatus = 23
        Time = 1240840034
        PlaceStatus =   PlaceStatus = 23
        Time = 1240840034
        LoginId = LoggedOut
}
```

Creating Dynamic Statistics

As mentioned above, there may be times when you want to get statistical information that has not already been defined in the Configuration Layer. In cases like that, you can use `RequestOpenStatisticEx`. Before you do, however, you should make sure you understand several topics covered in the Reporting Technical Reference 8.0 Overview and the Framework 8.1 Stat Server User's Guide, including the use of masks.

The first things you need to do in order to use `RequestOpenStatisticEx` are similar to what we did in the previous section. You will start by creating the request and specifying the statistic object and notification mode:

```
[C#]

var req = RequestOpenStatisticEx.Create();

req.StatisticObject = StatisticObject.Create();
req.StatisticObject.ObjectId = "Analyst001";
req.StatisticObject.ObjectType = StatisticObjectType.Agent;
req.StatisticObject.TenantName = "Resources";
req.StatisticObject.TenantPassword = "";

req.Notification = Notification.Create();
req.Notification.Mode = NotificationMode.Immediate;
req.Notification.Frequency = 15;
```

Now, instead of requesting a statistic type, you need to set up your own masks, as described in the section on "Metrics: Their Composition and Definition" in the Reporting Technical Reference 8.0 Overview. The following mask and statistic metric settings give the Current State for the agent mentioned above:

```
[C#]

var mainMask = ActionsMask.CreateDnActionMask();
mainMask.SetBit(DnActions.WaitForNextCall);
mainMask.SetBit(DnActions.CallDialing);
mainMask.SetBit(DnActions.CallRinging);
mainMask.SetBit(DnActions.NotReadyForNextCall);
mainMask.SetBit(DnActions.CallOnHold);
mainMask.SetBit(DnActions.CallUnknown);
mainMask.SetBit(DnActions.CallConsult);
mainMask.SetBit(DnActions.CallInternal);
mainMask.SetBit(DnActions.CallOutbound);
mainMask.SetBit(DnActions.CallInbound);
mainMask.SetBit(DnActions.LoggedOut);

var relMask = ActionsMask.CreateDnActionMask();

req.StatisticMetricEx = StatisticMetricEx.Create();
req.StatisticMetricEx.Category = StatisticCategory.CurrentState;
req.StatisticMetricEx.IntervalLength = 0;
req.StatisticMetricEx.MainMask = mainMask;
req.StatisticMetricEx.RelativeMask = relMask;
req.StatisticMetricEx.Subject = StatisticSubject.DNStatus;
```

Once you have set up the masks and the statistic metric, you can create a `ReferenceId` and send the request:

```
[C#]

req.ReferenceId = referenceIdFromRequestOpenStatistic;

Console.WriteLine("Sending:\n{0}", req);
var response =
        statServerProtocol.Request(req);
Console.WriteLine("Received:\n{0}", response);
```

Current Target State Events

You can use `RequestGetStatisticEx` and `RequestOpenStatisticEx` to set up the same type of current target state definitions that Universal Routing Server (URS) uses. (You can also set these up using Configuration Manager.) When this type of request has been sent, Stat Server sends some additional event types:

- `EventCurrentTargetStateSnapshot`
- `EventCurrentTargetStateTargetUpdated`
- `EventCurrentTargetStateTargetAdded`
- `EventCurrentTargetStateTargetRemoved`

The `Snapshot` event is returned in response to the open, while the `Updated` event is sent as state changes occur. In a situation where you open a `CurrentTargetState`-based statistic against an agent group, the Added and Removed messages occur when an agent is added to or removed from an agent group — it would behave in a similar fashion for place groups.

Here is the output from a typical request:

```
'EventCurrentTargetStateSnapshot' (17) attributes:
      TM_LENGTH [int] = 0
      USER_REQ_ID [int] = -1
      LONG_VALUE [int] = 0
      CURRENT_TARGET_STATE_INFO [CurrentTargetState] = CurrentTargetStateSnapshot (size=1) [
    [0] CurrentTargetStateInfo {
      AgentId = Analyst001
      AgentDbId = 101
      LoginId = null
      PlaceId = null
      PlaceDbId = 0
      Extensions = KVList:
'VOICE_MEDIA_STATUS' [int] = 0
'AGENT_VOICE_MEDIA_STATUS' [int] = 0
    }
  ]

      REQ_ID [int] = 5
      TM_SERVER [int] = 1245182089
```

Peeking at a Statistic

There may be times when you need to get immediate information on a statistic you have opened. For example, you may want to initialize a wallboard display. In that case, you can use `RequestPeekStatistic`. Note that Stat Server does not send a handshake event when you use this request, so you should use the Send method rather than the `Request` method when you use it. Note also that you need to use the `StatisticId` property to provide the `ReferenceId` of the `RequestOpenStatistic` or `RequestOpenStatisticEx` associated with the statistic you want information on:

> ### Tip
> If you use the `Request` method on a `RequestPeekStatistic`, your request will time out and receive null, rather than retrieving the desired information from Stat Server.

[C#]

```
var requestPeekStatistic = RequestPeekStatistic.Create();
requestPeekStatistic.StatisticId = 3;

Console.WriteLine("Sending:\n{0}", requestPeekStatistic);
statServerProtocol.Send(requestPeekStatistic);
```

Suspending Notification

Because there are times when you do not need to collect information on a statistic for a while, the Platform SDK has requests that allow you to suspend and resume notification. These requests are like the peek request in that Stat Server does not send a handshake event when you use them, so you should use the send method rather than the request method when you use these requests. Note also that you need to use the `StatisticId` property of these requests to provide the `ReferenceId` of the `RequestOpenStatistic` or `RequestOpenStatisticEx` associated with the statistic you want information on. Here is how to suspend notification:

[C#]

```
var requestSuspendNotification = RequestSuspendNotification.Create();
requestSuspendNotification.StatisticId = 3;

Console.WriteLine("Sending:\n{0}", requestSuspendNotification);
statServerProtocol.Send(requestSuspendNotification);
```

Use code like this to resume notification:

[C#]

```
var requestResumeNotification = RequestResumeNotification.Create();
requestResumeNotification.StatisticId = 3;

Console.WriteLine("Sending:\n{0}", requestResumeNotification);
statServerProtocol.Send(requestResumeNotification);
```

Closing the Statistic and the Connection

When you are finished communicating with Stat Server, you should close the statistics that you have opened and close the connection, in order to minimize resource utilization:

[C#]

```
var requestCloseStatistic = RequestCloseStatistic.Create();
requestCloseStatistic.StatisticId = 3;

Console.WriteLine("Sending:\n{0}", requestCloseStatistic);
statServerProtocol.Send(requestCloseStatistic);

...

statServerProtocol.BeginClose();
```

Interaction Server

You can use the Open Media Platform SDK to write Java or .NET applications that handle third-party work items in conjunction with the Genesys Interaction Server. You can also use it to work with servers that implement the Genesys External Service Protocol.

This document shows how to implement the basic functions you will need to write simple Interaction Server–based email applications. The first application is a simple media server that submits a new third-party work item. The second application enables an agent to receive a third-party work item, accept it for processing, and mark it done.

Java

Setting Up Interaction Server Protocol Objects

The first thing you need to do to use the Open Media Platform SDK is instantiate a Protocol object. To do that, you must supply information about the server you want to connect with. This example uses an `InteractionServerProtocol` object, supplying its URI, but you can also use name, host, and port information:

```
[Java]

InteractionServerProtocol interactionServerProtocol =
        new InteractionServerProtocol(
                new Endpoint(
                        InteractionServerUri));
```

After instantiating the `InteractionServerProtocol` object, you need to open a connection to Interaction Server:

```
[Java]

interactionServerProtocol.open();
```

Creating a Simple Media Server

The Open Media Platform SDK makes it easy to write a simple server that can submit third-party work items to Interaction Server. To write one, start by entering configuration information:

```
[Java]

// Enter configuration information here:
private String interactionServerName = "<server name>";
private String interactionServerHost = "<host>";
private int interactionServerport = <port>;
```

```
private int tenantId = 101;
private String inboundQueue = "<queue>";
private String mediaType = "<media type>";
// End of configuration information.
```

Now you will need to set up a protocol object:

```
[Java]

interactionServerUri = new Uri("tcp://"
        + interactionServerHost + ":"
        + interactionServerport);
InteractionServerProtocol interactionServerProtocol =
        new InteractionServerProtocol(
        new Endpoint(interactionServerName, interactionServerUri));
```

Once you have set up the protocol object, you can tell it the name of your application and let it know that it is a media server:

```
[Java]

interactionServerProtocol.setClientName("EntityListener");
interactionServerProtocol.setClientType(
        InteractionClient.MediaServer);
```

At this point, you can add user data associated with the new interaction:

```
[Java]

KeyValueCollection userData =
        new KeyValueCollection();

userData.add("Subject",
        "New Interaction Created by a Custom Media Server");
```

Now you can open the protocol object, and prepare the interaction to be submitted:

```
[Java]

try
{
        interactionServerProtocol.open();

        RequestSubmit requestSubmit = RequestSubmit.create(
                inboundQueue,
                mediaType,
                "Inbound");
        requestSubmit.setTenantId(tenantId);
        requestSubmit.setInteractionSubtype("InboundNew");
        requestSubmit.setUserData(userData);
```

If you use the `Request` method, you will receive a synchronous response containing a message from Interaction Server:

```
[Java]

Message response =
    interactionServerProtocol.request(requestSubmit);
System.out.println("Response: " + response.messageName() + ".\n\n");
```

Closing the Connection

Finally, when you are finished communicating with Interaction Server, you should close the connection to minimize resource utilization:

```
[Java]

interactionServerProtocol.close();
```

.NET

Setting Up Interaction Server Protocol Objects

The first thing you need to do to use the Open Media Platform SDK is instantiate a Protocol object. To do that, you must supply information about the server you want to connect with. This example uses an `InteractionServerProtocol` object, supplying its URI, but you can also use name, host, and port information:

```
[C#]

InteractionServerProtocol interactionServerProtocol =
        new InteractionServerProtocol(
                new Endpoint(
                        InteractionServerUri));
```

After instantiating the `InteractionServerProtocol` object, you need to open a connection to Interaction Server:

```
[C#]

interactionServerProtocol.Open();
```

Creating a Simple Media Server

The Open Media Platform SDK makes it easy to write a simple server that can submit third-party work items to Interaction Server. To write one, start by entering configuration information:

```
[C#]

// Enter configuration information here:
private string interactionServerName = "<server name>";
private string interactionServerHost = "<host>";
private int interactionServerport = <port>;
private int tenantId = 101;
private string inboundQueue = "<queue>";
private string mediaType = "<media type>";
// End of configuration information.
```

Now you will need to set up a protocol object:

```
[C#]

interactionServerUri = new Uri("tcp://"
        + interactionServerHost + ":"
        + interactionServerport);
InteractionServerProtocol interactionServerProtocol =
        new InteractionServerProtocol(
        new Endpoint(interactionServerName, interactionServerUri));
```

Once you have set up the protocol object, you can tell it the name of your application and let it know that it is a media server:

```
[C#]

interactionServerProtocol.ClientName = "EntityListener";
interactionServerProtocol.ClientType =
        InteractionClient.MediaServer;
```

At this point, you can add user data associated with the new interaction:

```
[C#]

KeyValueCollection userData =
        new KeyValueCollection();

userData.Add("Subject",
        "New Interaction Created by a Custom Media Server");
```

Now you can open the protocol object, and prepare the interaction to be submitted:

```
[C#]

try
{
        interactionServerProtocol.Open();

        RequestSubmit requestSubmit = RequestSubmit.Create(
                inboundQueue,
                mediaType,
                "Inbound");
        requestSubmit.TenantId = tenantId;
        requestSubmit.InteractionSubtype = "InboundNew";
        requestSubmit.UserData = userData;
```

If you use the `Request` method, you will receive a synchronous response containing a message from Interaction Server:

```
[C#]

IMessage response =
    interactionServerProtocol.Request(requestSubmit);
LogAreaRichTextBox.Text = LogAreaRichTextBox.Text
    + "Response: " + response.Name + ".\n\n";
```

Closing the Connection

Finally, when you are finished communicating with Interaction Server, you should close the connection to minimize resource utilization:

```
[C#]

interactionServerProtocol.Close();
```

Additional Topics

As support for the Platform SDKs continues to grow, new topics and examples that illustrate best-practice approaches to common tasks are being added to the documentation. For more information about using the Open Media Platform SDK, including functional code snippets, please read the following topics:

- Creating an E-Mail - This article discusses how to use the Open Media and Contacts Platform SDKs in conjuction to create outgoing e-mail messages. You can also apply the concepts illustrated here to other types of Interactions.

Universal Contact Server

You can use the Contacts Platform SDK to write Java or .NET applications that interact with the Genesys Universal Contact Server (UCS). This allows you to create applications that work with contacts, interactions, and standard responses in a variety of ways - either to create a full-featured agent desktop, or a simple application that forwards e-mail messages.

This document shows how to implement the basic functions you will need to write simple UCS-based applications.

When you are ready to write more complicated applications, take a look at the classes and methods described in the Platform SDK API Reference.

Java

Using the Contacts Protocols

Before using the Contacts Platform SDK, you should include `import` statements that allow access to the Platform SDK Commons and Contacts classes:

```
[Java]

import com.genesyslab.platform.commons.protocol.*;

import com.genesyslab.platform.contacts.protocol.*;
import com.genesyslab.platform.contacts.protocol.contactserver.*;
import com.genesyslab.platform.contacts.protocol.contactserver.events.*;
import com.genesyslab.platform.contacts.protocol.contactserver.requests.*;
```

Setting Up Universal Contact Server Protocol Objects

The first thing you need to do to use the Contacts Platform SDK is instantiate a `UniversalContactServerProtocol` object. To do that, you must supply information about the Universal Contact Server you want to connect with. This example uses the server's name, host, and port information, but you can also use just the URI of your Universal Contact Server:

```
[Java]

UniversalContactServerProtocol ucsConnection =
```

new UniversalContactServerProtocol(new Endpoint(universalContactServerURI));

It is a good practice to always set the application name at the same time that you instantiate a new protocol object. This application name will be used to identify where UCS requests came from.

This is also a good time to add event handlers to the protocol object. See the Event Handling section in this introductory material for code samples and details.

```
[Java]

// Set the ApplicationName property
ucsConnection.setApplicationName("IntroducingContactsPSDK");
```

After setting up your protocol object, the code to open a connection to the server is simple:

```
[Java]

ucsConnection.open();
```

> **Tip**
>
> Be sure to use proper error handling techniques in your code, especially when working with the protocol connection. To save space, these error handling steps are not shown in this example.

Inserting an Interaction

Now that the protocol connection is open, you are ready to start handling interactions. In this example, we will start by creating a new, outbound e-mail interaction using the `RequestInsertInteraction` request.

Creating a new e-mail interaction object takes a bit of planning. Before you can create and submit the request object, you need to create and configure the following objects:

- `InteractionAttributes` - Sets common attributes for this interaction, specifying details such as the media type and status. All interactions need these attributes to be configured.
- `EmailOutEntityAttributes` - Sets attributes that are specific to an outbound e-mail interaction. For outbound e-mail interactions, this includes the sending and receiving addresses. (The type of interaction you are creating will dictate which object to use here; for example, phone interactions require a `PhoneCallEntityAttributes` object instead of `EmailOutEntityAttributes`.)
- `InteractionContent` - Specifies the actual interaction content. This can be Text, MIME, StructuredText, or StructuredText with MIME content.

The following code snippet shows how each of these objects is configured for our simple outbound e-mail example:

```
[Java]

// Set common interaction attributes
InteractionAttributes attributes = new InteractionAttributes();
attributes.setTenantId(101);
attributes.setMediaTypeId("email");
attributes.setTypeId("Outbound");
attributes.setSubtypeId("OutboundRedirect");
attributes.setStatus(Statuses.Pending);
```

```
attributes.setSubject(subjectLine);
attributes.setQueueName(queueName);
attributes.setEntityTypeId(EntityTypes.EmailOut);

// Set entity-specific attributes
EmailOutEntityAttributes outEntityAttributes = new EmailOutEntityAttributes();
outEntityAttributes.setFromAddress(fromAddress);
outEntityAttributes.setToAddresses(forwardAddress);

// Set interaction content
InteractionContent content = new InteractionContent();
content.setText("E-mail message text...");
```

Tip

The `InteractionAttributes` class stores the `StartDate` property in UTC format. If no value is provided, UCS uses the current date.

Once you have configured the attributes and content for the interaction, it is easy to create and submit the new request:

```
[Java]

// Create the new interaction request
RequestInsertInteraction request = new RequestInsertInteraction();
request.setInteractionAttributes(attributes);
request.setEntityAttributes(outEntityAttributes);
request.setInteractionContent(content);

// Submit the request
EventInsertInteraction eventInsertIxn = ucsConnection.request(request);
```

Adding an Attachment

Now that you know how to create new e-mail interactions, it is the perfect time to learn how to add attachments to existing interactions. The process for this is much easier than creating a new interaction; you just need to create the request and specify the attachment properties as shown in the code snippet below. Once the request is ready, submit it to your UCS protocol object.

```
[Java]

RequestAddDocument request = new RequestAddDocument();
request.setInteractionId(eventInsertIxn.getInteractionId());
request.setDocumentId(strDocumentId);
request.setDescription(strDescription);
request.setMimeType(strMimeType);
request.setTheName(strName);
request.setTheSize(intSize);

EventAddDocument eventAddDocument = ucsConnection.request(request);
```

Note that before adding an attachment, you need to have the Interaction ID available. In our example, the Interaction ID was returned as part of the `EventInsertInteraction` from the previous section. Otherwise we would need to submit a `RequestGetInteractionContent` request and then take the Interaction ID from the resulting event.

Getting an Interaction from UCS

Now that we have created a new Interaction and submitted it to UCS, what happens next? The final task we will cover in this introduction is how to return the Interaction and any of its attachments for processing.

The structure of `RequestGetInteractionContent` is very basic: set the Interaction ID you are looking for, and then use the `IncludeAttachments` and `IncludeBinaryContent` properties to specify what type of content you want to be returned. In this example, we will return the attachment created previously and store it in an `Attachment` object for later use.

```
[Java]

RequestGetInteractionContent request = new RequestGetInteractionContent();
request.setInteractionId(eventInsertIxn.getInteractionId());
request.setIncludeAttachments(true);

EventGetInteractionContent eventGetIxnContent = ucsConnection.Request(request);

String subject = eventGetIxnContent.getInteractionAttributes().getSubject();
String key = eventGetIxnContent.getInteractionAttributes().getId();
if (eventGetIxnContent.getAttachments() != null)
{
        Attachment attachedFile = eventGetIxnContent.getAttachments().get(0);
}
```

Closing the Connection

Finally, when you are finished communicating with the server, you should close the connection and dispose of the object to minimize resource utilization:

```
[Java]

if (ucsConnection.getState() != ChannelState.Closed && ucsConnection.getState() !=
ChannelState.Closing)
        {
                ucsConnection.close();
        }
```

.NET

Using the Contacts Protocols

Before using the Contacts Platform SDK, you should include using statements that allow access to the Platform SDK Commons and Contacts namespaces:

```
[C#]

using Genesyslab.Platform.Commons.Protocols;
using Genesyslab.Platform.Contacts.Protocols;
using Genesyslab.Platform.Contacts.Protocols.ContactServer;
using Genesyslab.Platform.Contacts.Protocols.ContactServer.Requests;
using Genesyslab.Platform.Contacts.Protocols.ContactServer.Events;
```

Setting Up Universal Contact Server Protocol Objects

The first thing you need to do to use the Contacts Platform SDK is instantiate a `UniversalContactServerProtocol` object. To do that, you must supply information about the Universal Contact Server you want to connect with. This example uses the server's name, host, and port information, but you can also use just the URI of your Universal Contact Server:

```
[C#]

UniversalContactServerProtocol ucsConnection;
ucsConnection = new UniversalContactServerProtocol(new Endpoint("UCS", ucsHost, ucsPort));
```

It is a good practice to always set the application name at the same time that you instantiate a new protocol object. This application name will be used to identify where UCS requests came from.

This is also a good time to add event handlers to the protocol object. See the Event Handling article for details.

```
[C#]

// Set the ApplicationName property
ucsConnection.ApplicationName = "IntroducingContactsPSDK";

// Add event handlers
ucsConnection.Opened += new EventHandler(ucsConnection_Opened);
ucsConnection.Error += new EventHandler(ucsConnection_Error);
ucsConnection.Closed += new EventHandler(ucsConnection_Closed);
```

After setting up your protocol object, the code to open a connection to the server is simple:

```
[C#]

ucsConnection.Open();
```

> **Tip**
>
> Be sure to use proper error handling techniques in your code, especially when working with the protocol connection. To save space, these error handling steps are not shown

> in this example.

Inserting an Interaction

Now that the protocol connection is open, you are ready to start handling interactions. In this example, we will start by creating a new, outbound e-mail interaction using the `RequestInsertInteraction` request.

Creating a new e-mail interaction object takes a bit of planning. Before you can create and submit the request object, you need to create and configure the following objects:

- `InteractionAttributes` - Sets common attributes for this interaction, specifying details such as the media type and status. All interactions need these attributes to be configured.
- `EmailOutEntityAttributes` - Sets attributes that are specific to an outbound e-mail interaction. For outbound e-mail interactions, this includes the sending and receiving addresses. (The type of interaction you are creating will dictate which object to use here; for example, phone interactions require a `PhoneCallEntityAttributes` object instead of `EmailOutEntityAttributes`.)
- `InteractionContent` - Specifies the actual interaction content. This can be Text, MIME, StructuredText, or StructuredText with MIME content.

The following code snippet shows how each of these objects is configured for our simple outbound e-mail example:

```
[C#]

// Set common interaction attributes
InteractionAttributes attributes = new InteractionAttributes();
attributes.TenantId = 101;
attributes.MediaTypeId = "email";
attributes.TypeId = "Outbound";
attributes.SubtypeId = "OutboundRedirect";
attributes.Status = new NullableStatuses(Statuses.Pending);
attributes.Subject = subjectLine;
attributes.QueueName = queueName;
attributes.EntityTypeId = new NullableEntityTypes(EntityTypes.EmailOut);

// Set entity-specific attributes
EmailOutEntityAttributes outEntityAttributes = new EmailOutEntityAttributes();
outEntityAttributes.FromAddress = fromAddress;
outEntityAttributes.ToAddresses = forwardAddress;

// Set interaction content
InteractionContent content = new InteractionContent();
content.Text = "E-mail message text...";
```

> ### Tip
> The `InteractionAttributes` class stores the `StartDate` property in UTC format. If no

> value is provided, UCS uses the current date.

Once you have configured the attributes and content for the interaction, it is easy to create and submit the new request:

```
[C#]

// Create the new interaction request
RequestInsertInteraction request = new RequestInsertInteraction();
request.InteractionAttributes = attributes;
request.EntityAttributes = outEntityAttributes;
request.InteractionContent = content;

// Submit the request
EventInsertInteraction eventInsertIxn = ucsConnection.Request(request);
```

Adding an Attachment

Now that you know how to create new e-mail interactions, it is the perfect time to learn how to add attachments to existing interactions. The process for this is much easier than creating a new interaction; you just need to create the request and specify the attachment properties as shown in the code snippet below. Once the request is ready, submit it to your UCS protocol object.

```
[C#]

RequestAddDocument request = new RequestAddDocument();
request.InteractionId = eventInsertIxn.InteractionId;
request.DocumentId = strDocumentId;
request.Description = strDescription;
request.MimeType = strMimeType;
request.TheName = strName;
request.TheSize = intSize;

EventAddDocument eventAddDocument = ucsConnection.Request(request);
```

Note that before adding an attachment, you need to have the Interaction ID available. In our example, the Interaction ID was returned as part of the `EventInsertInteraction` from the previous section. Otherwise we would need to submit a `RequestGetInteractionContent` request and then take the Interaction ID from the resulting event.

Getting an Interaction from UCS

Now that we have created a new Interaction and submitted it to UCS, what happens next? The final task we will cover in this introduction is how to return the Interaction and any of its attachments for processing.

The structure of `RequestGetInteractionContent` is very basic: set the Interaction ID you are looking for, and then use the `IncludeAttachments` and `IncludeBinaryContent` properties to specify what type of content you want to be returned. In this example, we will return the attachment created previously and store it in an `Attachment` object for later use.

```
[C#]

RequestGetInteractionContent request = new RequestGetInteractionContent();
request.InteractionId = eventInsertIxn.InteractionId;
request.IncludeAttachments = true;

EventGetInteractionContent eventGetIxnContent = ucsConnection.Request(request);

String subject = eventGetIxnContent.InteractionAttributes.Subject;
String key = eventGetIxnContent.InteractionAttributes.Id;
if (eventGetIxnContent.Attachments != null)
{
        Attachment attachedFile = eventGetIxnContent.Attachments.Get(0);
}
```

Closing the Connection

Finally, when you are finished communicating with the server, you should close the connection and dispose of the object to minimize resource utilization:

```
[C#]

if (ucsConnection.State != ChannelState.Closed && ucsConnection.State != ChannelState.Closing)
        {
                ucsConnection.Close();
                ucsConnection.Dispose();
        }
```

Additional Topics

As support for the Platform SDKs continues to grow, new topics and examples that illustrate best-practice approaches to common tasks are being added to the documentation. For more information about using the Contacts Platform SDK, including functional code snippets, please read the following topics:

- Creating an E-Mail - This article discusses how to use the Open Media and Contacts Platform SDKs in conjuction to create outgoing e-mail messages.

Creating an E-Mail

Java

This article discusses the general process used to create e-mail messages, and provides suggestions about how you should work with those protocols.

Overview of Creating a New E-Mail Message

To create a new e-mail message, there are four basic steps you should follow:

1. Connect to Genesys Servers - Use the Protocol Manager Application Block to access the appropriate Genesys Servers.
2. Create a new Interaction - Request a new Interaction that will be used to manage the e-mail message within Interaction Server.
3. Store e-mail details in UCS - Once the Interaction is available, you can use the unique `InteractionId` that is returned to create a new UCS entry that contains details and contents for the e-mail message.
4. Place the Interaction in the appropriate queue - When both parts of the e-mail message have been stored, move the Interaction into the correct queue for processing.

A quick overview of these steps, and an outline of the key requests sent to Genesys servers, is shown below.

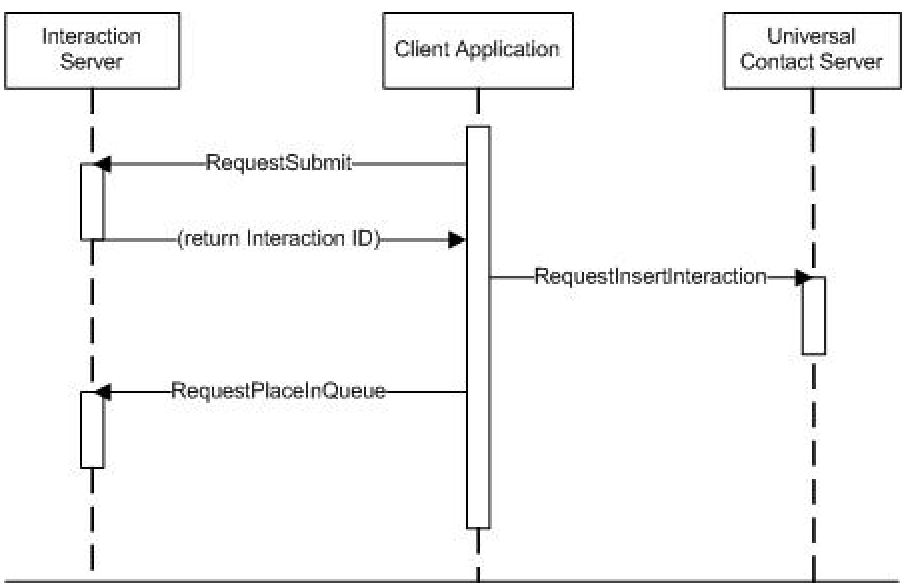

> **Tip**
>
> The order of the second and third steps can be reversed, if desired, as long as the final UCS entry contains the correct InteractionId value. In this case you would need to update the UCS entry after creating the new Interaction.

The following sections include code snippets that show one possible approach for handling each of these steps. The snippets have been simplified to focus only on code related to Genesys-specific functions.

Connecting to Genesys Servers

When creating and handling e-mail interactions, it is important to remember how e-mail messages are stored within the Genesys environment, and which Genesys servers you are interacting with.

Each e-mail message is stored as two separate pieces: an Interaction, and an entry in the Universal Contact Server (UCS) database. The e-mail is represented as an Interaction so that it can be sorted and processed using queues that have defined behavior. Even though e-mails are managed through Interaction Server, the actual contents and subject matter of each message must be stored in the UCS database. Any attempt to create or handle e-mail messages will require access to both Genesys Servers: Interaction Server (using the Open Media protocol) and UCS (using the Contacts Platform SDK protocol).

Before writing your e-mail application, some fairly standard code must be added to allow access to these Genesys servers. First, all necessary references and import statements must be added to your project. This includes the two specific protocols mentioned above, together with some common Genesys libraries and the Protocol Manager Application Block.

With those statements in place, we configure the Protocol Manager Application Block to handle communication with Genesys servers using the `ProtocolManagementServiceImpl` object, which is defined and configured as shown below.

```
[Java]

private InteractionServerProtocol interactionServerProtocol;
private UniversalContactServerProtocol contactServerProtocol;

public void connectToProtocols() throws URISyntaxException, ProtocolException
{
    Endpoint interactionServerEndpoint = new Endpoint(new URI("tcp://ixnServer:7005"));
    interactionServerProtocol = new InteractionServerProtocol(interactionServerEndpoint);
    interactionServerProtocol.setClientName("EmailSample");
    interactionServerProtocol.setClientType(InteractionClient.AgentApplication);

    Endpoint contactServerEndpoint = new Endpoint(new URI("tcp://ucsServer:7006"));
    contactServerProtocol = new UniversalContactServerProtocol(contactServerEndpoint);
    contactServerProtocol.setClientName("EmailSample");

    interactionServerProtocol.beginOpen();
    contactServerProtocol.beginOpen();
}
```

For more information about the Protocol Manager Application Block, see the Connecting to a Server article found in this guide.

Creating an Interaction

With connections to the Genesys servers established, we are ready to request a new Interaction that will represent our e-mail message in Interaction Server. You accomplish this by creating a new `RequestSubmit`, setting a few parameters to indicate that this Interaction represents an e-mail message, and then sending the request to Interaction Server with your `ProtocolManagementService` object.

```
[Java]

public void createInteraction(String ixnType, String ixnSubtype, String queue) throws
Exception
{
    RequestSubmit req = RequestSubmit.create();
    req.setInteractionType(ixnType);
    req.setInteractionSubtype(ixnSubtype);
    req.setQueue(queue);
    req.setMediaType("email");

    Message response = interactionServerProtocol.request(req);
    if(response == null || response.messageId() != EventAck.ID) {
          // For this sample, no error handling is implemented
          return;
    }

    EventAck event = (EventAck)response;
    mInteractionId = event.getExtension().getString("InteractionId");
}
```

A full list of properties that need to be set is included in the table below. Note that the `InteractionType` and `InteractionSubtype` properties must match existing business attributes, as specified in Configuration Server.

Property Name	Description
InteractionType	Interaction type for this e-mail message. Must match an Interaction Type Business Attribute, as specified in Configuration Server.
InteractionSubtype	Interaction subtype for this e-mail message. Must match an Interaction Subtype Business Attribute, as specified in Configuration Server.
Queue	Queue that this Interaction will be placed in initially. Must be defined in Configuration Server. When creating a new e-mail Interaction, the initial queue should not process the message (because additional information needs to be stored in UCS first).
MediaType	Primary media type of the interaction that is being submitted to Interaction Server. Intended for Media Server.

Once a response is received from Interaction Server, you can confirm that an `EventAck` response was returned and that the request was processed successfully. If an `EventError` response is returned instead, then you will need to implement some error handling code.

It is also important to save and track the `InteractionId` value of the newly created Interaction. This

ID needs to be specified in UCS entries that hold details related to the e-mail message, and is also required for moving the Interaction to an appropriate queue when you are ready to process the e-mail. In this example we are storing the `InteractionId` value in a simple variable named `mInteractionId`, which is assumed to be defined for your project. In larger samples (or full projects), a more robust way of tracking and handling Interactions may be required.

Storing E-Mail Details in UCS

With the ID of your newly created Interaction available, it is time to store details about the e-mail you are sending in the UCS database.

There are three types of information that must be stored in the UCS database:

- Interaction Attributes - Define details about the related Interaction for this information.
- Entity Attributes - Define where the e-mail message is coming from and going to. You will use `EmailOutEntityAttributes` for storing outbound e-mail messages, and `EmailInEntityAttributes` for storing inbound e-mail messages.
- Interaction Content - Define the actual contents of the email message, including the main text and any MIME attachments.

Creating and configuring a `RequestInsertInteraction` object with this information can be easily accomplished, as shown below.

```
[Java]

public void storeDetails(String ixnType, String ixnSubtype) throws Exception
{
    // Set Interaction Attributes
    InteractionAttributes ixnAttributes = new InteractionAttributes();
    ixnAttributes.setId(mInteractionId);
    ixnAttributes.setMediaTypeId("email");
    ixnAttributes.setTypeId(ixnType);
    ixnAttributes.setSubtypeId(ixnSubtype);
    ixnAttributes.setTenantId(101);
    ixnAttributes.setStatus(Statuses.Pending);
    ixnAttributes.setSubject("Sample e-mail subject");
    ixnAttributes.setEntityTypeId(EntityTypes.EmailOut);

    // Set Entity Attributes
    EmailOutEntityAttributes entityAttributes = new EmailOutEntityAttributes();
    entityAttributes.setFromAddress("sending@email.com");
    entityAttributes.setToAddresses("receiving@email.com");
    entityAttributes.setCcAddresses("copying@email.com");
    ...

    // Set Interaction Content
    InteractionContent content = new InteractionContent();
    content.setText("This is the e-mail body.");
    ...

    // Send the request
    RequestInsertInteraction req = new RequestInsertInteraction();
    req.setInteractionAttributes(ixnAttributes);
    req.setEntityAttributes(entityAttributes);
    req.setInteractionContent(content);

    contactServerProtocol.send(req);
```

```
}
```

A list of `InteractionAttributes` properties that need to be set for an email message is provided in the following table. The properties shown for `EmailOutEntityAttributes` and `InteractionContent` represent some of those most commonly used with email. Please check the documentation provided for each class to see a full list of available properties.

Interaction Attribute Name	Description
EntityTypeId	Indicates whether this is an outgoing or incoming e-mail.
Id	Interaction ID of the related Interaction record, created earlier.
MediaTypeId	Primary media type of the Interaction you are submitting to Interaction Server. Intended for Media Server.
Subject	Subject line for this e-mail message.
SubtypeId	Interaction subtype for this e-mail message. Must match an Interaction Subtype Business Attribute, as specified in Configuration Server.
Status	Current status of the e-mail message.
TenantId	ID of the Tenant where this e-mail belongs.
TypeId	Interaction type for this e-mail message. Must match an Interaction Type Business Attribute, as specified in Configuration Server.

Placing the Interaction in the Appropriate Queue

When an Interaction has been created to handle the e-mail, and all content has been stored in the UCS database, you are free to begin processing the message as you would process any normal Interaction. This is accomplished by moving the Interaction that you created into the appropriate queue for e-mail processing, as defined in Interaction Routing Designer.

```
[Java]

public void placeInQueue(String queue) throws Exception
{
    RequestPlaceInQueue req = RequestPlaceInQueue.create();
    req.setInteractionId(mInteractionId);
    req.setQueue(queue);

    interactionServerProtocol.send(req);
}
```

Replying to an E-Mail Message

Replying to an existing e-mail message follows the same basic process outlined above, but requires a few additional parameters to be set in your requests. These changes are described in the following subsections.

Changes to Creating an Interaction

When creating the Interaction, you need to specify one additional parameter before submitting your `RequestSubmit`. Take the `InteractionId` of the Interaction that represents the original e-mail message, and use that value as the `ParentInteractionId` parameter in your request, as shown below:

```
[Java]

RequestSubmit req = RequestSubmit.create();

...

req.setParentInteractionId = parentInteractionId;
```

The following table describes these additional attributes.

Attribute Name	Description
ParentInteractionId	`InteractionId` of a parent e-mail Interaction. Only set this value when replying to an existing e-mail message.

Changes to Storing E-Mail Details in UCS

When storing e-mail details in UCS, you need to specify values for three additional interaction attributes before sending your `RequestInsertInteraction`. These attributes (shown in the code snippet below) provide a link between the parent entry in UCS and any related children, as well as specifying a common thread ID.

```
[Java]

InteractionAttributes ixnAttributes = new InteractionAttributes();

...

ixnAttributes.setParentId(parentInteractionId);
ixnAttributes.setCanBeParent(False);
ixnAttributes.setThreadId(parentThreadId);
```

The table below describes these additional attributes.

Attribute Name	Description
CanBeParent	Boolean value that indicates whether this message can be a parent.
ParentId	Interaction ID for the parent e-mail Interaction.
ThreadId	Unique value that is shared between all UCS entries in an e-mail conversation.

Other Considerations

Although this introduction to creating and handling e-mail messages is not intended to be a comprehensive guide, it is useful to quickly introduce some other considerations and basic concepts regarding how requests are submitted and how errors should be handled.

- The first consideration to take into account is how you submit requests using the Protocol Management Application Block. In the code provided here, a simple send method is used to submit most requests without waiting for a response from the server. However, in more complicated samples or implementations you may need to process responses, or store and use values returned (such as the `InteractionId` in this example) once a request is processed.
  Please read the article on Event Handling provided in this document for a better understanding of how to handle incoming responses in both a synchronous and asynchronous fashion. This allows better error handling to be implemented if a request fails.
- A second consideration to be aware of is how records in Interaction Server and UCS are related when implementing error handling. If you have already created a new Interaction when your `RequestInsertInteraction` fails, then you will need to either resubmit the UCS record or delete the related Interaction by submitting a `RequestStopProcessing`. (If you reversed the steps shown here and created a UCS record first, then the same concept applies for removing that record when a new Interaction request fails.)

.NET

Overview of Creating a New E-Mail Message

To create a new e-mail message, there are four basic steps you should follow:

1. Connect to Genesys Servers - Use the Protocol Manager Application Block to access the appropriate Genesys Servers.
2. Create a new Interaction - Request a new Interaction that will be used to manage the e-mail message within Interaction Server.
3. Store e-mail details in UCS - Once the Interaction is available, you can use the unique InteractionId that is returned to create a new UCS entry that contains details and contents for the e-mail message.
4. Place the Interaction in the appropriate queue - When both parts of the e-mail message have been stored, move the Interaction into the correct queue for processing.

A quick overview of these steps, and an outline of the key requests sent to Genesys servers, is shown below.

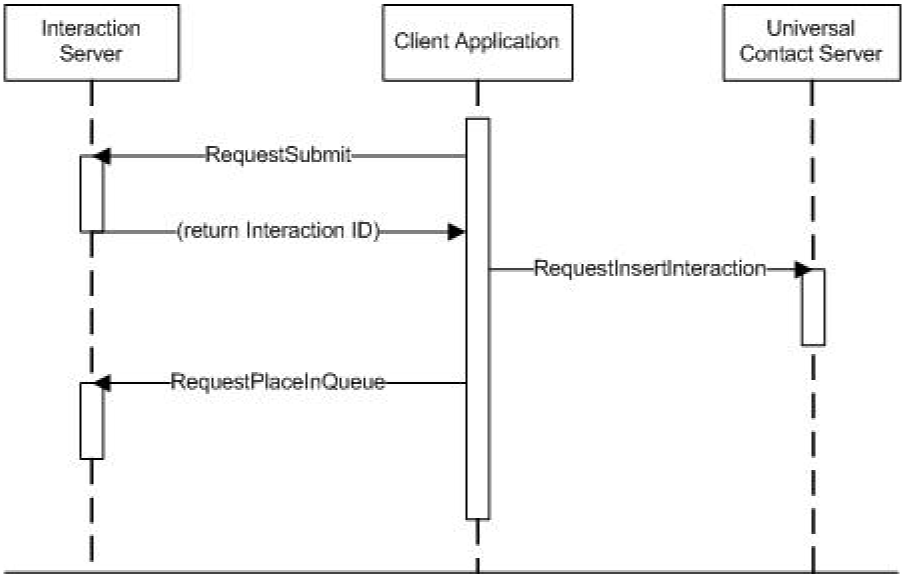

Tip

The order of the second and third steps can be reversed, if desired, as long as the final UCS entry contains the correct `InteractionId` value. In this case you would need to update the UCS entry after creating the new Interaction.

The following sections include code snippets that show one possible approach for handling each of these steps. The snippets have been simplified to focus only on code related to Genesys-specific functions.

Connecting to Genesys Servers

When creating and handling e-mail interactions, it is important to remember how e-mail messages are stored within the Genesys environment, and which Genesys servers you are interacting with.

Each e-mail message is stored as two separate pieces: an Interaction, and an entry in the Universal Contact Server (UCS) database. The e-mail is represented as an Interaction so that it can be sorted and processed using queues that have defined behavior. Even though e-mails are managed through Interaction Server, the actual contents and subject matter of each message must be stored in the UCS database. Any attempt to create or handle e-mail messages will require access to both Genesys Servers: Interaction Server (using the Open Media protocol) and UCS (using the Contacts Platform SDK protocol).

Before writing your e-mail application, some fairly standard code must be added to allow access to these Genesys servers. First, all necessary references and using statements must be added to your project.

```
[C#]

private InteractionServerProtocol interactionServerProtocol;
private UniversalContactServerProtocol contactServerProtocol;
```

```
public void ConnectToProtocols()
{
        var interactionServerEndpoint = new Endpoint(new Uri("tcp://ixnServer:7005"));
        interactionServerProtocol = new InteractionServerProtocol(interactionServerEndpoint);
        interactionServerProtocol.ClientName = "EmailSample";
        interactionServerProtocol.ClientType = InteractionClient.AgentApplication;

        var contactServerEndpoint = new Endpoint(new Uri("tcp://ucsServer:7006"));
        contactServerProtocol = new UniversalContactServerProtocol(contactServerEndpoint);
        contactServerProtocol.ClientName = "EmailSample";

        interactionServerProtocol.BeginOpen();
        contactServerProtocol.BeginOpen();
}
```

Creating an Interaction

With connections to the Genesys servers established, we are ready to request a new Interaction that will represent our e-mail message in Interaction Server. All you need to do to accomplish this is to create a new `RequestSubmit`, set a few parameters to indicate that this Interaction represents an e-mail message, and then use your `InteractionServerProtocol` object to send that request to Interaction Server.

Unlike other requests shown in this article, `RequestSubmit` is sent using the `BeginRequest` method so that we can receive and process the response from Interaction Server.

```
[C#]

public void CreateInteraction(string ixnType, string ixnSubtype, string queue)
{
var req = RequestSubmit.Create();
req.InteractionType = ixnType;
req.InteractionSubtype = ixnSubtype;
req.MediaType = "email";
req.Queue = queue;

interactionServerProtocol.BeginRequest(req, new AsyncCallback(OnCreateInteractionComplete),
null);
}
```

A full list of properties that need to be set is included in the following table. Note that the `InteractionType` and `InteractionSubtype` properties must match existing business attributes, as specified in Configuration Server.

Property Name	Description
InteractionSubtype	Interaction subtype for this e-mail message. Must match an Interaction Subtype Business Attribute, as specified in Configuration Server.
InteractionType	Interaction type for this e-mail message. Must match an Interaction Type Business Attribute, as specified in Configuration Server.
MediaType	Primary media type of the interaction that is being submitted to Interaction Server. Intended for Media

Property Name	Description
	Server.
Queue	Queue that this Interaction will be placed in initially. Must be defined in Configuration Server. When creating a new e-mail Interaction, the initial queue should not process the message (because additional information needs to be stored in UCS first).

Once a response is received from Interaction Server, you can confirm that an `EventAck` response was returned and that the request was processed successfully. If an `EventError` response is returned instead, then you will need to implement some error handling code.

You should also save and track the `InteractionId` value of the newly created Interaction. This ID needs to be specified in UCS entries that hold details related to the e-mail message, and is also required for moving the Interaction to an appropriate queue when you are ready to process the e-mail.

```
[C#]

private void OnCreateInteractionComplete(IAsyncResult result)
{
    var response = interactionServerProtocol.EndRequest(result);
    if (response == null || response.Id != EventAck.MessageId)
        // for this sample, no error handling is implemented
        return;

    var @event = response as EventAck;
    mInteractionId = (string)@event.Extension["InteractionId"];
}
```

In this example we are storing the `InteractionId` value in a simple variable named `mInteractionId`, which is assumed to be defined for your project. In larger samples (or full projects), a more robust way of tracking and handling Interactions may be required.

Storing E-Mail Details in UCS

With the ID of your newly created Interaction available, it is time to store details about the e-mail you are sending in the UCS database.

There are three types of information that must be stored in the UCS database:

- Interaction Attributes - Define details about the related Interaction for this information.
- Entity Attributes - Define where the e-mail message is coming from and going to. You will use EmailOutEntityAttributes for storing outbound e-mail messages, and EmailInEntityAttributes for storing inbound e-mail messages.
- Interaction Content - Define the actual contents of the email message, including the main text and any MIME attachments.

Creating and configuring a `RequestInsertInteraction` object with this information can be easily accomplished, as shown below.

```
[C#]
```

```
public void StoreDetails(string ixnType, string ixnSubtype)
{
    var req = new RequestInsertInteraction();
    req.InteractionAttributes = new InteractionAttributes()
    {
        Id = mInteractionId,
        MediaTypeId = "email",
        TypeId = ixnType,
        SubtypeId = ixnSubtype,
        TenantId = 101,
        Status = new NullableStatuses(Statuses.Pending),
        Subject = "Sample e-mail subject",
        EntityTypeId = new NullableEntityTypes(EntityTypes.EmailOut),
    };
    req.EntityAttributes = new EmailOutEntityAttributes()
    {
        FromAddress = "sending@email.com",
        ToAddresses = "receiving@email.com",
        CcAddresses = "copied@email.com",
        ...
    };
    req.InteractionContent = new InteractionContent()
    {
        Text = "This is the e-mail body.",
       ...
    };
    contactServerProtocol.Send(req);
}
```

A list of `InteractionAttributes` properties that need to be set for an email message is provided in the following table. The properties shown for `EmailOutEntityAttributes` and `InteractionContent` represent some of those most commonly used with email. Please check the documentation provided for each class to see a full list of available properties.

Interaction Attribute Name	Description
EntityTypeId	Indicates whether this is an outgoing or incoming e-mail.
Id	Interaction ID of the related Interaction record, created earlier.
MediaTypeId	Primary media type of the Interaction you are submitting to Interaction Server. Intended for Media Server.
Subject	Subject line for this e-mail message.
SubtypeId	Interaction subtype for this e-mail message. Must match an Interaction Subtype Business Attribute, as specified in Configuration Server.
Status	Current status of the e-mail message.
TenantId	ID of the Tenant where this e-mail belongs.
TypeId	Interaction type for this e-mail message. Must match an Interaction Type Business Attribute, as specified in Configuration Server.

Placing the Interaction in the Appropriate Queue

When an Interaction has been created to handle the e-mail, and all content has been stored in the UCS database, you are free to begin processing the message as you would process any normal Interaction. This is accomplished by moving the Interaction that you created into the appropriate queue for e-mail processing, as defined in Interaction Routing Designer.

```
[C#]

public void PlaceInQueue(string queue)
{
    var req = RequestPlaceInQueue.Create();
    req.InteractionId = mInteractionId;
    req.Queue = queue;

    interactionServerProtocol.Send(req);
}
```

Replying to an E-Mail Message

Replying to an existing e-mail message follows the same basic process outlined above, but requires a few additional parameters to be set in your requests. These changes are described in the following subsections.

Changes to Creating an Interaction

When creating the Interaction, you need to specify one additional parameter before submitting your `RequestSubmit`. Take the `InteractionId` of the Interaction that represents the original e-mail message, and use that value as the `ParentInteractionId` parameter in your request, as shown below:

```
[C#]

var req = RequestSubmit.Create();
...
req.ParentInteractionId = parentInteractionId;
```

The following table describes these additional attributes.

Attribute Name	Description
ParentInteractionId	InteractionId of a parent e-mail Interaction. Only set this value when replying to an existing e-mail message.

Changes to Storing E-Mail Details in UCS

When storing e-mail details in UCS, you need to specify values for three additional interaction attributes before sending your `RequestInsertInteraction`. These attributes (shown in the code snippet below) provide a link between the parent entry in UCS and any related children, as well as specifying a common thread ID.

```
[C#]

var req = new RequestInsertInteraction();
...
req.InteractionAttributes.ParentId = parentInteractionId;
req.InteractionAttributes.CanBeParent = False;
req.InteractionAttributes.ThreadId = parentThreadId;
```

The following table describes these additional attributes.

Attribute Name	Description
CanBeParent	Boolean value that indicates whether this message can be a parent.
ParentId	Interaction ID for the parent e-mail Interaction.
ThreadId	Unique value that is shared between all UCS entries in an e-mail conversation.

Other Considerations

Although this introduction to creating and handling e-mail messages is not intended to be a comprehensive guide, it is useful to quickly introduce some other considerations and basic concepts regarding how requests are submitted and how errors should be handled.

The first consideration to take into account is how you submit requests. In the code provided here, a simple Send method is used to submit most requests without waiting for a response from the server. However, for more complicated samples or implementations you should consider using the `BeginRequest` method with a callback handler instead.

Using `BeginRequest` allows requests to be submitted without waiting for a response, but provides the ability to confirm the result and response of each request. This allows better error handling to be implemented if a request fails. *Creating an Interaction* uses the `BeginRequest` method and a callback handler to capture the InteractionID that is returned.

A second consideration to be aware of is how records in Interaction Server and UCS are related when implementing error handling. If you have already created a new Interaction and then the `RequestInsertInteraction` fails, you need to either resubmit the UCS record or delete the related Interaction by submitting a `RequestStopProcessing`. (If you reversed those steps and created a UCS record first, then the same idea must be applied if the request to create a new Interaction fails.)

Chat

You can use the Web Media Platform SDK to write Java or .NET applications that use the Genesys Web Media Server's chat, e-mail and voice callback protocols. These applications can range from the simple to the advanced.

This article shows how to implement the basic functions you will need to write a simple Web Media Server application. It provides code snippets to illustrate how the `FlexChat` protocol can be used to support a simple chat application.

Java

Importing the Web Media Protocols

Before using the Web Media Platform SDK, you will need to import the appropriate packages. Since we will be using the FlexChat protocol, we will use the following import statements:

```
[Java]
import com.genesyslab.platform.webmedia.protocol.*;
import com.genesyslab.platform.webmedia.protocol.flexchat.*;
import com.genesyslab.platform.webmedia.protocol.flexchat.events.*;
import com.genesyslab.platform.webmedia.protocol.flexchat.requests.*;
```

Setting Up Web Media Protocol Objects

When interacting with existing chat sessions, you will need to store session-specific details including a secure key and user ID. Additional objects that will be needed include a `FlexChatProtocol` object (for sending and receive messages) and a `FlexTranscript` object (used to store and interact with the chat transcript).

```
[Java]
private String mSecureKey = null;
private String mUserId = null;
private FlexTranscript mTranscript = null;
private FlexChatProtocol mFlexChatProtocol = null;
```

To use the Web Media Platform SDK, you first need to instantiate a protocol object by supplying information about the Web Media Server you want to connect with. This example specifies values for the name, host, and port values:

```
[Java]
mFlexChatProtocol = new FlexChatProtocol(new Endpoint("FlexChat", "<hostname>", <port>));
Thread mListenerThread = new ListenForEventsThread(mFlexChatProtocol);
```

Note that you have to provide a string when you create the `FlexChatProtocol` object. This string should be unique for each protocol used in your application. It might be a good idea to use the name of the server's application object from the configuration layer, which guarantees uniqueness as well as clearly identifying which server you are communicating with.

After instantiating the `FlexChatProtocol` object, you need to open the connection to the Web Media Server:

```
[Java]
mFlexChatProtocol.open();
```

Note that you should always use proper error handling techniques in your code, especially when working with the protocol connection. To save space, these error handling steps are not shown in this example.

Logging in to Chat Server

```
[Java]
// filter the request based on our configured application name
KeyValueCollection kvUserData = new KeyValueCollection();
kvUserData.addObject("FirstName", "John");
kvUserData.addObject("LastName", "Smith");
kvUserData.addObject("EmailAddress", "john.smith@email.com");
RequestLogin reqLogin = RequestLogin.create(strNickName, 0, kvUserData);
Message msg = mFlexChatProtocol.request(reqLogin);
```

After successfully logging in to Chat Server, a message is returned that includes some important information: the Secure Key and User ID values. You will use these values when sending messages to the Chat Server, so remember to keep track of them for later.

```
[Java]
if (msg != null && msg.messageId() == EventStatus.ID)
{
        EventStatus status = (EventStatus)msg;
        if (status.getRequestResult() == RequestResult.Success)
        {
                mSecureKey = status.getSecureKey();
                mUserId = status.getUserId();
        }
}
```

Updating a Chat Session

By creating a `RequestRefresh` object, you can either check for updates or send new text to an existing chat session. The following sample shows how to create a `RequestRefresh` object, send it to the Chat Server, and process the result.

```
[Java]
RequestRefresh reqRefresh = RequestRefresh.create(mUserId, mSecureKey,
        mTranscript.getLastPosition() + 1, MessageText.create("text", message));
Message msg = mFlexChatProtocol.request(reqRefresh);
if (msg != null && msg.messageId() == EventStatus.ID)
{
        EventStatus status = (EventStatus)msg;
```

```
        if (status.getRequestResult() == RequestResult.Success)
        {
                processTranscript(status.getFlexTranscript());
        }
}
```

Working with Restricted Characters

Due to server-side requirements, the XML-based Webmedia Platform SDK protocols (BasicChat, FlexChat, Callback and Email) do not support illegal characters in string values. See http://www.w3.org/TR/2000/REC-xml-20001006#NT-Char for the allowable character range.

The Platform SDK protocols do not change user data by default, but the following options are available if you want to replace illegal characters:

(1) Include code in your application to configure the protocol connection. For example:

```
[Java]
EmailProtocol protocol = new EmailProtocol(new Endpoint("emailServer", HOST, PORT));
PropertyConfiguration conf = new PropertyConfiguration();
conf.setBoolean(WebmediaChannel.OPTION_NAME_REPLACE_ILLEGAL_UNICODE_CHARS, true);
// "replacement" value is optional: if it is not specified - illegal characters will be
removed
conf.setOption(WebmediaChannel.OPTION_NAME_ILLEGAL_UNICODE_CHARS_REPLACEMENT, "?");
protocol.configure(conf);
protocol.open();
```

(2) Set specific JVM properties for the client application using `webmediaprotocol.jar`. For example:

```
[Java]
"-Dcom.genesyslab.platform.WebMedia.Email.replace-illegal-unicode-chars=true"
```

or

```
[Java]
"-Dcom.genesyslab.platform.WebMedia.Email.replace-illegal-unicode-chars=true
-Dcom.genesyslab.platform.WebMedia.Email.illegal-unicode-chars-replacement=?"
```

Using JVM system properties will affect all protocol connections for the specified Webmedia protocol. Using specific connection configuration values will only affect the specified protocol instance(s), and will take priority over JVM settings.

If no replacement character or string is specified, then illegal characters will be removed (that is, replaced with an empty string).

Values are extracted independently for the two methods listed above. If you enable character replacement using the `PropertyConfiguration` class without specifying a replacement value, but a replacement value is already specified through the JVM system properties, then characters will be replaced without verifying the enabling option in the JVM properties. It is recommended to use both options while writing connection configuration code.

Logging out of a Chat Session

When a client is ready to log out from the existing chat session, build a `RequestLogout` object and send it to the Chat Server.

```
[Java]
RequestLogout reqLogout = RequestLogout.create(mUserId, mSecureKey,
        mTranscript.getLastPosition());
Message msg = mFlexChatProtocol.request(reqLogout);
if (msg != null && msg.messageId() == EventStatus.ID)
{
        EventStatus status = (EventStatus)msg;
        if (status.getRequestResult() == RequestResult.Success)
        {
                processTranscript(status.getFlexTranscript());
        }
}
```

Disconnecting from a Chat Server

Finally, when you are finished communicating with the Chat Server, you should close the connection to minimize resource utilization.

```
[Java]
mFlexChatProtocol.close();
```

.NET

Using the Web Media Protocols

Before using the Web Media Platform SDK, you should include `using` statements that allow access to types from the Platform SDK Commons and Web Media namespaces. For the `FlexChat` protocol, we use the following statements:

```
[C#]
using Genesyslab.Platform.Commons.Collections;
using Genesyslab.Platform.Commons.Connection;
using Genesyslab.Platform.Commons.Protocols;

using Genesyslab.Platform.WebMedia.Protocols;
using Genesyslab.Platform.WebMedia.Protocols.FlexChat;
using Genesyslab.Platform.WebMedia.Protocols.FlexChat.Events;
using Genesyslab.Platform.WebMedia.Protocols.FlexChat.Requests;
```

Setting Up Web Media Protocol Objects

When interacting with existing chat sessions, you will need to store session-specific details including a secure key and user ID. Additional objects that will be needed include a `FlexChatProtocol` object (for sending and receive messages) and a `FlexTranscript` object (used to store and interact with the chat transcript).

```
[C#]
private string secureKey;
private string userId;
private FlexTranscript flexTranscript;
private FlexChatProtocol flexChatProtocol;
```

To use the Web Media Platform SDK, you first need to instantiate a `Protocol` object by supplying information about the Web Media Server you want to connect with. This example specifies values for the name, host, and port values:

```
[C#]
flexChatProtocol = new FlexChatProtocol(new Endpoint("Flex_Chat_Server", "<hostname>",
<port>));
```

Note that you have to provide a string when you create the `FlexChatProtocol` object. This string should be unique for each protocol used in your application. It might be a good idea to use the name of the server's application object from the configuration layer, which guarantees uniqueness as well as clearly identifying which server you are communicating with.

After instantiating the `FlexChatProtocol` object, you need to open the connection to the Web Media Server:

```
[C#]
flexChatProtocol.Open();
```

You should always use proper error handling techniques in your code, especially when working with the protocol connection. To save space, these error handling steps are not shown in this example.

Logging in to Chat Server

```
[C#]
// filter the request based on our configured application name
KeyValueCollection kvUserData = new KeyValueCollection();
kvUserData.Add("FirstName", "John");
kvUserData.Add("LastName", "Smith");
kvUserData.Add("EmailAddress", "john.smith@email.com");
RequestLogin reqLogin = RequestLogin.Create("reqLogin", 0, kvUserData);
EventStatus msg = this.flexChatProtocol.Request(reqLogin) as EventStatus;
```

After successfully logging in to Chat Server, a message is returned that includes some important information: the Secure Key and User ID values. You will use these values when sending messages to the Chat Server, so remember to keep track of them for later.

```
[C#]
if (msg != null && msg.Id == EventStatus.MessageId)
{
    if (msg.RequestResult == RequestResult.Success)
```

```
    {
        secureKey = msg.SecureKey;
        userId = msg.UserId;
    }
}
```

Updating a Chat Session

By creating a `RequestRefresh` object, you can either check for updates or send new text to an existing chat session. The following sample shows how to create a `RequestRefresh` object, send it to the Flex Chat protocol, and process the result.

```
[C#]
RequestRefresh reqRefresh = RequestRefresh.Create(
        userId, secureKey, flexTranscript.LastPosition + 1, MessageText.Create(""));
EventStatus msg = this.flexChatProtocol.Request(reqJoin) as EventStatus;
if (msg != null && msg.Id == EventStatus.MessageId)
{
        if (msg.RequestResult == RequestResult.Success)
        {
                ProcessTranscript(msg.FlexTranscript);
        }
}
```

Working with Restricted Characters

Due to server-side requirements, the XML-based Web Media Platform SDK protocols (`BasicChat`, `FlexChat`, `Callback` and `Email`) do not support illegal characters in string values. See http://www.w3.org/TR/2000/REC-xml-20001006#NT-Char for the allowable character range.

The Platform SDK protocols do not change user data by default, but if you want to replace illegal characters then you can include code in your application to configure the protocol connection. For example:

```
[C#]
EmailProtocol protocol = new EmailProtocol(new Endpoint("emailServer", HOST, PORT));
// Note: to use the PropertyConfiguration class, ensure that your using
// statements include Genesyslab.Platform.Commons.Connection
PropertyConfiguration conf = new PropertyConfiguration();
conf.SetBoolean(WebmediaChannel.OptionNameReplaceIllegalUnicodeChars, true);
// "replacement" value is optional: if it is not specified - illegal characters will be
removed
conf.SetOption(WebmediaChannel.OptionNameIllegalUnicodeCharsReplacement, "?");
protocol.Configure(conf);
protocol.Open();
```

Using specific connection configuration values in this manner will only affect the specified protocol instance(s).

If no replacement character or string is specified, then illegal characters will be removed (that is, replaced with an empty string).

Logging out of a Chat Session

When a client is ready to log out from the existing chat session, build a `RequestLogout` object and send it to your Flex Chat protocol.

```
[C#]
RequestLogout reqLogout = RequestLogout.Create(userId, secureKey,
flexTranscript.LastPosition);
IMessage msg = flexChatProtocol.Request(reqLogout);
if (msg != null && msg.Id == EventStatus.MessageId)
{
    if ((msg as EventStatus).RequestResult == RequestResult.Success)
    {
                  ProcessTranscript((msg as EventStatus).FlexTranscript);
    }
}
```

Disconnecting from a Chat Server

Finally, when you are finished communicating with the Chat Server, you should close the connection to minimize resource utilization.

```
[C#]
flexChatProtocol.Close();
```

Outbound

Java

You can use the Outbound Contact Platform SDK to write Java or .NET applications that work with the Genesys Outbound Contact Server. These applications can range from the simple to the advanced. This document shows how to implement the basic functions you will need to write a simple Outbound Contact application. It is organized to show the kind of structure you will probably use to write your own applications.

Setting Up an OutboundServerProtocol Object

The first thing you need to do to use the Outbound Contact Platform SDK is instantiate a `OutboundServerProtocol` object. To do that, you must supply information about the Outbound Contact Server you want to connect with. This example uses the URI of the server, but you can also use name, host, and port information:

```
[Java]

OutboundServerProtocol outboundServerProtocol =
        new OutboundServerProtocol(
                new Endpoint(
                        outboundServerUri));
```

After instantiating the protocol object, you need to open the connection to the server:

```
[Java]

outboundServerProtocol.open();
```

Closing the Connection

Finally, when you are finished communicating with the Outbound Contact Server, you should close the connection to minimize resource utilization:

```
[Java]

outboundServerProtocol.close();
```

.NET

You can use the Outbound Contact Platform SDK to write Java or .NET applications that work with the

Genesys Outbound Contact Server. These applications can range from the simple to the advanced. This document shows how to implement the basic functions you will need to write a simple Outbound Contact application. It is organized to show the kind of structure you will probably use to write your own applications.

Setting Up an OutboundServerProtocol Object

The first thing you need to do to use the Outbound Contact Platform SDK is instantiate a `OutboundServerProtocol` object. To do that, you must supply information about the Outbound Contact Server you want to connect with. This example uses the URI of the server, but you can also use name, host, and port information:

```
[C#]

OutboundServerProtocol outboundServerProtocol =
        new OutboundServerProtocol(
                new Endpoint(
                        outboundServerUri));
```

After instantiating the `OutboundServerProtocol` object, you need to open the connection to the Outbound Contact Server:

```
[C#]

outboundServerProtocol.Open();
```

Closing the Connection

Finally, when you are finished communicating with the Outbound Contact Server, you should close the connection to minimize resource utilization:

```
[C#]

outboundServerProtocol.Close();
```

Management Layer

You can use the Management Platform SDK to write Java or .NET applications that interact with the Genesys Message Server, Solution Control Server and Local Control Agents (LCAs). Most people will want to use this SDK to make their applications visible to the Genesys Management Layer so they can monitor them with Solution Control Server.

This document shows how to implement the basic functions you will need to write a simple voice application. It is organized to show the kind of structure you will probably use to write your own applications.

> **Important**
> This article was originally written before improvements to the protocol layer and message handling cause the Message Broker and Protocol Manager Application Blocks to be deprecated. Although still included with this release for backwards compatibility, Genesys no longer recommends using those Application Blocks in your applications. The concepts discussed in this article still apply, but please refer to the updated Connecting to a Server and Event Handling articles for an understanding of how messages and protocols should be managed.

Java

Making Your Application Visible to the Genesys Management Layer

A Genesys Local Control Agent (LCA) runs on each host in the Genesys environment, enabling the Management Layer to monitor and control the applications running on that host. This section shows how to use the LCA running on your own host to make your application visible to the Genesys Management Layer.

Connecting to the Local Control Agent

As mentioned previously, the Platform SDK uses a message-based architecture to connect to Genesys servers. Genesys recommends that you use the Protocol Manager Application Block to handle these connections. We also recommend that you use the Message Broker Application Block for your message handling. Here is how to use Protocol Manager to connect to the LCA. (This discussion is based on the material in the article on Connecting to a Server Using the Protocol Manager Application Block.)

Here are the import statements you need for Protocol Manager:

```
[Java]

import com.genesyslab.platform.applicationblocks.commons.protocols.LcaConfiguration;
import
com.genesyslab.platform.applicationblocks.commons.protocols.ProtocolManagementServiceImpl;
```

After you have set up your import statements, the first thing you need to do is create a Protocol Management Service object:

```
[Java]

protocolManagementServiceImpl = new ProtocolManagementServiceImpl();
```

Then you need to create an LCA Configuration object:

```
[Java]

lcaConfiguration = new LcaConfiguration("Lca_App");
```

Note that the name you supply with the Configuration object ("Lca_App" in this example) is arbitrary.

Now you can configure the connection. This sample uses the default LCA port of 4999. It also sets the execution mode to Backup:

```
[Java]

try {
        lcaConfiguration.setUri(new URI("tcp://localhost:4999"));
} catch (URISyntaxException e) {
        e.printStackTrace();
}
lcaConfiguration.setClientName("Generic_Server_Backup");
lcaConfiguration.setExecutionMode(ApplicationExecutionMode.Backup);
```

Once you have configured your connection, you can register your `LcaConfiguration` object with the Protocol Management Service, and open your connection to the LCA:

```
[Java]

protocolManagementServiceImpl.register(lcaConfiguration);
try {
        protocolManagementServiceImpl.getProtocol("Lca_App").open();
} catch (InterruptedException e) {
        e.printStackTrace();
} catch (ProtocolException e) {
        e.printStackTrace();
}
```

Updating the Application Status

When you need to update the status of your application, send a `RequestUpdateStatus`. Here is how to indicate that the application is running:

```
[Java]

RequestUpdateStatus requestUpdateStatus = RequestUpdateStatus.create();
requestUpdateStatus
        .setApplicationName(lcaConfiguration.getClientName());
requestUpdateStatus.setControlStatus(ApplicationStatus.Running.asInteger());
```

```
try {
        protocolManagementServiceImpl.getProtocol("Lca_App").send(
                requestUpdateStatus);
} catch (ProtocolException e) {
        e.printStackTrace();
}
```

The LCA will not return an event when you change the application status. So for this particular task, you will not need any more code.

Execution Mode and Event Handling

As mentioned, the LCA will not return an event when you change the application status. But when you change the execution mode — for example, from `Primary` to `Backup` — you will receive an `EventChangeExecutionMode`. Unlike most events you receive in the Platform SDK, this event *requires* a response from your application. If the Management Layer does not know that your application is expecting to work in `Primary` mode, for example, it cannot rely on the stability of the Genesys environment.

> **Important**
> If you do not respond within the configured timeout period, your application will be terminated by the Management Layer.

After receiving the `EventChangeExecutionMode`, your application must send a `ResponseExecutionModeChanged` to indicate to the Management Layer that you are now ready to run in the new execution mode.

In order to handle these events, you need to set up the Message Broker Application Block, as outlined here. (Note that this information is based on the Event Handling Using the Message Broker Application Block article, which discusses this process in more detail.)

First, here are the import statements:

```
[Java]

import com.genesyslab.platform.applicationblocks.commons.Action;
import com.genesyslab.platform.applicationblocks.commons.broker.BrokerServiceFactory;
import com.genesyslab.platform.applicationblocks.commons.broker.EventBrokerService;
import com.genesyslab.platform.applicationblocks.commons.broker.MessageIdFilter;
```

Now you can create an Event Broker Service object and register your event handler:

```
[Java]

eventBrokerService = BrokerServiceFactory
        .CreateEventBroker(protocolManagementServiceImpl.getReceiver());
eventBrokerService.register(new ChangeExecutionModeHandler(),
        new MessageIdFilter(EventChangeExecutionMode.ID));
```

Here is a sample of the handler you might set up for the `EventChangeExecutionMode`. This handler includes your `ResponseExecutionModeChanged`:

```
[Java]

class ChangeExecutionModeHandler implements Action {

        public void handle(Message obj) {
                EventChangeExecutionMode eventChangeExecutionMode =
(EventChangeExecutionMode) obj;
                if (eventChangeExecutionMode != null) {
                        System.out.println("eventChangeExecutionMode:\n"
                                + eventChangeExecutionMode + "\n\n");
                        ApplicationExecutionMode mode = eventChangeExecutionMode
                                .getExecutionMode();
                        lcaConfiguration.setExecutionMode(mode);
                        ResponseExecutionModeChanged response = ResponseExecutionModeChanged
                                .create(mode);
                        System.out.println("Sending response: " + response);
                        try {
                                protocolManagementServiceImpl.getProtocol("Lca_App").send(
                                        response);
                        } catch (ProtocolException e) {
                                e.printStackTrace();
                        }
                }
        }
}
```

This is how you might send a request to change your application's execution mode from Backup to `Primary`:

```
[Java]

RequestUpdateStatus requestUpdateStatus = RequestUpdateStatus.create();
requestUpdateStatus
        .setApplicationName(lcaConfiguration.getClientName());
requestUpdateStatus.setExecutionMode(ApplicationExecutionMode.Primary);
try {
        protocolManagementServiceImpl.getProtocol("Lca_App").send(
                requestUpdateStatus);
} catch (ProtocolException e) {
        e.printStackTrace();
}
```

Once you have sent this request, the LCA will return an `EventChangeExecutionMode`, to which your application will respond with the `ResponseExecutionModeChanged` shown above.

Closing the Connection

When you are finished, close the connection to the LCA:

```
[Java]

try {
        protocolManagementServiceImpl.getProtocol("Lca_App").close();
} catch (ProtocolException e) {
        e.printStackTrace();
}
```

Monitoring Your Application with Solution Control Server

Solution Control Server can be used to monitor applications running in the Genesys environment. Here is how to obtain information about hosts and applications.

Connecting to Solution Control Server

As with the LCA example above, you need to set up a Protocol Management Service:

```
[Java]

protocolManagementServiceImpl = new ProtocolManagementServiceImpl();
```

Now you can create an `ScsConfiguration` object and supply the necessary parameters. The `ClientName` is the name of a Solution Control application that has been set up in the Configuration Layer, while the `ClientId` is the `DBID` of that application:

```
[Java]

scsConfiguration = new ScsConfiguration("Scs_App");
try {
        scsConfiguration.setUri(new URI(uriString));
} catch (URISyntaxException e) {
        e.printStackTrace();
}
scsConfiguration.setClientName(applicationName);
scsConfiguration.setClientId(applicationDbid);
scsConfiguration.setUserName(userName);
```

At this point, register your Configuration object with the Protocol Management Service and open the connection to Solution Control Server:

```
[Java]

protocolManagementServiceImpl.register(scsConfiguration);
try {
        protocolManagementServiceImpl.getProtocol("Scs_App").open();
} catch (InterruptedException e) {
        e.printStackTrace();
} catch (ProtocolException e) {
        e.printStackTrace();
}
```

Setting Up Event Handling

You will need to set up some event handling code, since Solution Control Server will return `EventInfo` or `EventError` messages in response to your requests for information. The code for this is similar to the LCA-related code shown above:

```
[Java]

eventBrokerService = BrokerServiceFactory
        .CreateEventBroker(protocolManagementServiceImpl.getReceiver());
.
```

```
.
.
eventBrokerService.register(new ScsEvent(),
        new MessageIdFilter(EventInfo.ID));
.
.
.
class ScsEvent implements Action {

        public void handle(Message obj) {
                EventInfo eventInfo =
                        (EventInfo) obj;
                if (eventInfo != null) {
                        // Handle this event
                }
        }
}
```

Requesting Application Information

Here is how to request the status of an application, using its `DBID`:

```
[Java]

RequestGetApplicationInfo requestGetApplicationInfo =
        RequestGetApplicationInfo.create(applicationDbid);
try {
        protocolManagementServiceImpl.getProtocol("Scs_App").send(requestGetApplicationInfo);
} catch (ProtocolException e) {
        e.printStackTrace();
}
```

When you send this request, you will receive an `EventInfo` that includes the status of the application:

```
'EventInfo' (1) attributes:
        attr_cfg_obj_type [int] = 9 [Application]
        attr_obj_live_status [int] = 1 [Stopped]
        attr_client [int] = 1384
        attr_message [str] = "APP_STATUS_STOPPED"
        attr_cfg_obj_id [int] = 174
        attr_live_status_second [int] = [output suppressed]
        attr_ref_id [int] = 2
        attr_app_work_mode [int] = 2 [Exiting]
```

If you want to be notified when the status of an application changes, send a `RequestSubscribe`.

```
[Java]

RequestSubscribe requestSubscribeApp = RequestSubscribe.create();
requestSubscribeApp.setControlObjectType(ControlObjectType.Application);
requestSubscribeApp.setControlObjectId(applicationDbid);
try {
        protocolManagementServiceImpl.getProtocol("Scs_App").send(requestSubscribeApp);
} catch (ProtocolException e) {
        e.printStackTrace();
}
```

Whenever the application's status changes, you will receive an `EventInfo` that informs you of the new status.

Requesting Host Information

You can also request information about the status of a host. But in this case, you must issue a `RequestSubscribe` before you will receive any information about the host. Here is how:

```
[Java]

RequestSubscribe requestSubscribe = RequestSubscribe.create();
requestSubscribe.setControlObjectType(ControlObjectType.Host);
requestSubscribe.setControlObjectId(hostDbid);
try {
        protocolManagementServiceImpl.getProtocol("Scs_App").send(requestSubscribe);
} catch (ProtocolException e) {
        e.printStackTrace();
}

RequestGetHostInfo requestGetHostInfo =
        RequestGetHostInfo.create(hostDbid);
try {
        protocolManagementServiceImpl.getProtocol("Scs_App").send(requestGetHostInfo);
} catch (ProtocolException e) {
        e.printStackTrace();
}
```

If you just send the `RequestSubscribe`, you will be notified any time the host status changes. If you also send the `RequestGetHostInfo`, you will also receive an immediate notification of the host's status, whether it has changed or not. Here is a sample of the information you will receive.

```
'EventInfo' (1) attributes:
        attr_cfg_obj_type [int] = 10 [Host]
        attr_obj_live_status [int] = 2 [StopTransition]
        attr_client [int] = 1920
        attr_message [str] = "HOST_STATUS_RUNNING"
        attr_cfg_obj_id [int] = 114
        attr_ref_id [int] = 3
```

Once you have subscribed to a host, you can send a `RequestGetHostInfo` at any time to receive information about its status.

Closing the Connection

When you are finished, close the connection to Solution Control Server:

```
[Java]

try {
        protocolManagementServiceImpl.getProtocol("Scs_App").close();
} catch (ProtocolException e) {
        e.printStackTrace();
} catch (IllegalStateException e) {
        e.printStackTrace();
} catch (InterruptedException e) {
        e.printStackTrace();
}
```

Sending a Log Message to Message Server

You can easily send log messages to Message Server using the Management Platform SDK. The following discussion will use a `MessageServerProtocol` object, rather than using the Protocol Manager Application Block.

First you need to create the Protocol object:

```
[Java]

MessageServerProtocol messageServerProtocol =
        new MessageServerProtocol(
                new Endpoint(new URI(serverUri)));
```

Now you can configure the Protocol object and open the connection to Message Server:

```
[Java]

messageServerProtocol.setClientType
        (ConfServerClientType.ThirdPartyApp.ordinal());
messageServerProtocol.setClientName("Third_Party_App");
messageServerProtocol.setClientHost("hostname");

messageServerProtocol.open();
```

At this point, you are ready to create a `RequestLogMessage`:

```
[Java]

RequestLogMessage request = RequestLogMessage.create();
request.setEntryId(9600);
request.setEntryCategory(LogCategory.Alarm);
request.setEntryText("Message Text...");
request.setLevel(LogLevel.Alarm);
request.setTime(new Date());
```

Once you have created the request, you can send the request to Message Server. When you are finished, you should close the connection:

```
[Java]

messageServerProtocol.send(request);
.
.
.
messageServerProtocol.close();
```

.NET

Making Your Application Visible to the Genesys Management Layer

A Genesys Local Control Agent (LCA) runs on each host in the Genesys environment, enabling the Management Layer to monitor and control the applications running on that host. This section shows how to use the LCA running on your own host to make your application visible to the Genesys Management Layer.

Connecting to the Local Control Agent

As mentioned previously, the Platform SDK uses a message-based architecture to connect to Genesys servers. Genesys recommends that you use the Protocol Manager Application Block to handle these connections. We also recommend that you use the Message Broker Application Block for your message handling. Here is how to use Protocol Manager to connect to the LCA. (This discussion is based on the material in the article on Connecting to a Server Using the Protocol Manager Application Block.)

Here is the `using` statement you need for Protocol Manager:

[C#]

```
using Genesyslab.Platform.ApplicationBlocks.Commons.Protocols;
```

After you have set up your using statement, the first thing you need to do is create a Protocol Management Service object:

[C#]

```
protocolManagementService = new ProtocolManagementService();
```

Then you need to create an LCA Configuration object:

[C#]

```
LcaConfiguration lcaConfiguration = new LcaConfiguration("Lca_App");
```

Note that the name you supply with the Configuration object ("Lca_App" in this example) is arbitrary.

Now you can configure the connection. This sample uses the default LCA port of 4999. It also sets the application status to `Initializing` and the execution mode to Backup:

[C#]

```
lcaConfiguration.Uri = new Uri("tcp://localhost:4999");
lcaConfiguration.ClientName = applicationName;
lcaConfiguration.ControlStatus = (int) ApplicationStatus.Initializing;
lcaConfiguration.ExecutionMode = ApplicationExecutionMode.Backup;
```

Once you have configured your connection, you can register your `LcaConfiguration` object with the Protocol Management Service, and open your connection to the LCA:

[C#]

```
protocolManagementService.Register(lcaConfiguration);
```

```
protocolManagementService["Lca_App"].Open();
```

Updating the Application Status

When you need to update the status of your application, send a `RequestUpdateStatus`. Here is how to indicate that the application is running:

```
[C#]

RequestUpdateStatus requestUpdateStatus = RequestUpdateStatus.Create();
requestUpdateStatus.ApplicationName = lcaConfiguration.ClientName;
requestUpdateStatus.ControlStatus = (int) ApplicationStatus.Running;
protocolManagementService["Lca_App"].Send(requestUpdateStatus);
```

The LCA will not return an event when you change the application status. So for this particular task, you will not need any more code.

Execution Mode and Event Handling

As mentioned, the LCA will not return an event when you change the application status. But when you change the execution mode — for example, from `Primary` to `Backup` — you will receive an `EventChangeExecutionMode`. Unlike most events you receive in the Platform SDK, this event requires a response from your application. If the Management Layer does not know that your application is expecting to work in `Primary` mode, for example, it cannot rely on the stability of the Genesys environment.

> **Important**
> If you do not respond within the configured timeout period, your application will be terminated by the Management Layer.

After receiving the `EventChangeExecutionMode`, your application must send a `ResponseExecutionModeChanged` to indicate to the Management Layer that you are now ready to run in the new execution mode.

In order to handle these events, you need to set up the Message Broker Application Block, as outlined here. (Note that this information is based on the Event Handling Using the Message Broker Application Block article, which discusses this process in more detail.)

First, here is the using statement:

```
[C#]

using Genesyslab.Platform.ApplicationBlocks.Commons.Broker;
```

Now you can create an Event Broker Service object and register your event handler:

```
[C#]

eventBrokerService =
        BrokerServiceFactory.CreateEventBroker(protocolManagementService.Receiver);
eventBrokerService.Register(this.OnEventChangeExecutionMode,
```

```
        new MessageIdFilter(EventChangeExecutionMode.MessageId));
```

Here is a sample of the handler you might set up for the EventChangeExecutionMode. This handler includes your ResponseExecutionModeChanged:

```
[C#]

private void OnEventChangeExecutionMode(IMessage theMessage)
{
                EventChangeExecutionMode eventChangeExecutionMode = theMessage as
EventChangeExecutionMode;
                if (eventChangeExecutionMode != null)
                {
                        ApplicationExecutionMode mode =
eventChangeExecutionMode.ExecutionMode;
                        lcaConfiguration.ExecutionMode = mode;
                        ResponseExecutionModeChanged response =
ResponseExecutionModeChanged.Create(mode);
                        Console.WriteLine("Sending response: " + response);
                        protocolManagementService["Lca_App"].Send(response);
                }
        }
```

This is how you might send a request to change your application's execution mode from Backup to Primary:

```
[C#]

RequestUpdateStatus requestUpdateStatus =
        RequestUpdateStatus.Create();
requestUpdateStatus.ApplicationName = lcaConfiguration.ClientName;
requestUpdateStatus.ExecutionMode =
        ApplicationExecutionMode.Primary;
protocolManagementService["Lca_App"].Send(requestUpdateStatus);
```

Once you have sent this request, the LCA will return an EventChangeExecutionMode, to which your application will respond with the ResponseExecutionModeChanged shown above.

Closing the Connection

When you are finished, close the connection to the LCA:

```
[C#]

protocolManagementService["Lca_App"].Close();
```

Monitoring Your Application with Solution Control Server

Solution Control Server can be used to monitor applications running in the Genesys environment. Here is how to obtain information about hosts and applications.

Connecting to Solution Control Server

As with the LCA example above, you need to set up a Protocol Management Service:

[C#]

```
protocolManagementService = new ProtocolManagementService();
```

Now you can create an `ScsConfiguration` object and supply the necessary parameters. The `ClientName` is the name of a Solution Control application that has been set up in the Configuration Layer, while the `ClientId` is the DBID of that application:

[C#]

```
scsConfiguration = new ScsConfiguration("Scs_App");
scsConfiguration.Uri = new Uri(uriString);
scsConfiguration.ClientId = applicationDbid;
scsConfiguration.ClientName = applicationName;
scsConfiguration.UserName = userName;
```

At this point, register your Configuration object with the Protocol Management Service and open the connection to Solution Control Server:

[C#]

```
protocolManagementService.Register(scsConfiguration);
protocolManagementService["Scs_App"].Open();
```

Setting Up Event Handling

You will need to set up some event handling code, since Solution Control Server will return `EventInfo` or `EventError` messages in response to your requests for information. The code for this is similar to the LCA-related code shown above:

[C#]

```
eventBrokerService =
        BrokerServiceFactory.CreateEventBroker(protocolManagementService.Receiver);
.
.
.
eventBrokerService.Register(OnEventInfo);
.
.
.
private void OnEventInfo(IMessage theMessage)
{
        EventInfo eventInfo = theMessage as EventInfo;
        if (eventInfo != null)
        {
                // Handle this event
        }
}
```

Requesting Application Information

Here is how to request the status of an application, using its DBID:

[C#]

```
RequestGetApplicationInfo requestGetApplicationInfo =
RequestGetApplicationInfo.Create(applicationDbid);
```

```
protocolManagementService["Scs_App"].Send(requestGetApplicationInfo);
```

When you send this request, you will receive an `EventInfo` that includes the status of the application:

```
'EventInfo' ('1')
message attributes:
attr_app_work_mode [int] = 0 [Primary]
attr_client [int] = 660
attr_ref_id [int] = 4
attr_message [str] = "APP_STATUS_RUNNING"
attr_obj_live_status [int] = 6 [Running]
attr_cfg_obj_id [int] = 109
attr_cfg_obj_type [int] = 9 [Application]
```

If you want to be notified when the status of an application changes, send a `RequestSubscribe`.

```
[C#]

RequestSubscribe requestSubscribeApp = RequestSubscribe.Create();
requestSubscribeApp.ControlObjectType = ControlObjectType.Application;
requestSubscribeApp.ControlObjectId = applicationDbid;

protocolManagementService["Scs_App"].Send(requestSubscribeApp);
```

Whenever the application's status changes, you will receive an `EventInfo` that informs you of the new status.

Requesting Host Information

You can also request information about the status of a host. But in this case, you must issue a `RequestSubscribe` before you will receive any information about the host. Here is how:

```
[C#]

RequestSubscribe requestSubscribeHost = RequestSubscribe.Create();
requestSubscribe.ControlObjectType = ControlObjectType.Host;
requestSubscribe.ControlObjectId = hostDbid;

protocolManagementService["Scs_App"].Send(requestSubscribeHost);

RequestGetHostInfo requestGetHostInfo = RequestGetHostInfo.Create();
requestGetHostInfo.ControlObjectId = hostDbid;

protocolManagementService["Scs_App"].Send(requestGetHostInfo);
```

If you just send the `RequestSubscribe`, you will be notified any time the host status changes. If you also send the `RequestGetHostInfo`, you will also receive an immediate notification of the host's status, whether it has changed or not. Here is a sample of the information you will receive.

```
'EventInfo' ('1')
message attributes:
attr_client [int] = 660
attr_ref_id [int] = 3
attr_message [str] = "HOST_STATUS_RUNNING"
attr_obj_live_status [int] = 2 [StopTransition]
attr_cfg_obj_id [int] = 111
attr_cfg_obj_type [int] = 10 [Host]
```

Once you have subscribed to a host, you can send a `RequestGetHostInfo` at any time to receive information about its status.

Closing the Connection

When you are finished, close the connection to Solution Control Server:

```
[C#]

protocolManagementService["Scs_App"].Close();
```

Sending a Log Message to Message Server

You can easily send log messages to Message Server using the Management Platform SDK. The following discussion will use a `MessageServerProtocol` object, rather than using the Protocol Manager Application Block.

First you need to create the Protocol object:

```
[C#]

MessageServerProtocol messageServerProtocol = new MessageServerProtocol(
        new Endpoint(new Uri("tcp://host:4321")));
```

Now you can configure the Protocol object and open the connection to Message Server:

```
[C#]

messageServerProtocol.ClientType = (int) ConfServerClientType.ThirdPartyApp;
messageServerProtocol.ClientName = "Third_Party_App";
messageServerProtocol.ClientHost = "hostname";

messageServerProtocol.Open();
```

At this point, you are ready to create a `RequestLogMessage`:

```
[C#]

RequestLogMessage requestLogMessage = RequestLogMessage.Create();
requestLogMessage.EntryId = 9600;
requestLogMessage.EntryCategory = LogCategory.Alarm;
requestLogMessage.EntryText = "Message Text...";
requestLogMessage.Level = LogLevel.Alarm;
requestLogMessage.Time = new DateTime();
```

Once you have created the request, you can send the request to Message Server. When you are finished, you should close the connection:

```
[C#]

messageServerProtocol.Send(requestLogMessage);
.
.
.
messageServerProtocol.Close();
```

LCA Hang-Up Detection Support

This page provides:

- an overview and list of requirements for the LCA Hang-Up Detection Support feature
- design details explaining how this feature works
- code examples showing how to implement LCA Hang-Up Detection Support in your applications

Introduction to LCA Hang-up Detection Support

Beginning with release 8.1, the Platform SDKs now allow user-developed application to include hang-up detection functionality.

The Genesys Management Layer relies on Local Control Agent (LCA) to monitor and control applications. An open connection between LCA and Genesys applications is typically used to determine which applications are running or stopped. However, if an application that has stopped responding still has a connection to LCA then it could appear to be running correctly - preventing Management Layer from switching over to a backup application or taking other actions to restore functionality.

Hang-up detection allows Local Control Agent (LCA) to detect unresponsive Genesys applications by checking for regular heartbeat messages. When an unresponsive application is found, pre-configured actions can be taken - including triggering alarms or restarting the application.

Note: Hang-up detection functionality has been available in the Genesys Management Layer since release 8.0.1. For more information, refer to the Framework 8.0 Management Layer User's Guide. For details about related configuration options, refer to the Framework 8.0 Configuration Options Reference Manual.

Two levels of hang-up detection are available: implicit and explicit.

Implicit Hang-up Detection

The easiest form of hang-up detection to implement is implicit hang-up detection.

In this scenario, application status is monitored through the connection between your application and LCA. This functionality can be extended by adding a requirement that your application periodically interacts with LCA (either responding to ping request or sending its own heart-beat messages) as a necessary condition of application liveliness.

This simple form of hang-up detection can be implemented internally by using the `LocalControlAgentProtocol` to connect to LCA. In this case, existing applications only need to be rebuilt with a version of `LocalControlAgentProtocol` that supports hang-up detection functionality - no coding changes are required - and given the appropriate configuration options in Genesys Management Layer.

Explicit Hang-up Detection

Explicit hang-up detection offers more robust protection from applications that may become unresponsive, but is also more complex.

The periodic interaction that is monitored by implicit hang-up detection only confirms that your application can interact with LCA. In most cases this means that the application is able to communicate with other apps and that the thread responsible for communicating with LCA is still active. However, multi-threaded applications may contain other threads that are blocked or have stopped responding without interrupting communication with LCA. Explicit hang-up detection allows you to determine when only part of your application hangs-up by monitoring individual threads in the application.

In addition to allowing your application to register (or unregister) individual threads to be monitored, explicit hang-up detection also allows your application to stop or delay the monitoring process. Threads that execute synchronous functions (which can block thread execution for some extended periods) or other features that prevent accurate monitoring should take advantage of this feature.

Feature Overview

- To maintain backwards compatibility, hang-up detection must be explicitly enabled in the application configuration.
- Implicit hang-up detection can be used for applications that do not require complex monitoring functionality. No code changes are required, just rebuild your application using the new version of `LocalControlAgentProtocol`.
- Explicit hang-up detection requires minimal application participation - enabling monitoring, registering and unregistering execution threads, and providing heartbeats. Most hang-up detection functionality is implemented within the Management Layer component, while all timing information (such as maximum allowed period between heartbeats) is configured through Genesys Management Layer.

System Requirements

Genesys Management Layer:

- Release 8.0.1 or later

Platform SDK for .NET:

- Management SDK protocol release 8.1 or later
- .NET Framework 3.5
- Visual Studio 2008 (required for .NET project files)

Platform SDK for Java:

- Management SDK protocol release 8.1 or later

- J2SE 5.0 or Java 6 SE runtime

Design Details

This section provides an overview of the main classes and interfaces used to add thread monitoring functionality for Explicit hang-up detection. Before using the classes and methods described here, be sure that you have implemented basic LCA Integration in your application using `LocalControlAgentProtocol`.

Although the details of thread monitoring implementation are slightly differently for Java and .NET, the basic idea is the same: to create and update a thread monitoring table that LCA can use to confirm the status of your application.

Note that for implicit hang-up detection you are only required to rebuild your application and make adjustments to the configuration options in Genesys Management Layer; the details described below are not required for simple application monitoring.

Thread Monitoring Table

The new thread monitoring functions described below allow `LocalControlAgentProtocol` to create and maintain a thread monitoring table within the application. This table tracks basic thread status.

Sample Thread Monitoring Table

OS Thread ID	Logical Thread ID	Thread Class	Heartbeat Counter	Flags
0	«main»	1	444345	active
1	«pool_1»	2	354354	suspend
2	«pool_2»	2	432432	deleted
3	«pool_3»	2	434323	active
4	«DB_store»	3	31212	active
....	....	....	....	....

Each row corresponds to a monitored thread. Columns of the table are:

- OS Thread ID—The OS-specific thread ID, used for thread identification during monitoring. OS thread ID is not passed by application but is received directly from system.
- Logical Thread ID – Application logical thread ID (or logical name, in Java). Used for logging and thread identification.
- Thread Class—Thread class integer. This value is only meaningful within the scope of the application; threads with the same thread class value in a different application can have different roles. Examples of thread classes might be the main loop thread, pool threads, or special threads (such as external authentication threads in ConfigServer).
- Heartbeat Counter—Cumulative counter of `Heartbeat()` calls made by the corresponding thread. Incrementing this value is the main way to indicate that the thread is still alive.

NOTE: This value is initialized with a random value when the thread is registered for monitoring. This prevents incorrect hang-up detection if threads are created and terminated with high frequency, leading to repeating OS thread IDs.

- Flag—Special flags.
  - Suspended/Resumed—Corresponds to the state of thread monitoring.
  - Deleted—Used internally to notify LCA that a thread was unregistered from monitoring.

.NET Implementation

ThreadMonitoring Class

The `ThreadMonitoring` class is defined in the `Genesyslab.Diagnostics` namespace of Genesyslab.Core.dll. This class contains the following public static methods:

- `Register(int threadClass, string threadLogicId)`—enables monitoring for this thread
- `Unregister()`—removes this thread from monitoring
- `Heartbeat()`—increases heartbeat counter for this thread (indicating that thread is still alive)
- `SuspendMonitoring()`—suspend monitoring for this thread
- `ResumeMonitoring()`—resumes monitoring for this thread

Note: Each method should be called from within the thread that is being monitored.

When a thread is registered for monitoring, the following parameters are included:

- `threadClass`—Any positive integer that represents the type of thread, allowing you to specify different monitoring settings for groups of threads within an application.
- `threadLogicId`—A logical, descriptive thread ID that is independent from thread ID provided by OS. This value is used for thread identification within LCA and for logging purposes. This ID should be unique within the application.

PerformanceCounter Constants

The following String constants (names) are defined in the `ThreadMonitoring` class:

```
public const string CategoryName = "Genesyslab PSDK .NET";
public const string HeartbeatCounterName = "Thread Heartbeat";
public const string StateCounterName = "Thread State";
public const string ProcessIdCounterName = "ProcessId";
public const string OsThreadIdCounterName = "OsThreadId";
```

The Platform SDK thread monitoring functionality uses these constants to manage PerformanceCounter values. In addition to these custom performance counters, you can also use standard ones, such as those defined in `Thread` category: "% Processor Time", "% User Time", etc.

See MSDN<ref>MSDN PerformanceCounter Class (http://msdn.microsoft.com/en-us/library/system.diagnostics.performancecounter.aspx)</ref> for details about performance counters.

Note: Use of PerformanceCounters is optional, and is not required for LCA hang-up detection functionality.

Java Implementation

ThreadHeartbeatCounter class

The `ThreadHeartbeatCounter` class is defined in the `com.genesyslab.platform.commons.threading` package, located within commons.jar. This class is designed as a JMX<ref>JMX: Java Management Extensions (http://java.sun.com/javase/technologies/core/mntr-mgmt/javamanagement/)</ref> MBean and implements the public ThreadHeartbeatCounterMBean interface which is accessible through Java management framework.

There is no public constructor for the `ThreadHeartbeatCounter` class; each thread that you want to monitor should create its own instance with following static method:

```
public static ThreadHeartbeatCounter createThreadHeartbeatCounter(
            String threadLogicalName,
            int threadClass);
```

When a thread is registered for monitoring, the following parameters are included:

- `threadLogicalName`—A logical, descriptive thread name that is used to identify the thread within LCA and for logging purposes. This name should be unique within the application.
- `threadClass`—Any positive integer that represents the type of thread, allowing you to specify different monitoring settings for groups of threads within an application.

One key difference from thread monitoring using .NET is the need to create a monitoring object instance. The lifecycle of this object, including `MBeanServer` registration, is supported by the parent class `PSDKMBeanBase` and is shown in the five steps below:

1. Start monitoring a thread:

```
ThreadHeartbeatCounter monitor =
  ThreadHeartbeatCounter.createThreadHeartbeatCounter(
                threadId, threadClass);
monitor.initialize();
```

2. Notify LCA that thread is still alive (increase heartbeat counter):

```
monitor.alive();
```

3. Suspend monitoring of this thread:

```
monitor.setActive(false);
```

4. Resume monitoring of this thread:

```
monitor.setActive(true);
```

5. Finish monitoring and unregister this thread:

```
monitor.unregister();
```

Note: Each of these methods must be called from within the thread that is being monitored.

Once a `ThreadHeartbeatCounter` object is unregistered, that instance cannot be reused. To begin monitoring that thread again (or any other) you first need to create a new instance of the thread monitoring object.

ThreadHeartbeatCounterMBean interface

The `ThreadHeartbeatCounterMBean` interface is intended to present an open API to the JMX MBean. This interface contains the following publicly accessible methods:

```
public long getThreadSystemId();
public String getLogicalName();
public int getThreadClass();
public void setThreadClass(int newThreadClass);
public int getHeartbeatCounter();
public void setActive(boolean isActive);
public boolean isActive();
```

These methods are "MBean client-side" methods and are used by LCA protocol to get actual information about the thread for the monitoring table. They also allow users to change the thread class and suspend or resume thread monitoring (using `setActive(false/true)`) of a particular thread at application runtime.

References

<references />

Handle Application "Graceful Stop" with the LCA Protocol

Graceful stop is operation that allow application to complete current request handling before actual stop in case when such handling can require significant (time greater than application stop timeout) time.

Two new states are related to this command:

- SUSPENDING – This state means that an application has understood the command from LCA, so that the application does NOT accept new requests (unless specified by Management Layer) and will complete current requests. This status should be reported by an application as the result of Suspend command from LCA. If the application does no support graceful stop then it can just ignore the Suspend command; no state changes should be reported in this case.
- SUSPENDED – This state means that an application has completed handling current requests and can be stopped without any impact to its client.

For applications which DO support graceful stop, the scenario is as follows:

1. SCI issues command "Stop application graceful"
2. SCS receives command, sets up a suspended state timer and sends the Suspend application command to LCA
3. LCA receives Suspend command and resends it to application (application receives `EventSuspendApplication`)
4. Since the application supports this feature, it reports SUSPENDING state with `RequestUpdateStatus` and start behave accordingly
5. SCS receives a status update through LCA and cancels the timer set at point (2).
6. From this point, the application has unlimited time to complete handling requests (of course, it can also be stopped by usual stop command)
7. When the application completes handling requests, it reports the status SUSPENDED (with `RequestUpdateStatus`)
8. SCS receive the status update through LCA and stops the application with the usual stop command (`EventChangeExecutionMode`)

For applications which do NOT support graceful stop, the scenario is as follows:

1. SCI issues the command: "Stop application graceful"
2. SCS receives the command, sets up a suspended state timer and sends the Suspend application command to LCA
3. LCA receives Suspend command and resends it to application
4. Since the application does not support this feature, the Suspend command is ignored
5. Suspended state timer set at point (2) expires in SCS. SCS determines that the application does not support Graceful stop.

6. SCS issues an ordinary stop command to application (`EventChangeExecutionMode`).

Please also note that message between SCI-SCS and SCS-LCA are not same.

If your application should support graceful stop please check:

- If application receives Suspend command from LCA.
- If application correctly report SUSPENDING/SUSPENDED states.
- If application can be stopped by usual stop command.

If your application should NOT support graceful stop please check:

- If application can be stopped by usual stop command (in this case Graceful stop is equal to usual Stop command with some delay)

> **Tip**
>
> This feature is new for 8.0 Genesys Management Framework, so, all the involved components (SCI, SCS, LCA) should be 8.0+ versions (checked with 8.0.3x).

Routing Server

Many types of interactions enter a modern contact center, and each of them can have many possible destinations. Universal Routing Server (URS) helps them get to the right place at the right time by enabling you to create customized *routing strategies* — sets of instructions that tell URS how to handle interactions. These routing strategies can be as simple or complex as you need them to be. URS uses routing strategies to send interactions from one *target* to another, as needed, until the interactions have been successfully processed.

The Routing Platform SDK allows you to write .NET applications that combine logic from your custom application with the router-based logic of URS, in order to solve many common interaction-related tasks.

This document tells you where you can go to get more information about URS. It also contains a brief overview of the features of the Routing Platform SDK, followed by code snippets that show how to implement the basic functions you will need to write applications that work with URS.

Universal Routing Server Overview

The best way to start learning about Universal Routing Server (URS) is by getting a copy of the Universal Routing 8.1 Reference Manual. This book tells you how to work with routing strategies, objects, functions, options, and statistics. It also includes a detailed list of *Related Documentation Resources*, which discusses other sources of information that can be useful when you are working with Genesys Universal Routing.

After becoming familiar with the information in the *Universal Routing Reference Manual* and related documentation, you can start using the routing API that is exposed by the Platform SDK. As you learn about Genesys routing, it is important to keep in mind that the main purpose of the Platform SDK routing API is to work as a complement to the complex capabilities already available from URS, not to act as a replacement. This API makes it easier to resolve difficult interaction-related tasks by combining the capabilities of URS with logic from your custom application.

To create routing strategies, you use either Genesys Composer, which lets you create SCXML-based strategies, or Interaction Routing Designer (IRD), which creates strategies in the Genesys IRL routing language. Once the URS environment is established, you then use the Platform SDK routing API to give your application control over which routing strategies are selected under a given set of circumstances or what criteria URS uses to choose a particular routing target. For example, your application can select statistics for URS to use in determining which agent group would be the best one to route a particular interaction to.

Two Types of Router API Usage

Platform SDK lets you use two different methodologies in working with URS. The first method involves a **standalone router**. When you use the standalone router method, all of the interaction processing logic, including media control, is handled by the router. This method can be used by calling `RequestLoadStrategy`.

The second method is called the **router-behind** API. This method can be used when you want your application to handle media control, such as attached data or treatments, rather than leaving that up to the router. With this method, the router is normally used only to select resources.

The code snippets in this article include some requests that work with standalone routers and some that work with the router-behind API.

Java

Connecting to Universal Routing Server

The Platform SDK uses a message-based architecture to connect to Genesys servers. In general, Genesys recommends that you use the Protocol Manager Application Block to handle these connections. We also recommend that you use the Message Broker Application Block for your message handling. However, the current version of the Protocol Manager Application Block does not support Universal Routing Server (URS). Because of this, the following code samples show how to connect to URS by using the native protocol object that is part of the Routing Platform SDK.

You can modify the Protocol Manager Application Block to support URS by following the instructions in the section on *Supporting New Protocols* in the article on Using the Protocol Manager Application Block. After doing that, you can use the material in the article on Connecting to a Server to modify your code to use the Protocol Manager Application Block. In the meantime, here is how to connect to URS using a native protocol object.

First set up `import` statements for the routing namespaces:

```
[Java]

import com.genesyslab.platform.routing.protocol.routingserver.*;
import com.genesyslab.platform.routing.protocol.routingserver.requests.*;
```

After you have set up your import statements, you need to create a `RoutingServerProtocol` object:

```
[Java]

RoutingServerProtocol protocol =
        new RoutingServerProtocol(
                new Endpoint(
                        name, host, port));
protocol.setClientName(clientName);
protocol.setClientType(clientType);
```

Then you can open your connection to URS:

```
[Java]

protocol.open();
```

Message Handling

Once you have set up your server connection, you can set up Message Broker, the application block that handles the events returned by URS. This section gives you an idea of how to do that, based on the information in the Event Handling article. First, here are the import statements:

```
[Java]

import com.genesyslab.platform.applicationblocks.commons.Action;
import com.genesyslab.platform.applicationblocks.commons.broker.BrokerServiceFactory;
import com.genesyslab.platform.applicationblocks.commons.broker.EventBrokerService;
import com.genesyslab.platform.applicationblocks.commons.broker.MessageIdFilter;
```

Now you can create an Event Broker Service object:

```
[Java]

eventBrokerService = BrokerServiceFactory
        .CreateEventBroker(protocolManagementServiceImpl.getReceiver());
```

For each message you want to receive, you must set up a handler class. Here is a sample:

```
[Java]

class EventInfoHandler implements Action {

        public void handle(Message obj) {
                EventInfo eventInfo = (EventInfo) obj;
                if (eventInfo != null) {
                        System.out.println("EventInfo:\n"
                                + eventInfo.toString());
                }
                ...
        }
}
```

This handler processes any `EventInfo` messages you receive. The handler classes for other messages will have a similar structure. In order to work properly, each handler must be registered with the Event Broker Service. Here is how to register your `EventInfo` handler:

```
[Java]

eventBrokerService.register(new EventInfoHandler(),
        new MessageIdFilter(EventInfo.ID));
```

Working with URS

As mentioned above, there are two basic methods for using the Platform SDK to work with URS. This section contains examples of both the standalone router and router-behind APIs.

Standalone Router

The Routing Platform SDK allows you to control which routing strategy is executed on a given routing point, while leaving everything else to the routing server. To use this methodology, which is known as "standalone router," issue a RequestLoadStrategy that specifies the routing point and the associated T-Server, and also the location of the routing strategy. Once the routing strategy has been loaded, all interactions arriving on the specified routing point will be processed with that strategy.

The following snippet shows how to do this:

```
[Java]

RequestLoadStrategy requestLoadStrategy = RequestLoadStrategy.create();
requestLoadStrategy.setTServer("TheT-Server");
requestLoadStrategy.setRoutingPoint("TheRoutingPoint");
requestLoadStrategy.setPath("<Path to the strategy>");

Message response = protocol.request(requestLoadStrategy);
```

URS will respond to your request with an EventInfo, an example of which is shown here:

```
'EventInfo' (2) attributes:
                R_Message [str] = "ATTENTION: Strategy has been loaded from ooo-file."
                R_cdn_status [int] = 1 [Loaded]
                R_cdn [str] = "RP_sip1"
                R_ErrorCode [int] = 0 [NoError]
                R_tserver [str] = "TServerSip1"
                R_refID [int] = 1
                R_time [str] = "06/30/2011 10:00:29"
                R_path [str] = "<Path>"
```

You can use RequestNotify to check which routing points have been loaded:

```
[Java]

RequestNotify requestNotify =
    RequestNotify.create();

protocol.send(requestNotify);
```

This request will also return an EventInfo similar to the one shown above.

When you want to stop using the routing strategy you have loaded, for example, if you want to start using a different one, you can issue a RequestReleaseStrategy:

```
[Java]

RequestReleaseStrategy requestReleaseStrategy =
    RequestReleaseStrategy.create();
requestReleaseStrategy.setTServer("TheT-Server");
requestReleaseStrategy.setRoutingPoint("TheRoutingPoint");

Message response = protocol.request(requestReleaseStrategy);
```

Router-Behind API

The router-behind method allows your application code to handle media control. The following example shows how to execute a strategy using RequestExecuteStrategy. This request is different from RequestLoadStrategy in that it only executes a strategy one time, instead of associating a particular strategy with a routing point. To use RequestExecuteStrategy, specify the routing strategy you want to execute and the tenant (contact center) in whose environment the strategy is to be executed, as shown here:

```
[Java]

RequestExecuteStrategy requestExecuteStrategy =
        RequestExecuteStrategy.create();
requestExecuteStrategy.setStrategy("TheRoutingStrategyName");
requestExecuteStrategy.setTenant("TheTenant");

Message response = protocol.request(requestExecuteStrategy);
```

It is important to remember that it can often take a considerable amount of time to process a routing strategy. If your request is correctly formatted and can be executed by URS, then you will immediately receive an `EventExecutionInProgress`. This is a very simple event that only returns the reference ID of your request, as you can see here:

```
'EventExecutionInProgress' ('199')
message attributes:
R_refID [int]   = 2
```

Once your request has successfully executed, you will receive an `EventExecutionAck`. Here is an example of the kind of output you might receive from an EventExecutionAck:

```
'EventExecutionAck' ('200')
message attributes:
R_result [bstr] = KVList:
      'DN' [str] = "701"
      'CUSTOMER_ID' [str] = "TenantForTest"
      'TARGET' [str] = "701_sip@StatServer1.A"
      'SWITCH' [str] = "SipSwitch"
      'NVQ' [int] = 1
      'PLACE' [str] = "701"
      'AGENT' [str] = "701_sip"
      'ACCESS' [str] = "701"
      'VQ' [str] = "1234"
```

Context = ComplexClass(OperationContext):

```
      UserData [bstr] = KVList:
             'ServiceObjective' [str] = ""
             'ServiceType' [str] = "default"
             'CBR-Interaction_cost' [str] = ""
             'RTargetTypeSelected' [str] = "0"
             'CBR-IT-path_DBIDs' [str] = ""
             'RVQDBID' [str] = ""
             'RTargetPlaceSelected' [str] = "701"
             'RTargetAgentSelected' [str] = "701_sip"
             'CBR-actual_volume' [str] = ""
             'RStrategyName' [str] = "##GetTarget"
             'RRequestedSkillCombination' [str] = ""
             'RTargetRuleSelected' [str] = ""
             'RStrategyDBID' [str] = ""
             'RRequestedSkills' [bstr] = KVList:

             'CustomerSegment' [str] = "default"
             'RTargetObjSelDBID' [str] = "984"
             'RTargetObjectSelected' [str] = "701_sip"
             'RTenant' [str] = "TenantForTest"
             'RVQID' [str] = ""
             'CBR-contract_DBIDs' [str] = ""
R_refID [int]   = 0
```

If you have any syntax errors, your request will not execute and you will receive an `EventError`. Here is an example of an EventError:

```
'EventError' (1) attributes:
                R_cdn_status [int] = 0 [Released]
                R_cdn [str] = ""
                R_ErrorCode [int] = 4 [NotAvailable]
                R_tserver [str] = ""
                R_refID [int] = 1
                R_time [str] = ""
                R_path [str] = "<Path>"
```

If, on the other hand, URS has a problem executing your request, you will receive an EventExecutionError, an example of which is shown here:

```
'EventExecutionError' ('201')
message attributes:
R_result [bstr] = KVList:
                                'Reason' [str] = "Rejected"
Context         = ComplexClass(OperationContext):
                UserData [bstr] = KVList:
                                                'PegRejected' [int] = 1
R_refID [int]   = 2
```

There may be times when you want URS to pick a routing target for you. You can use RequestFindTarget for that purpose. As shown in the sample below, you can use a statistic to aid in this selection:

```
[Java]

RequestFindTarget requestFindTarget =
    RequestFindTarget.create();
requestFindTarget.setTenant("TheTenant");
requestFindTarget.setTargets("TheTargetList");
requestFindTarget.setTimeout(5);
requestFindTarget.setStatistic("TheStatistic");
requestFindTarget.setStatisticUsage(StatisticUsage.Max);
requestFindTarget.setVirtualQueue("TheQueue");
requestFindTarget.setPriority(1);
requestFindTarget.setMediaType("TheMediaType");

Message response = protocol.request(requestFindTarget);
```

You can also have URS fetch statistical information for you directly, in case you want to know more about the current conditions in your contact center, perhaps in preparation for a RequestFindTarget. The following example shows how to do this, using RequestGetStatistic:

```
[Java]

RequestGetStatistic requestGetStatistic =
        RequestGetStatistic.create();
requestGetStatistic.setTenant("TheTenant");
requestGetStatistic.setTargets("TheTargetList");
requestGetStatistic.setStatistic("StatAgentsBusy");

Message response = protocol.request(requestGetStatistic);
```

Both RequestFindTarget and RequestGetStatistic return the same messages as RequestExecuteStrategy.

Closing the Connection

When you are finished communicating with URS, you should close the connection, in order to minimize resource utilization:

```
[Java]

protocol.close();
```

.NET

Connecting to Universal Routing Server

The Platform SDK uses a message-based architecture to connect to Genesys servers. In general, Genesys recommends that you use the Protocol Manager Application Block to handle these connections. We also recommend that you use the Message Broker Application Block for your message handling. However, the current version of the Protocol Manager Application Block does not support Universal Routing Server (URS). Because of this, the following code samples show how to connect to URS by using the native protocol object that is part of the Routing Platform SDK.

You can modify the Protocol Manager Application Block to support URS by following the instructions in the section on *Supporting New Protocols* in the article on Using the Protocol Manager Application Block. After doing that, you can use the material in the article on Connecting to a Server to modify your code to use the Protocol Manager Application Block. In the meantime, here is how to connect to URS using a native protocol object.

First set up using statements for the routing namespaces:

```
[C#]

using Genesyslab.Platform.Routing.Protocols;
using Genesyslab.Platform.Routing.Protocols.RoutingServer;
using Genesyslab.Platform.Routing.Protocols.RoutingServer.Events;
using Genesyslab.Platform.Routing.Protocols.RoutingServer.Requests;
```

After you have set up your using statements, you need to create a `RoutingServerProtocol` object:

```
[C#]

RoutingServerProtocol protocol =
        new RoutingServerProtocol(
                new Endpoint(
                        name, host, port));
protocol.ClientName = clientName;
protocol.ClientType = clientType;
```

Then you can open your connection to URS:

```
[C#]

protocol.Open();
```

Message Handling

Once you have set up your server connection, you can set up Message Broker, the application block that handles the events returned by URS. This section gives you an idea of how to do that, based on the information in the Event Handling article.

First, here is the using statement:

```
[C#]

using Genesyslab.Platform.ApplicationBlocks.Commons.Broker;
```

Now you can create an Event Broker Service object:

```
[C#]

eventBrokerService =
        BrokerServiceFactory.CreateEventBroker(protocol);
```

For each message you want to receive, you must set up a handler. Here is a sample:

```
[C#]

private void OnEventExecutionAck(IMessage theMessage)
{
        EventExecutionAck eventExecutionAck =
                theMessage as EventExecutionAck;
        if (eventExecutionAck != null)
        {
                writeToLogArea("EventExecutionAck:\n"
                        + eventExecutionAck + "\n");
                ...
        }
}
```

This handler processes any `EventExecutionAck` messages you receive. The handler classes for other messages will have a similar structure. In order to work properly, each handler must be registered with the Event Broker Service. Here is how to register handlers for three of the URS events:

```
[C#]

eventBrokerService.Register(this.OnEventExecutionAck, new
MessageIdFilter(EventExecutionAck.MessageId));
eventBrokerService.Register(this.OnEventInfo, new MessageIdFilter(EventInfo.MessageId));
eventBrokerService.Register(this.OnEventError, new MessageIdFilter(EventError.MessageId));
```

Working with URS

As mentioned above, there are two basic methods for using the Platform SDK to work with URS. This section contains examples of both the standalone router and router-behind APIs.

Standalone Router

The Routing Platform SDK allows you to control which routing strategy is executed on a given routing point, while leaving everything else to the routing server. To use this "standalone router" methodology, issue a `RequestLoadStrategy` that specifies the routing point and the associated T-Server, and also the location of the routing strategy. Once the routing strategy has been loaded, all interactions arriving on the specified routing point will be processed with that strategy.

The following snippet shows how to do this:

```
[C#]

RequestLoadStrategy requestLoadStrategy =
    RequestLoadStrategy.Create();
requestLoadStrategy.TServer = "TheT-Server";
requestLoadStrategy.RoutingPoint = "TheRoutingPoint";
requestLoadStrategy.Path = "<Path to the strategy>";

IMessage response = protocol.Request(requestLoadStrategy);
```

URS will respond to your request with an `EventInfo`, an example of which is shown here:

```
'EventInfo' (2) attributes:
                R_Message [str] = "ATTENTION: Strategy has been loaded from ooo-file."
                R_cdn_status [int] = 1 [Loaded]
                R_cdn [str] = "RP_sip1"
                R_ErrorCode [int] = 0 [NoError]
                R_tserver [str] = "TServerSip1"
                R_refID [int] = 1
                R_time [str] = "06/30/2011 10:00:29"
                R_path [str] = "<Path>"
```

You can use `RequestNotify` to check which strategies have been loaded to routing points:

```
[C#]

RequestNotify requestNotify =
    RequestNotify.Create();

protocol.Send(requestNotify);
```

This request will also return an `EventInfo` similar to the one shown above.

When you want to stop using the routing strategy you have loaded, for example, if you want to start using a different one, you can issue a `RequestReleaseStrategy`:

```
[C#]

RequestReleaseStrategy requestReleaseStrategy =
    RequestReleaseStrategy.Create();
requestReleaseStrategy.TServer = "TheT-Server";
requestReleaseStrategy.RoutingPoint = "TheRoutingPoint";

IMessage response = protocol.Request(requestReleaseStrategy);
```

Router-Behind API

The router-behind method allows your application code to handle media control. The following example shows how to execute a strategy using `RequestExecuteStrategy`. This request is different from `RequestLoadStrategy` in that it only executes a strategy one time, instead of associating a particular strategy with a routing point. To use `RequestExecuteStrategy`, specify the routing strategy you want to execute and the tenant (contact center) in whose environment the strategy is to be executed, as shown here:

```
[C#]

RequestExecuteStrategy requestExecuteStrategy =
        RequestExecuteStrategy.Create();
requestExecuteStrategy.Strategy = "TheRoutingStrategyName";
requestExecuteStrategy.Tenant = "TheTenant";

IMessage response = protocol.Request(requestExecuteStrategy);
```

It is important to remember that it can often take a considerable amount of time to process a routing strategy. If your request is correctly formatted and can be executed by URS, then you will immediately receive an `EventExecutionInProgress`. This is a very simple event that only returns the reference ID of your request, as you can see here:

```
'EventExecutionInProgress' ('199')
message attributes:
R_refID [int]   = 2
```

Once your request has successfully executed, you will receive an `EventExecutionAck`. Here is an example of the kind of output you might receive from an `EventExecutionAck`:

```
'EventExecutionAck' ('200')
message attributes:
R_result [bstr] = KVList:
                                'DN' [str] = "701"
                                'CUSTOMER_ID' [str] = "TenantForTest"
                                'TARGET' [str] = "701_sip@StatServer1.A"
                                'SWITCH' [str] = "SipSwitch"
                                'NVQ' [int] = 1
                                'PLACE' [str] = "701"
                                'AGENT' [str] = "701_sip"
                                'ACCESS' [str] = "701"
                                'VQ' [str] = "1234"
Context         = ComplexClass(OperationContext):
                UserData [bstr] = KVList:
                                                'ServiceObjective' [str] = ""
                                                'ServiceType' [str] = "default"
                                                'CBR-Interaction_cost' [str] = ""
                                                'RTargetTypeSelected' [str] = "0"
                                                'CBR-IT-path_DBIDs' [str] = ""
                                                'RVQDBID' [str] = ""
                                                'RTargetPlaceSelected' [str] = "701"
                                                'RTargetAgentSelected' [str] = "701_sip"
                                                'CBR-actual_volume' [str] = ""
                                                'RStrategyName' [str] = "##GetTarget"
                                                'RRequestedSkillCombination' [str] = ""
                                                'RTargetRuleSelected' [str] = ""
                                                'RStrategyDBID' [str] = ""
                                                'RRequestedSkills' [bstr] = KVList:

                                                'CustomerSegment' [str] = "default"
                                                'RTargetObjSelDBID' [str] = "984"
                                                'RTargetObjectSelected' [str] = "701_sip"
                                                'RTenant' [str] = "TenantForTest"
                                                'RVQID' [str] = ""
                                                'CBR-contract_DBIDs' [str] = ""
R_refID [int]   = 0
```

If you have any syntax errors, your request will not execute and you will receive an EventError. Here is an example of an `EventError`:

```
'EventError' (1) attributes:
                R_cdn_status [int] = 0 [Released]
                R_cdn [str] = ""
                R_ErrorCode [int] = 4 [NotAvailable]
                R_tserver [str] = ""
                R_refID [int] = 1
                R_time [str] = ""
                R_path [str] = "<Path>"
```

If, on the other hand, URS has a problem executing your request, you will receive an `EventExecutionError`, an example of which is shown here:

```
'EventExecutionError' ('201')
message attributes:
R_result [bstr] = KVList:
                                'Reason' [str] = "Rejected"
Context         = ComplexClass(OperationContext):
                UserData [bstr] = KVList:
                                                'PegRejected' [int] = 1
R_refID [int]   = 2
```

There may be times when you want URS to pick a routing target for you. You can use `RequestFindTarget` for that purpose. As shown in the sample below, you can use a statistic to aid in this selection:

```
[C#]

RequestFindTarget requestFindTarget =
    RequestFindTarget.Create();
requestFindTarget.Tenant = "TheTenant";
requestFindTarget.Targets = "TheTargetList";
requestFindTarget.Timeout = 5;
requestFindTarget.Statistic = "TheStatistic";
requestFindTarget.StatisticUsage = StatisticUsage.Max;
requestFindTarget.VirtualQueue = "TheQueue";
requestFindTarget.Priority = 1;
requestFindTarget.MediaType = "TheMediaType";

IMessage response = protocol.Request(requestFindTarget);
```

You can also have URS fetch statistical information for you directly, in case you want to know more about the current conditions in your contact center, perhaps in preparation for a `RequestFindTarget`. The following example shows how to do this, using `RequestGetStatistic`:

```
[C#]

RequestGetStatistic requestGetStatistic =
    RequestGetStatistic.Create();
requestGetStatistic.Tenant = "TheTenant";
requestGetStatistic.Targets = "TheTargetList";
requestGetStatistic.Statistic = "StatAgentsBusy";

IMessage response = protocol.Request(requestGetStatistic);
```

Both `RequestFindTarget` and `RequestGetStatistic` return the same messages as `RequestExecuteStrategy`.

Closing the Connection

When you are finished communicating with URS, you should close the connection, in order to minimize resource utilization:

[C#]

```
protocol.Close();
```

Component Overviews

Component Overviews

- Using the Log Library
- Using the Switch Policy Library

Using the Log Library

Java

The purpose of the Platform SDK Log Library is to present an easy-to-use API for logging messages in custom-built applications. Depending on how you configure your logger, you can quickly and easily have log messages of different verbose levels written to any of the following targets:

- Genesys Message Server
- Console
- Specified log files

This document provides some key considerations about how to configure and use this component, as well as code examples that will help you work with the Platform SDK Log Library in your own projects.

Introduction to Loggers

When working with custom Genesys loggers, the first step is to understand the basic process of creating and maintaining your logger. How do you create a logger instance? What configuration options should you use, and how and when can those options be changed? What is required to clean up once the logger is no longer useful?

Luckily, the main functions and lifecycle of a logger are easy to understand. The following list outlines the basic process required to create and maintain your customized logger. For more detailed information, check the *Lifecycle of a Logger section*, or the specific code examples related to *Creating a Logger*, *Customizing your Logger*, *Using your Logger*, or *Cleaning Up Your Code*.

1. Use the `LoggerFactory` to create a `RootLogger` instance.
2. Reconfigure the default `RootLogger` settings, if desired.
    - Create a `LogConfiguration` instance to enable and configure log targets. Depending on the setting you assign, log messages can be sent to the Console, to a Genesys Message Server, or to one or more user-defined log files.
    - `LoggerPolicy` property gives you control over how log messages are created and formatted, or allows you to overwrite `MessageServerProtocol` properties.
3. Use your logger for logging messages.
4. Dispose of the `RootLogger` instance when it is no longer needed. (You can also close the logger if it will be reused in the future.)

Lifecycle of a Logger

There are two possible states for a logger, as shown in the lifecycle diagram below.

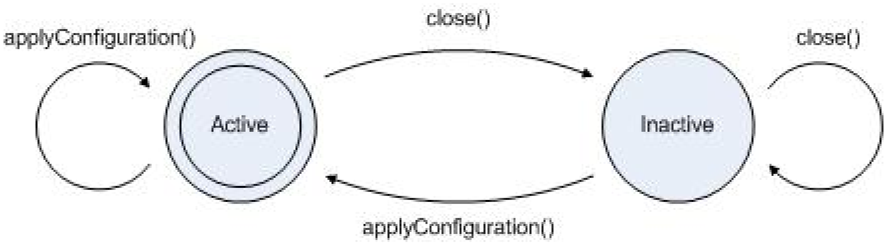

Your logger begins in the active state once it is created. If you did not specify any configuration options during creation, then all messages with a verbose level at least equal to `VerboseLevel.Trace` are logged to the Console by default.

You can use the `applyConfiguration` method to change logger configuration settings from either an active or inactive state:

- If the logger is active when you call this method, then all messages being processed will be handled before the logger is stopped and reconfigured. Note that any file targets (from both the old and new configurations) are not closed automatically when the logger is reconfigured, although file targets *can* be closed and a new log file segment started if the new `Segmentation` settings require this.
- If the logger is inactive when you call this method, then it is activated after the new configuration settings are applied.

You can use the `close` method to make the logger inactive without disposing of it. All messages being processed when this method is called are processed before the logger is stopped. Once the logger is inactive, no further messages are processed until after the `applyConfiguration` method is called.

> **Important**
>
> If your logger is connected to Message Server, the logger does not manage the lifecycle of the `MessageServerProtocol` instance. You must manage and close that connection manually.

Creating a Default Logger

This section provides simple code examples that show how to quickly create a logger with the default configuration and use it as part of your application.

Creating a Logger

As always, the starting point is ensuring that you have the necessary Platform SDK libraries referenced and declared in your project. For logging functionality described in this article, that includes the following packages:

- import com.genesyslab.platform.commons.log.*;
- import com.genesyslab.platform.commons.collections.KeyValueCollection;
- import com.genesyslab.platform.management.protocol.messageserver.LogLevel;
- import com.genesyslab.platform.logging.*;
- import com.genesyslab.platform.logging.configuration.*;
- import com.genesyslab.platform.logging.utilities.*;
- import com.genesyslab.platform.logging.runtime.LoggerException;

Once your project is properly configured and coding about to begin, the first task is to create an instance of the `RootLogger` class. This is easy to accomplish with help from the `LoggerFactory` - the only information you need to provide is a logger name that can be used later to configure targets for filtering logging output.

```
[Java]

try{
        // Create a logger instance:
        RootLogger logger = new LoggerFactory("myLoggerName").getRootLogger();
}
catch(LoggerException e){
        // Handle exceptions...
}
```

The default behavior for this logger is to send all messages of `Trace` verbose and higher to the Console. You can change this behavior by using a `LogConfiguration` instance to change configuration settings and then applying those values to the `RootLogger` instance, as shown below, but for now we will accept the default values.

Using your Logger

With a logger created and ready for use, the next step is to generate some custom log messages and ensure that your logger is working correctly.

One way to generate log messages is with the `write` method. Depending on the parameters used with this method, message formatting can either be provided by templates extracted from LMS files, or through the settings that you configure in a `LogEntry` instance. For the example below, the only formatting come from the `LogEntry` parameter.

```
[Java]

LogEntry logEntry;
logEntry = new LogEntry("Sample Internal message.");
logEntry.setId(CommonMessage.GCTI_INTERNAL.getId());
logger.write(logEntry);

logEntry.setMessage("Sample Debug message.");
logEntry.setId(CommonMessage.GCTI_DEBUG.getId());
logger.write(logEntry);
```

You can also generate log messages by using one of the methods listed in the following table.

Message Level	Available Methods
Debug	• debug(Object arg0) • debug(Object arg0, Throwable arg1) • debugFormat(String arg0, object arg1)
Info	• info(Object arg0) • info(Object arg0, Throwable arg1) • infoFormat(String arg0, object arg1)
Interaction	• warn(Object arg0) • warn(Object arg0, Throwable arg1) • warnFormat(String arg0, object arg1)
Error	• error(Object arg0) • error(Object arg0, Throwable arg1) • errorFormat(String arg0, object arg1)
Alarm	• fatalError(Object arg0) • fatalError(Object arg0, Throwable arg1) • fatalErrorFormat(String arg0, object arg1)

These methods do not use any external templates or formatting, relying entirely on the information passed into them. In the examples below, the messages are logged at `Info` and `Debug` verbose levels, without any changes or formatting.

```
[Java]

logger.info("Sample Info message.");
logger.debug("Sample Debug message.");
```

Cleaning Up Your Code

Once you have finished logging messages with your logger, there are two options available: you can close the logger if you want it to be available for reuse later, or dispose of the logger if your application doesn't need it any longer. (Note that you do not have to close a logger before disposing of it.)

Once closed, a logger remains in an inactive state until either the `ApplyConfiguration` method is called or you dispose of the object, as shown in the *Lifecycle of a Logger* diagram above.

```
[Java]
```

```
// closing the logger
logger.close();
// disposing of the logger
logger = null;
```

Customizing your Logger

Now that you know how to create and use a generic logger, it is time to look at some of the configuration options available to alter the behavior of your logger.

The `LogConfiguration` class allows you change application details, specify targets (including Genesys Message Server) for your log messages, and adjust the verbose level you want to report on. You can apply these changes to either a new logger that is created with the `LogFactory`, or to an existing logger by using the `ApplyConfiguration` method.

> **Tip**
>
> `MessageHeaderFormat` property in the `LogConfiguration` class has no effect on records in the Message Server Database due to message server work specificity.

Creating a LogConfiguration to Specify Targets and Verbose Levels

The first step to configuring the settings for your logger is creating an instance of the `LogConfigurationImpl` class and setting some basic parameters that describe your application.

```
[Java]

LogConfigurationImpl logConfigImpl = new LogConfigurationImpl();
logConfigImpl.setApplicationHost("myHostname");
logConfigImpl.setApplicationName("myApplication");
logConfigImpl.setApplicationId(10);
logConfigImpl.setApplicationType(20);
logConfigImpl.setVerbose(VerboseLevel.ALL);
```

Additional `LogConfiguration` properties that can be configured to specify the name of an application-specific LMS file (`MessageFile`) and whether timestamps should use local or UTC format (`TimeUsage`). These steps aren't shown here for brevity; refer to the API Reference for details.

> **Tip**
>
> If logging to the network, timestamps for log entries always use UTC format to avoid confusion. In this case the `TimeUsage` setting specified by your `LogConfiguration` is ignored.

The next step is to assign this implementation to an actual `LogConfiguration` instance. Once you do that, you can specify the target locations where log messages will be sent and the verbose levels accepted by each individual target. (Only messages with a level greater than or equal to the verbose setting will be logged.)

```
[Java]

LogConfiguration config = logConfigImpl;

// configure logging to console
config.getTargets().getConsole().setEnabled(true);
config.getTargets().getConsole().setVerbose(VerboseLevel.TRACE);

// configure logging to system events log
config.getTargets().getNetwork().setEnabled(true);
config.getTargets().getNetwork().setVerbose(VerboseLevel.STANDARD);
```

Adding files to your logger requires one extra step: creating and configuring a FileConfiguration instance that provides details about each log file to be used. For example:

```
[Java]

// add logging to Log file "Log\fulllog" - for all messages
FileConfiguration file = new FileConfigurationImpl(true, VerboseLevel.ALL, "Log/fulllogfile");
file.setMessageHeaderFormat(MessageHeaderFormat.FULL);
config.getTargets().getFiles().add(file);

// add logging to Log file "Log\infolog" - for Info (and higher) messages
file = new FileConfigurationImpl(true, VerboseLevel.TRACE, "Log/infologfile");
file.setMessageHeaderFormat(MessageHeaderFormat.SHORT);
config.getTargets().getFiles().add(file);
```

> **Warning**
> Each file added as a target must have a unique name. If two or more items are added to the file collection with the same name, only one file target will be created with the lowest specified verbose level. Other settings will be taken from one of the items using the same filename, but there is no way to predict which item will be used.

In the example above, the first line of code ensures that your logger will process messages for all verbose levels - but each target location has its own setting afterwards that specifies what level of messages can be logged by that target. You also can enable or disable individual logging targets by changing and then reapplying the settings in the `LogConfiguration` instance.

Once you have created and configured the `LogConfiguration` instance, all that remains is to apply those settings to your logger. The following code shows how you can apply these settings to either a new Logger instance, or an already existing logger.

```
[Java]

// applying new configuration to an existing logger
logger.applyConfiguration(config);
```

For more information about using `ApplyConfiguration`, see the *Lifecycle of a Logger* section above and the API Reference entry for that method.

Alternative Ways to Create a LogConfiguration

Another way to create a `LogConfiguration` instance is by parsing a `KeyValueCollection` that contains the appropriate settings. A brief code example of how to accomplish this is provided below.

```
[Java]

KeyValueCollection kvConfig = new KeyValueCollection();
// verbose level of logger will be VerboseLevel.All
kvConfig.addString("verbose","all");

// enable output of info (and higher) messages to console
kvConfig.addString("trace","stdout");

// add file target for debug debug output
kvConfig.addString("debug","Log/dbglogfile");

// Parse the created keyValueCollection. Messages generated during parsing are logged to
Console.
LogConfiguration config = LogConfigurationFactory.parse(kvConfig, (ILogger)new
Log4JLoggerFactoryImpl ());
```

Finally, you can also create a `LogConfiguration` by parsing an `org.w3c.dom.Element` that contains the appropriate settings. This Element can be created manually, or obtained from a `CfgApplication` object.

Dealing with Sensitive Log Data

There are two optional filters included as part of the common Platform SDK functionality that can be used to handle senstive log data. These are not part of the Log Library but are discussed here to help ensure sensitive data is properly considered and handled in any custom applications involving logging.

- Hiding Data in Logs - The first option to protect sensitive data is to prevent it from being printed to log files at all.
- Adding Predefined Prefix/Postfix Strings - The second option does not hide the sensitive information directly, but adds user-defined strings around values for selected key-value pairs. This makes it easy for you to locate and removed sensitive data in case log files need to be shared or distributed for any reason.

For more information about these filters, refer to the `KeyValueOutputFilter` documentation in this API Reference.

Hiding Data in Logs

In the code except below, the `KeyValuePrinter` class is used to hide any value of a key-value pair where the key is "Password":

```
[Java]

KeyValueCollection kvOptions = new KeyValueCollection();
KeyValueCollection kvData = new KeyValueCollection();
kvData.addString("Password", KeyValuePrinter.HIDE_FILTER_NAME);
```

```
KeyValuePrinter hidePrinter = new KeyValuePrinter(kvOptions, kvData);
KeyValuePrinter.setDefaultPrinter(hidePrinter);

KeyValueCollection col = new KeyValueCollection();
col.addString("Password", "secretPassword");
```

As result, the KeyValueCollection log output will have the "secretPassword" value printed as "*****". Values for other keys will display as usual.

Adding Predefined Prefix/Postfix Strings

The `PrefixPostfixFilter` class is designed to give you the ability to wrap parts of the log with predefined prefix/postfix strings. This makes it possible to easily filter out sensitive information from an already-printed log file when such a necessity arises.

In the code except below, the KeyValuePrinter is set to wrap "Password" key-value pairs in the "<###" (prefix), "###>" (postfix) strings:

```
[Java]

KeyValueCollection kvData = new KeyValueCollection();
KeyValueCollection kvPPfilter = new KeyValueCollection();
KeyValueCollection kvPPOptions = new KeyValueCollection();
kvPPfilter.addString(KeyValuePrinter.CUSTOM_FILTER_TYPE, "PrefixPostfixFilter");
kvPPOptions.addString(PrefixPostfixFilter.KEY_PREFIX_STRING, "<###");
kvPPOptions.addString(PrefixPostfixFilter.VALUE_POSTFIX_STRING, "###>");
kvPPOptions.addString(PrefixPostfixFilter.KEY_POSTFIX_STRING, ">");
kvPPOptions.addString(PrefixPostfixFilter.VALUE_PREFIX_STRING, "<");
kvPPfilter.addList(KeyValuePrinter.CUSTOM_FILTER_OPTIONS, kvPPOptions);
kvData.addList("Password", kvPPfilter);
KeyValuePrinter.setDefaultPrinter(
        new KeyValuePrinter(new KeyValueCollection(), kvData));

KeyValueCollection col = new KeyValueCollection();
col.addString("test", "secretPassword");
```

As result, the `KeyValueCollection` log output will have the "Password-secretPassword" key-value printed as "<###Password-secretPassword###>", leaving all other key-values printed as normal.

.NET

The purpose of the Platform SDK Log Library is to present an easy-to-use API for logging messages in custom-built applications. Depending on how you configure your logger, you can quickly and easily have log messages of different verbose levels written to any of the following targets:

- Genesys Message Server
- Console
- .NET Trace
- Windows System Log (Application Log only)
- Specified log files

This document provides some key considerations about how to configure and use this component, as well as code examples that will help you work with the Platform SDK Log Library for .NET in your own projects.

Introduction to Loggers

When working with custom Genesys loggers, the first step is to understand the basic process of creating and maintaining your logger. How do you create a logger instance? What configuration options should you use, and how and when can those options be changed? What is required to clean up once the logger is no longer useful?

Luckily, the main functions and lifecycle of a logger are easy to understand. The following list outlines the basic process required to create and maintain your customized logger. For more detailed information, check the *Lifecycle of a Logger section*, or the specific code examples related to *Creating a Logger*, *Customizing your Logger*, *Using your Logger*, or *Cleaning Up Your Code*.

1. Use the LoggerFactory to create an ILogger instance.
2. Reconfigure the default ILogger settings, if desired.
   - NetworkProtocol property allows you to specify a MessageSeverProtocol instance. This will let your ILogger instance send messages to Genesys Message Server.
   - LoggerPolicy property gives you control over how log messages are created and formatted, or allows you to overwrite MessageServerProtocol properties.
   - LogConfiguration class allows you to configure other aspects of your ILogger instance, which are applied with the ApplyConfiguration method.
3. Use the logger for logging messages.
4. Dispose of the ILogger instance when it is no longer needed. (You can also close the logger if it will be reused in the future.)

Lifecycle of a Logger

There are two possible states for a logger, as shown in the lifecycle diagram below.

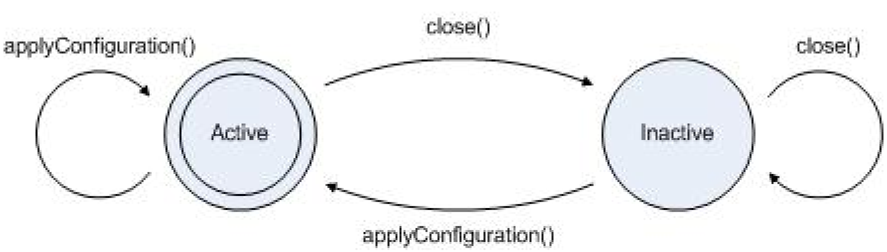

Your logger begins in the active state once it is created. If you did not specify any configuration options during creation, then all messages with a verbose level at least equal to `VerboseLevel.Trace` are logged to the Console by default.

You can use the `ApplyConfiguration` method to change logger configuration settings from either an active or inactive state:

- If the logger is active when you call this method, then all messages being processed will be handled before the logger is stopped and reconfigured. Note that any file targets (from both the old and new configurations) are not closed automatically when the logger is reconfigured, although file targets can be closed and a new log file segment started if the new `Segmentation` settings require this.
- If the logger is inactive when you call this method, then it is activated after the new configuration settings are applied.

You can use the `Close` method to make the logger inactive without disposing of it. All messages being processed when this method is called are processed before the logger is stopped. Once the logger is inactive, no further messages are processed until after the `ApplyConfiguration` method is called.

> **Important**
> If your logger is connected to Message Server, the logger does not manage the lifecycle of the `MessageServerProtocol` instance. You must manage and close that connection manually.

Creating a Default Logger

This section provides simple code examples that show how to quickly create a logger with the default configuration and use it as part of your application.

Creating a Logger

As always, the starting point is ensuring that you have the necessary Platform SDK libraries referenced and declared in your project. For logging functionality, that includes the following namespaces:

- `Genesyslab.Platform.Commons.Logging`
- `Genesyslab.Platform.Logging`
- `Genesyslab.Platform.Logging.Configuration`
- `Genesyslab.Platform.Logging.Utilities`

Once your project is properly configured and coding about to begin, the first task is to create an instance of the `ILogger` class. This is easy to accomplish with help from the `LoggerFactory` - the only information you need to provide is a logger name that can be used later to configure targets for filtering logging output.

[C#]

```
IRootLogger logger = LoggerFactory.CreateRootLogger("myLoggerName");
```

The default behavior for this logger is to send all messages of `Trace` verbose and higher to the

Console. You can change this behavior by using the `ILogConfiguration` interface to pass configuration settings into the `LoggerFactory`, as shown below, but for this example we will accept the default values.

Using your Logger

Now that your logger is created and ready for use, the next step is to generate some custom log messages and ensure that the logger is working correctly.

One way to generate log messages is with the `Write` method. Message formatting is provided either by templates extracted from LMS files or directly from a `LogEntry` parameter, depending on what information you pass into the method. For the example below, LMS file templates provide formatting.

```
[C#]

//log the message with standard id: "9999|STANDARD|GCTI_INTERNAL|Internal error '%s' occurred"
//formatting template is extracted from LMS file
logger.Write((int)CommonMessage.GCTI_INTERNAL, "Sample Internal message.");

//log the message with standard id: "9900|DEBUG|GCTI_DEBUG|%s"
//formatting template is extracted from LMS file
logger.Write((int)CommonMessage.GCTI_DEBUG, "Sample Debug message.");
```

You can also generate log messages by using one of the methods listed in the following table.

Message Level	Available Methods
Debug	• Debug(object message) • Debug(object message, Exception exception) • DebugFormat(string format, params object [] args)
Info	• Info(object message) • Info(object message, Exception exception) • InfoFormat(string format, params object [] args)
Interaction	• Warn(object message) • Warn(object message, Exception exception) • WarnFormat(string format, params object [] args)
Error	• Error(object message) • Error(object message, Exception exception) • ErrorFormat(string format, params object [] args)

Message Level	Available Methods
Alarm	• FatalError(object message) • FatalError(object message, Exception exception) • FatalErrorFormat(string format, params object [] args)

These methods do not use any external templates or formatting, relying entirely on the information passed into them. In the examples below, the messages are logged at `Info` and `Debug` verbose levels, without any changes or formatting.

```
[C#]

logger.Info("Sample Info message.");
logger.Debug("Sample Debug message.");
```

Cleaning Up Your Code

Once you have finished logging messages with your logger, there are two options available: you can close the logger if you want it to be available for reuse later, or dispose of the logger if your application doesn't need it any longer. (Note that you do not have to close a logger before disposing of it.)

Once closed, a logger remains in an inactive state until either the `ApplyConfiguration` or `Dispose` method is called, as shown in the lifecycle diagram above.

```
[C#]

//closing the logger
logger.Close();
...
//disposing of the logger
logger.Dispose();
```

Customizing your Logger

Now that you know how to create and use a generic logger, it is time to look at some of the configuration options available to alter the behavior of your logger.

The `LogConfiguration` class allows you change application details, specify targets (including Genesys Message Server) for your log messages, and adjust the verbose level you want to report on. You can apply these changes to either a new logger that is created with the `LogFactory`, or to an existing logger by using the `ApplyConfiguration` method.

> **Tip**
> The `setMessageHeaderFormat` method has no effect on records in the Message Server

Database due to message server work specificity.

Creating a LogConfiguration to Specify Targets and Verbose Levels

The first step to configuring the settings for your logger is creating an instance of the `LogConfiguration` class and setting some basic parameters that describe your application.

```
[C#]

LogConfiguration config = new LogConfiguration
{
    ApplicationHost = "myHostname",
    ApplicationName = "myApplication",
    ApplicationId = 10,
    ApplicationType = 20
};
```

Additional `LogConfiguration` properties that can be configured to specify the name of an application-specific LMS file (`MessageFile`) and whether timestamps should use local or UTC format (`TimeUsage`). These steps aren't shown here for brevity; refer to the API Reference for details.

Tip

If logging to the network, as described in *Logging Messages to Genesys Message Server*, timestamps for log entries always use UTC format to avoid confusion. In this case the `TimeUsage` setting specified by your `LogConfiguration` is ignored.

The next step is to specify the target locations where log messages are recorded, and to configure the verbose levels for the logger and for individual targets. (Only messages with a level greater than or equal to the verbose setting will be logged.)

```
[C#]

config.Verbose = VerboseLevel.All;

//configure logging to console
config.Targets.Console.IsEnabled = true;
config.Targets.Console.Verbose = VerboseLevel.Trace;

//configure logging to system events log
config.Targets.System.IsEnabled = true;
config.Targets.System.Verbose = VerboseLevel.Standard;

//add logging to Log file "Log\fulllog" - for all messages
config.Targets.Files.Add(new FileConfiguration(true, VerboseLevel.All, "Log/fulllogfile"));
//add logging to Log file "Log\infolog" - for Info (and higher) messages
config.Targets.Files.Add(new FileConfiguration(true, VerboseLevel.Trace, "Log/infologfile"));
```

In the example above, the first line of code ensures that your logger will process messages for all verbose levels - but each target location has its own setting afterwards that specifies what level of messages can be logged by that target. You also can enable or disable individual logging targets by changing and then reapplying the settings in the LogConfiguration.

> **Warning**
> Each file added as a target must have a unique name. If two or more items are added to the file collection with the same name, only one file target will be created with the lowest specified verbose level. Other settings will be taken from one of the items using the same filename, but there is no way to predict which item will be used.

Once you have created and configured the LogConfiguration instance, all that remains is to apply those settings to your logger. The following code shows how you can apply these settings to either a new Logger instance, or an already existing logger.

```
[C#]

//applying new configuration to an existing logger
logger.ApplyConfiguration(config);
...
//apply new configuration to a new logger when it is created
IRootLogger newlogger = LoggerFactory.CreateRootLogger("NewLoggerName", config);
```

For more information about using ApplyConfiguration, see the logger lifecycle section above and the API Reference entry for that method.

Alternative Ways to Create a LogConfiguration

Another way to create a LogConfiguration is by parsing a KeyValueCollection that contains the appropriate settings. A brief code example of how to accomplish this is provided below.

```
[C#]

KeyValueCollection kvConfig = new KeyValueCollection();
//verbose level of logger will be VerboseLevel.All
kvConfig.Add("verbose","all");
//enable output of info (and higher) messages to console
kvConfig.Add("trace","stdout");
//add file target for debug debug output
kvConfig.Add("debug","Log/dbglogfile");
//Parse the created keyValueCollection. Messages generated during parsing are logged to
Console.
LogConfiguration config = LogConfigurationFactory.Parse(kvConfig, new ConsoleLogger());
```

Finally, you can also create a LogConfiguration by parsing an XElement that contains the appropriate settings, as shown below.

```
[C#]

XElement xElementConfig =
        new XElement("CfgApplication",
                new XElement("options",
                        new XElement("list_pair",
                                new XAttribute("key","log"),
                                        XElement.Parse("<str_pair key=\"verbose\" value =
\"all\"/>"),
                                        XElement.Parse("<str_pair key=\"trace\" value =
```

```
\"stdout\"/>"),
                                        XElement.Parse("<str_pair key=\"debug\" value = \"Log/
dbglogfile\"/>"),
                        )
                )
        );
LogConfiguration config = LogConfigurationFactory.Parse(xElementConfig, new ConsoleLogger());
```

Although the `XElement` can be created manually (as shown above), it is much more likely that it will be obtained from a `CfgApplication` object. The following code example illustrates how this can be done.

```
[C#]

ConfService confservice=null;
//...
//initializing the ConfService
//...
CfgApplication cfgApp = confservice.RetrieveObject<CfgApplication>(
        new CfgApplicationQuery{Name = "Sample Application"});
XElement xElementConfig = cfgApp.ToXml();
LogConfiguration config = LogConfigurationFactory.Parse(xElementConfig, new ConsoleLogger());
```

Logging Messages to Genesys Message Server

Creating a connection with Genesys Message Server is similar to setting other targets for your logger, but contains a couple of additional steps. Several new settings are required to determine how your logger handles buffering and spooling when sending log messages over the network. Once that is complete, you also have to create (and manage) a protocol object that connects to Message Server.

The following example shows how this can be accomplished. For details and additional information about the properties being configured, refer to the appropriate API Reference entries.

```
[C#]

LogConfiguration config = new LogConfiguration {Verbose = VerboseLevel.All};
config.Targets.Network.IsEnabled = true;
config.Targets.Network.Verbose = VerboseLevel.All;
config.Targets.Network.Buffering = Buffering.On|Buffering.KeepOnProtocolChange;
config.Targets.Network.SpoolFile = "temp/spool";

//create and open connection to message server
MessageServerProtocol protocol = new MessageServerProtocol(
        new Endpoint(myApplication, myHostname, myPort));
protocol.Open();

IRootLogger logger = LoggerFactory.CreateRootLogger("mySample");
logger.NetworkProtocol = protocol;
logger.ApplyConfiguration(config);
```

It is important to remember that any connection created to Message Server is not managed automatically by the Logger lifecycle. You are responsible to manage and dispose of the connection manually.

Dealing with Sensitive Log Data

There are two optional filters included as part of the common Platform SDK functionality that can be used to handle sensitive log data. These are not part of the Log Library for .NET, but are discussed here to help ensure sensitive data is properly considered and handled in any custom applications involving logging.

- Hiding Data in Logs - The first option to protect sensitive data is to prevent it from being printed to log files at all.
- Adding Predefined Prefix/Postfix Strings - The second option does not hide the sensitive information directly, but adds user-defined strings around values for selected key-value pairs. This makes it easy for you to locate and removed sensitive data in case log files need to be shared or distributed for any reason.

Hiding Data in Logs

In the code except below, the `KeyValuePrinter` class is used to hide any value of a key-value pair where the key is "Password":

```
[C#]

KeyValueCollection kvOptions = new KeyValueCollection();
KeyValueCollection kvData = new KeyValueCollection();
kvData["Password"] = KeyValuePrinter.HideFilterName;
KeyValuePrinter hidePrinter = new KeyValuePrinter(kvOptions, kvData);
KeyValuePrinter.DefaultPrinter = hidePrinter;

KeyValueCollection col = new KeyValueCollection();
col["Password"] = "secretPassword";
```

As result, the `KeyValueCollection` log output will have the "secretPassword" value printed as "*****". Values for other keys will display as usual.

Adding Predefined Prefix/Postfix Strings

The `PrefixPostfixFilter` class is designed to give you the ability to wrap parts of the log with predefined prefix/postfix strings. This makes it possible to easily filter out sensitive information from an already-printed log file when such a necessity arises.

In the code except below, the `KeyValuePrinter` is set to wrap "Password" key-value pairs in the "`<###`" (prefix), "`###>`" (postfix) strings:

```
[C#]

KeyValueCollection kvData = new KeyValueCollection();
KeyValueCollection kvPPfilter = new KeyValueCollection();
KeyValueCollection kvPPOptions = new KeyValueCollection();
kvPPfilter[KeyValuePrinter.CustomFilterType] = typeof(PrefixPostfixFilter).FullName;
kvPPOptions[PrefixPostfixFilter.KeyPrefixString] = "<###";
kvPPOptions[PrefixPostfixFilter.ValuePrefixString] = "<";
kvPPOptions[PrefixPostfixFilter.ValuePostfixString] = "###>";
kvPPOptions[PrefixPostfixFilter.KeyPostfixString] = ">";
kvPPfilter[KeyValuePrinter.CustomFilterOptions] = kvPPOptions;
kvData["Password "] = kvPPfilter;
```

```
KeyValuePrinter.DefaultPrinter = new KeyValuePrinter(new KeyValueCollection(), kvData);
KeyValueCollection col = new KeyValueCollection();
col["Password"] = "myPassword";
```

As result, the `KeyValueCollection` log output will have the "Password-secretPassword" key-value printed as "<###Password-secretPassword###>", leaving all other key-values printed as normal.

Using the Switch Policy Library

This document shows how to add simple T-Server functionality to your applications by using the Switch Policy Library.

The Platform SDK Switch Policy Library (SPL) can be used in applications that need to perform agent-related switch activity with a variety of T-Servers, without knowing beforehand what kinds of T-Servers will be used. It simplifies these applications by indicating which switch functions are available at any given time and also by showing how you can use certain switch features in your applications. However, if your application works with only one kind of T-Server, you may want to have your application communicate directly with the T-Server, rather than using SPL.

Switch Policy Library Overview

Some telephony applications need to work with more than one type of switch. Unfortunately, however, one switch may not perform a particular telephony function in the same way as another switch. This means that it can be useful to have an abstraction layer of some kind when working with multiple switches, so that you do not need custom code for each switch that is used by the application. The Switch Policy Library is designed with just this kind of abstraction in mind.

Java

Setting Up Switch Policy Library

SPL should be used by your agent desktop applications as a library, which means that it would be located within the agent desktop application shown above. The application can call SPL for guidance on how to send requests to or process events from your T-Server, as shown in the *Code Samples* section.

SPL is driven by an XML-based configuration file that supports many commonly-used switches in performing agent-related functions. Your application can query SPL to determine whether a particular feature is supported for the switch you want to work with. If a feature you need is not supported for the switches you need to work with, you can make a copy of the default configuration file and modify it as needed.

> **Important**
> Genesys does not support modifications to the SPL configuration file. Any modifications you make are performed at your own risk.

A copy of the default configuration file is included inside the Switch Policy Library JAR file. You can extract the XML configuration file from `switchpolicy.jar`, modify it, and pass it as an argument to the corresponding method of the `SwitchPolicyServiceFactory` factory class.

Code Samples

This section contains examples of how to perform useful functions with SPL.

These samples each require a valid instance of the `ISwitchPolicyService`, which can be created as shown here:

```
[Java]

ISwitchPolicyService service =
        SwitchPolicyServiceFactory.createSwitchPolicyService();
```

> **Tip**
>
> The `DN` classes specified below implement the `IDNContext` interface, while the `Party` classes implement the `IPartyContext` interface, and the `Call` classes implement the `ICallContext` interface.

Customizing the XML Configuration File

The following code samples create a service using the default configuration. As noted above, Genesys does not support modifications to the default SPL configuration file. Should you decide to assume the risk of creating a custom XML configuration file, your application can access this file as shown here:

```
[Java]

public void serviceCreationWithParent(ApplicationContext parent) {
        final String file_path =
                "<Path to XML Configuration File>";
        FileSystemResource resource =
                new FileSystemResource(file_path);
        ISwitchPolicyService service =
                SwitchPolicyServiceFactory.createSwitchPolicyService(parent, resource);
}
```

Get A Phone Set Configuration

On some switches, phone sets are presented as more than one *Directory Number* (DN). These DNs may also have different types, such as *Position* and *Extension*. Because these configurations vary by switch type, an application needs to know how the phone set configuration for a particular switch is structured. For example, it needs to know how many DNs are used to represent a phone set, and what their types are. To retrieve this phone set configuration information, perform the following steps:

1. Create an instance of `PhoneSetConfigurationContext`, specifying the switch type.

2. Call `ISwitchPolicyService.getPolicy`, using this `PhoneSetConfigurationContext`.
3. Analyze the returned `PhoneSetConfigurationPolicy`. The `PhoneSetConfigurationPolicy.getConfigurations` method will return all possible phone set configurations for the specified switch.

The following code snippet shows how to do this:

```
[Java]

ISwitchPolicyService service =
        SwitchPolicyServiceFactory.createSwitchPolicyService();
PhoneSetConfigurationContext context =
        new PhoneSetConfigurationContext("SwitchName");
PhoneSetConfigurationPolicy policy =
        service.getPolicy(PhoneSetConfigurationPolicy.class, context);
System.out.println(policy);
```

Get Phone Set Availability Information

When working with a phone set, additional information about the included DNs may be required. This could include information about which of the DNs should be available to the end user (for example, which ones should be visible in the user interface), which of them is callable, and which number (the `Callable Number`) the application should use to reach the agent who is logged into the phone set. To retrieve this phone set availability information, perform the following steps:

1. Create an instance of `DNAvailabilityContext` and populate it with the following required information:
    1. Specify the switch type.
    2. Specify the Agent ID.
    3. Fill the DN collection with valid implementations of `IDNContext`.
2. Call `ISwitchPolicyService.getPolicy`, using this `DNAvailabilityContext`.
3. Analyze the returned `DNAvailabilityPolicy`. The `DNAvailabilityPolicy.getDNStatuses` method will return availability information for each DN in the request.

The following code snippet shows how to do this:

```
[Java]

String extDN = "1001";
String posDN = "2001";
String agentID = "9999";
// logout, in service
DNAvailabilityContext context = new DNAvailabilityContext(switchname);
context.setAgentId(agentID);
DNContextStub ext = new DNContextStub(); // implements IDNContext interface
ext.setIdentifier(extDN);
ext.setAgentStatus(AgentStatus.LOGOUT);
ext.setServiceStatus(ServiceStatus.IN_SERVICE);
ext.setType(AddressType.DN);

DNContextStub pos = new DNContextStub(); // implements IDNContext interface
pos.setIdentifier(posDN);
pos.setAgentStatus(AgentStatus.LOGOUT);
pos.setServiceStatus(ServiceStatus.IN_SERVICE);
pos.setType(AddressType.Position);
```

```
ArrayList<IDNContext> dns = new ArrayList<IDNContext>();
dns.add(ext);
dns.add(pos);
context.setDNs(dns);

// here service is correctly initialized instance of ISwitchPolicyService
DNAvailabilityPolicy policy = service.getPolicy(DNAvailabilityPolicy.class, context);
System.out.println(policy);
```

Get Function Availability Information for the Current Context

Some switches differ in when they allow certain functions to be performed. Also, some functions can always be performed on certain switches, while others may be impossible to perform. For example, `RequestMergeCalls` can never be performed on some switches. For other functions, whether or not the function can be performed varies depending on context. For example, on some switches `RequestReleaseCall` can only be used when a call is in a *Held, Dialing*, or *Established* state, while on other switches it is also possible to release a call when it is in a *Ringing* state. In addition to this, on some switches the phone set is presented as more than one *Directory Number* (DN) and each DN can have a different type, such as *Position* and *Extension*. Some functions are allowed for both types, while some other functions may be restricted to a certain DN type. To retrieve this kind of function availability information for the current context, perform the following steps:

1. Create an instance of `FunctionHandlingContext` and populate it with the following required information:
    1. Specify the switch type.
    2. Specify the request by calling the `setMessage` method.
    3. Describe the context as fully as possible.
2. Call `ISwitchPolicyService.getPolicy`, using this `FunctionHandlingContext`.
3. Analyze the returned `FunctionAvailabilityPolicy`. If the specified request is possible in the given context, the `getIsFunctionAvailable` method will return true. However, if the request is not supported, SPL will return null.

The following code snippet shows how to do this:

```
[Java]

DNContextStub dn = new DNContextStub();// implements IDNContext
dn.setIdentifier("1001");
dn.setType(AddressType.DN);
dn.setAgentStatus(AgentStatus.READY);
dn.setServiceStatus(ServiceStatus.IN_SERVICE);

DNContextStub otherdn = new DNContextStub();
otherdn.setIdentifier("2001");
otherdn.setType(AddressType.DN);
otherdn.setAgentStatus(AgentStatus.READY);
otherdn.setServiceStatus(ServiceStatus.IN_SERVICE);

PartyContextStub mainparty = new PartyContextStub();// implements IPartyContext
mainparty.setIdentifier("9841");
mainparty.setStatus(PartyStatus.ESTABLISHED);
mainparty.setIsConferencing(true);
mainparty.setIsTransferring(true);
```

```
mainparty.setDN(dn);

PartyContextStub otherParty = new PartyContextStub();
otherParty.setIdentifier(mainparty.getIdentifier());
otherParty.setStatus(PartyStatus.ESTABLISHED);
otherParty.setIsConferencing(true);
otherParty.setIsTransferring(true);
otherParty.setDN(otherdn);

CallContextStub call = new CallContextStub();// Implements ICallContext
call.setStatus(CallStatus.ESTABLISHED);
call.setDestination(mainparty);
call.setOrigination(otherParty);
call.setIdentifier(mainparty.getIdentifier());
call.setConferencing(true);
call.setTransferring(true);
call.setParties(Arrays.<IPartyContext> asList(mainparty, otherParty));

for (String swtype : new String[] { swtypeA4400Classic, swtypeA4400Emul, swtypeA4400Subs }) {
        for (CallType callType : GEnum.valuesBy(CallType.class)) {
                FunctionHandlingContext context = new FunctionHandlingContext(swtype);
                context.setMessage(RequestHoldCall.create());
                context.setDN(dn);
                mainparty.setCallType(callType);
                otherParty.setCallType(callType);
                call.setCallType(callType);
                context.setParty(mainparty);
                context.setCall(call);
                FunctionAvailabilityPolicy policy =
service.getPolicy(FunctionAvailabilityPolicy.class, context);
                System.out.println(policy);
        }
}
```

Get Instructions On How To Implement a Feature

Some switches differ in how certain features can be accessed. The majority of their features may map directly to individual switch functions, but this is not always so. For example, for some switches it is not possible to log the agent out while the agent is in the ready state. So, the feature which implements agent logout for these switches would require two steps:

1. Make sure the agent is in a NotReady state
2. Log the agent out

SPL implements a *feature handler* for each feature that it supports. To create and run a feature handler, perform the following steps:

1. Create a new instance of FunctionHandlingContext and populate it with the following required information:
    1. Specify the switch type.
    2. Specify the request by calling the setMessage method. This step can be omitted if a feature handler is going to be created by using the featureName parameter in the ISwitchPolicyService.createFeatureHandler(String featureName, FunctionHandlingContext context) method.
    3. Provide a valid Protocol instance by calling the setProtocol method.

    4. Describe the context as fully as possible.

2. Call `ISwitchPolicyService.createFeatureHandler` and pass this `FunctionHandlingContext`, either alone or with the name of the feature.
3. Call the `beginExecute` method of `IFeatureHandler` on the returned handler, passing the same instance of `FunctionHandlingContext`.
4. The remainder of the processing depends on the implementation, but the general approach is to perform the following actions while the status of the handler is `Executing`:
    1. Receive event from T-Server.
    2. Update `FunctionHandlingContext` based on the received event.
    3. Assign the received event by calling the `setMessage` method of your `FunctionHandlingContext` instance.
    4. Call the handle method of `IFeatureHandler` passing with it the updated `FunctionHandlingContext`.

The following code snippet shows how to do this:

```
[Java]

public void LogoutReadyAgent(Protocol protocol,
        ISwitchPolicyService service, String thisDN, String switchType)
        throws IllegalStateException, InterruptedException,
        SwitchPolicyException {

        FunctionHandlingContext context =
                new FunctionHandlingContext(switchType);
        context.setMessage(RequestAgentLogout.create(thisDN));
        DNContextStub dn = new DNContextStub(); // implements IDNContext interface
        dn.setAgentStatus(AgentStatus.READY);
        dn.setAgentWorkMode(AgentWorkMode.Unknown);
        dn.setIdentifier(thisDN);
        dn.setType(AddressType.DN);
        dn.setServiceStatus(ServiceStatus.IN_SERVICE);
        context.setDN(dn);
        context.setProtocol(protocol);

        IFeatureHandler handler = service.createFeatureHandler(context);
        if (handler != null) {
                handler.beginExecute(context);
                while (handler.getStatus() == FeatureStatus.EXECUTING) {
                        Message message = (Message) protocol.receive();
                        // update context due to received message
                        // .....
                        context.setMessage(message);
                        handler.handle(context);
                }
        }
}
```

Get Instructions On How To Accomplish Complex Functionality

Your application may sometimes need access to functionality that depends on the switch type. For example, when an application receives events from the T-Server, the way a given event's fields are used can depend on both the call scenario and the switch type. To retrieve this information, perform the following steps:

1. Create a `MessageHandlingContext` and populate it with the following required information:
    1. Name of switch.
    2. Name of handler.
2. Call `ISwitchPolicyService.createMessageHandler`, pass this context into it, and receive the resulting `IMessageHandler`.
3. Call the `IMessageHandler.handle` method on the received handler.

The following code snippet shows how to do this:

```
[Java]

        EventRinging msgRinging =
                EventRinging.create(TimeStamp.create(1249566176, 312000));

        KeyValueCollection p = new KeyValueCollection();
        p.addInt("BusinessCall", 0);
        p.addInt("GCTI_BUSINESS_CALL", 0);
        p.addString("GCTI_SUB_THIS_DN", "18101");
        p.addString("GCTI_SUB_OTHER_DN", "18100");
        p.addString("GCTI_OTHER_DEVICE_NAME", "18100");
        p.addString("GCTI_PARTY_NAME", "18100");
        msgRinging.setExtensions(p);
        msgRinging.setEventSequenceNumber(0x0000000000000399);
        msgRinging.setOtherDN("11100");
        msgRinging.setOtherDNRole(DNRole.RoleOrigination);//
        msgRinging.setOtherTrunk(521);
        msgRinging.setThisDNRole(DNRole.RoleOrigination);//
        msgRinging.setThisDN("11101");
        msgRinging.setDNIS("18101");
        msgRinging.setThisTrunk(522);
        msgRinging.setCallUuid("BTMT3AJVT17QPE364J2DV9V5I000005P");
        msgRinging.setConnID(new ConnectionId("022701b746b29021"));
        msgRinging.setCallID(3648);
        msgRinging.setCallType(CallType.Internal);
        msgRinging.setNetworkCallID(0x1ee07a4a400e0100l);
        msgRinging.setCallState(0);
        msgRinging.setAgentID("18101");
        msgRinging.setPropagatedCallType(CallType.Internal);

        MessageHandlingContext context = new MessageHandlingContext(Switchname);
        context.setHandlerName("OtherDN");
        IMessageHandler othDNH = service.createMessageHandler(context);
        String otherDN = (String) othDNH.handle(msgRinging); System.out.println(otherDN);
```

Add Logging Support

You can add support for logging by providing an application context with a registered `ILogger` bean. This logger will be used by the Switch Policy Library. Here is a code sample:

```
[Java]

AnnotationConfigApplicationContext context =
        new AnnotationConfigApplicationContext();
context.register(ConsoleLogger.class);

// ConsoleLogger implements ILogger interface
context.refresh();
```

```
ISwitchPolicyService service =
        SwitchPolicyServiceFactory.createSwitchPolicyService(context);
```

.NET

Setting Up Switch Policy Library

SPL should be used by your agent desktop applications as a library, which means that it would be located within the agent desktop application shown above. The application can call SPL for guidance on how to send requests to or process events from your T-Server, as shown in the *Code Samples* section.

SPL is driven by an XML-based configuration file that supports many commonly-used switches in performing agent-related functions. Your application can query SPL to determine whether a particular feature is supported for the switch you want to work with. If a feature you need is not supported for the switches you need to work with, you can make a copy of the default configuration file and modify it as needed.

> **Important**
> Genesys does not support modifications to the SPL configuration file. Any modifications you make are performed at your own risk.

A copy of the default configuration file is included inside the Switch Policy Library DLL. There is also a copy in the Bin directory of the Platform SDK installation package. If you need to modify the configuration file, you can use the `app.config` file for SPL to point to your copy.

Code Samples

This section contains examples of how to perform useful functions with SPL.

These samples each require a valid instance of the `ISwitchPolicyService`, which can be created as shown here:

```
[C#]

ISwitchPolicyService policyService =
        SwitchPolicyFactory.CreateSwitchPolicyService();
```

> **Tip**

> The `DN` classes specified below implement the `IDNContext` interface, while the `Party` classes implement the `IPartyContext` interface, and the `Call` classes implement the `ICallContext` interface.

Get A Phone Set Configuration

On some switches, phone sets are presented as more than one *Directory Number* (DN). These DNs may also have different types, such as *Position* and *Extension*. Because these configurations vary by switch type, an application needs to know how the phone set configuration for a particular switch is structured. For example, it needs to know how many DNs are used to represent a phone set, and what their types are. To retrieve this phone set configuration information, perform the following steps:

1. Create an instance of `PhoneSetConfigurationContext`, specifying the switch type.
2. Call `ISwitchPolicyService.GetPolicy`, using this `PhoneSetConfigurationContext`.
3. Analyze the returned `PhoneSetConfigurationPolicy`. The `PhoneSetConfigurationPolicy.Configurations` property will contain all possible phone set configurations for the specified switch.

The following code snippet shows how to do this:

```
[C#]

PhoneSetConfigurationContext context =
        new PhoneSetConfigurationContext("SomeSwitch");
PhoneSetConfigurationPolicy policy =
switchPolicyService.GetPolicy<PhoneSetConfigurationPolicy>(context);
foreach (PhoneSetConfiguration configuration in policy.Configurations)
{
        Console.WriteLine(configuration);
}
```

Get Phone Set Availability Information

When working with a phone set, additional information about the included DNs may be required. This could include information about which of the DNs should be available to the end user (for example, which ones should be visible in the user interface), which of them is callable, and which number (the Callable Number) the application should use to reach the agent who is logged into the phone set. To retrieve this phone set availability information, perform the following steps:

1. Create an instance of `DNAvailabilityContext` and populate it with the following required information:
   - Specify the switch type.
   - Specify the Agent ID.
   - Fill the DN collection with valid implementations of `IDNContext`.
2. Call `ISwitchPolicyService.GetPolicy`, using this `DNAvailabilityContext`.
3. Analyze the returned `DNAvailabilityPolicy`. The `DNAvailabilityPolicy.DNStatuses` property will contain availability information for each DN in the request.

The following code snippet shows how to do this:

```
[C#]

private static void DemonstrateDNAvailability(ISwitchPolicyService service)
{
        DNAvailabilityContext dnacontext =
                new DNAvailabilityContext("SomeSwitch");
        dnacontext.AgentId = "AgentLogin1000";
        dnacontext.DNs.Add(new Dn
        {
                AgentStatus = AgentStatus.Ready,
                        Identifier = "1000",
                        ServiceStatus = ServiceStatus.InService,
                        Type = AddressType.DN
        });
        dnacontext.DNs.Add(new Dn
        {
                AgentStatus = AgentStatus.Ready,
                        Identifier = "2000",
                        ServiceStatus = ServiceStatus.InService,
                        Type = AddressType.Position
        });

        DNAvailabilityPolicy dnpolicy =
                service.GetPolicy<DNAvailabilityPolicy>(dnacontext);
        DisplayInColor(dnpolicy, ConsoleColor.Red);
}
```

Get Function Availability Information for the Current Context

Some switches differ in when they allow certain functions to be performed. Also, some functions can always be performed on certain switches, while others may be impossible to perform. For example, `RequestMergeCalls` can never be performed on some switches. For other functions, whether or not the function can be performed varies depending on context. For example, on some switches `RequestReleaseCall` can only be used when a call is in a Held, Dialing, or Established state, while on other switches it is also possible to release a call when it is in a Ringing state. In addition to this, on some switches the phone set is presented as more than one *Directory Number* (DN) and each DN can have a different type, such as *Position* and *Extension*. Some functions are allowed for both types, while some other functions may be restricted to a certain DN type. To retrieve this kind of function availability information for the current context, perform the following steps:

1. Create an instance of `FunctionHandlingContext` and populate it with the following required information:
   - Specify the switch type.
   - Specify the request by setting the `Message` property.
   - Describe the context as fully as possible.
2. Call `ISwitchPolicyService.GetPolicy`, using this `FunctionHandlingContext`.
3. Analyze the returned `FunctionAvailabilityPolicy`. If the specified request is possible in the given context, the `IsFunctionAvailable` property will be true. However, if the request is not supported, SPL will return null.

The following code snippet shows how to do this:

[C#]

```
foreach (string switchType in new[] { swTypeA4400Classic, swTypeA4400emul, swTypeA4400Subs })
{
        DNContext dn = new DNContext //implements IDNContext
        {
                Identifier = "1001",
                Type = AddressType.DN,
                AgentStatus = AgentStatus.Ready,
                ServiceStatus = ServiceStatus.InService,
                DndStatus = FunctionStatus.Off,
                ForwardStatus = FunctionStatus.Off
        };

        DNContext otherDN = new DNContext
        {
                Identifier = "2001",
                Type = AddressType.DN,
                AgentStatus = AgentStatus.Ready,
                ServiceStatus = ServiceStatus.InService,
                DndStatus = FunctionStatus.Off,
                ForwardStatus = FunctionStatus.Off
        };

        foreach (CallType callType in Enum.GetValues(typeof(CallType)))
        {
                PartyContext mainParty = new PartyContext //implements IPartyContext
                {
                        Identifier = "1002",
                        Status = PartyStatus.Established,
                        CallType = callType,
                        IsConferencing = true,
                        IsTransferring = true,
                        DN = dn
                };

                PartyContext otherParty = new PartyContext
                {
                        Identifier = "1002",
                        CallType = callType,
                        DN = otherDN,
                        IsConferencing = true,
                        IsTransferring = true,
                        Status = PartyStatus.Established
                };
                CallContextStub ccontext = new CallContextStub //implements ICallContext
                {
                        CallType = callType,
                        Destination = mainParty,
                        Origination = otherParty,
                        Identifier = "1002",
                        IsConferencing = true,
                        IsTransferring = true,
                        Parties = new List<IPartyContext>{mainParty,otherParty},
                        Parent = null//no parentCall - our call is solitary call.
                };

                FunctionHandlingContext context = new FunctionHandlingContext(switchType)
                {
                        Message = RequestHoldCall.Create(),
                        DN = dn,
                        Party = mainParty,
                        Call = ccontext
```

```
                };
                FunctionAvailabilityPolicy policy =
service.GetPolicy<FunctionAvailabilityPolicy>(context);

                Console.WriteLine(policy);
        }
}
```

Get Instructions On How To Implement a Feature

Some switches differ in how certain features can be accessed. The majority of their features may map directly to individual switch functions, but this is not always so. For example, for some switches it is not possible to log the agent out while the agent is in the ready state. So, the feature which implements agent logout for these switches would require two steps:

1. Make sure the agent is in a `NotReady` state
2. Log the agent out

SPL implements a feature handler for each feature that it supports. To create and run a feature handler, perform the following steps:

1. Create a new instance of `FunctionHandlingContext` and populate it with the following required information:
   - Specify the switch type.
   - Specify the request by setting the `Message` property. This step can be omitted if the feature handler is created by using the `featureName` parameter in the `ISwitchPolicyService.CreateFeatureHandler(String featureName, FunctionHandlingContext context)` method.
   - Provide a valid `IProtocol` instance as the value of the `Protocol` property.
   - Describe the context as fully as possible.
2. Call the `ISwitchPolicyService.CreateFeatureHandler` and pass this `FunctionHandlingContext`, either alone or with the name of the feature.
3. Call the `BeginExecute` method on the returned handler, passing the same instance of `FunctionHandlingContext`.
4. The remainder of the processing depends on the implementation, but the general approach is to perform the following actions while the status of the handler is `Executing`:
   1. Receive event from TServer.
   2. Update `FunctionHandlingContext` based on the received event.
   3. Assign the received event to the `Message` property of your `FunctionHandlingContext` instance.
   4. Call the Handle method of `IFeatureHandler` passing with it the updated `FunctionHandlingContext`.

The following code snippet shows how to do this:

```
[C#]

private static void LoginReadyAgent(IProtocol protocol,
        ISwitchPolicyService service, string thisdn, string agentID)
```

```
{
        FunctionHandlingContext context = new FunctionHandlingContext("SomeSwitch");
        RequestAgentLogin requestAgentLogin = RequestAgentLogin.Create();
        requestAgentLogin.ThisDN = thisdn;
        requestAgentLogin.AgentID = agentID;
        requestAgentLogin.AgentWorkMode = AgentWorkMode.AutoIn;
        context.Message = requestAgentLogin;
        context.Protocol = protocol;

        IFeatureHandler loginHandler = service.CreateFeatureHandler(context);

        if(loginHandler == null)
        {
                protocol.Send(requestAgentLogin);
                // Process the incoming events for the scenario
                return;
        }

        // Processing feature handler
        loginHandler.BeginExecute(context);
        while (loginHandler.Status == FeatureStatus.Executing)
        {
                context.Message = context.Protocol.Receive();
                // Update the context based on the received T-Server event
                loginHandler.Handle(context);
        }
}
```

Get Instructions On How To Accomplish Complex Functionality

Your application may sometimes need access to functionality that depends on the switch type. For example, when an application receives events from the T-Server, the way a given event's fields are used can depend on both the call scenario and the switch type. To retrieve this information, perform the following steps:

1. Create a `MessageHandlingContext` and populate it with the following required information:
   - Name of switch.
   - Name of handler.
2. Call `ISwitchPolicyService.CreateMessageHandler`, pass this context into it, and receive the resulting `IMessageHandler`.
3. Call the `IMessageHandler.Handle` method on the received handler.

The following code snippet shows how to do this:

```
[C#]

private static void DemonstrateMessageHandler(ISwitchPolicyService service)
{
        EventRinging message = EventRinging.Create();
        message.ThirdPartyDN = "12345";
        message.DNIS = "18009870987";
        message.CallType = CallType.Internal;
        message.OtherDN = "9875";
        MessageHandlingContext context35 =
                new MessageHandlingContext("AlcatelA4400DHS3::Classic")
                { HandlerName = "OtherDN" };
        IMessageHandler handler = service.CreateMessageHandler(context35);
```

```
        string res = (string)handler.Handle(message);
        DisplayInColor(res, ConsoleColor.Yellow);
}
```

Add Logging Support

To add logging support, carry out the following steps:

1. Create an instance of `IUnityContainer` and register an anonymous instance or type mapping for the `ILogger` interface.
2. Pass the `IUnityContainer` created during the previous step to the factory method, which creates an instance of `ISwitchPolicyService`.

The following code snippet shows how to do this:

```
[C#]

IUnityContainer root = new UnityContainer();
root.RegisterInstance(new ConsoleLogger());
ISwitchPolicyService service =
        SwitchPolicyFactory.CreateSwitchPolicyService(root);
```

SPL also provides the following options:

- Your application can log the topmost messages into a distinct log. To use this option, call the `CreateSwitchPolicyService(IUnityContainer container, ILogger logger)` method of the `SwitchPolicyServiceFactory` class. The passed logger (if it is not null) will be used for logging the topmost messages.
- You can configure any switch container to use a specific logger. Objects created by the Unity container (feature handlers, policy providers and so on) can use the container to resolve the `ILogger` for further logging.

> **Tip**
>
> the classes provided by SPL resolve the `ILogger` (if there is one) at creation time. So, if your application changes the `ILogger` resolution rule for the root container that was previously passed into the `SwitchPolicyService` constructor after the corresponding method call, this will not affect:
>
> - Existing instances.
> - Objects which are created in the container(s), for which special `ILogger` mapping rule is configured.

Supported Functions

As mentioned above, SPL is driven by a configuration file that makes it possible to support a wide variety of switch functions. The following table shows functions that are supported by SPL at installation time, using the default configuration file.

Switch Functions Supported by SPL At Installation Time

Switch Function	Description
DN and Agent Functions	
RequestAgentLogin	Logs in the agent specified by the AgentId parameter to the ACD group specified by the parameter.
RequestAgentLogout	Logs the agent out of the ACD group specified by the Queue parameter.
RequestAgentNotReady	Sets a state in which the agent is not ready to receive calls. The agents telephone set is specified by the DN parameter; the ACD group into which the agent is logged is specified by the Queue parameter.
RequestAgentReady	Sets a state in which the agent is ready to receive calls. The agents phone set is specified by the DN parameter; the ACD group into which the agent is logged is specified by the Queue parameter.
RequestCallForwardCancel	Sets the Forwarding feature to Off for the telephony object that is specified by the DN parameter.
RequestCallForwardSet	Sets the Forwarding feature to On for the telephony object that is specified by the DN parameter.
RequestCancelMonitoring	A request by a supervisor to cancel monitoring the calls delivered to the agent. If this request is successful, T-Server distributes EventMonitoringCancelled to all clients registered on the supervisor's and agent's DNs.
RequestMonitorNextCall	A request by a supervisor to monitor (be automatically conferenced in as a party on) the next call delivered to an agent. Supervisors can request to monitor one subsequent call or all calls until the request is explicitly canceled. If a request is successful, EventMonitoringNextCall is distributed to all clients registered on the supervisor's and agent's DNs. Supervisors start monitoring each call in Mute mode. To speak, they must execute the function
RequestSetDNDOff	Sets the Do-Not-Disturb (DND) feature to Off for the telephony object specified by the DN parameter.
RequestSetDNDOn	Sets the Do-Not-Disturb (DND) feature to On for the telephony object specified by the DN parameter.
RequestSetMuteOff	On an existing conference call, cancels the Mute mode for the party specified by the DN parameter.
RequestSetMuteOn	On an existing conference call, sets Mute mode for the party specified by the DN parameter.
Call Handling Functions	
RequestAlternateCall	On behalf of the telephony object specified by the DN parameter, places the active call specified by

Switch Function	Description
	thecurrent_conn_id parameter on hold and connects the call specified by the held_conn_id parameter.
RequestAnswerCall	Answers the alerting call specified by the conn_id parameter.
RequestAttachUserData	On behalf of the telephony object specified by the DN parameter, attaches the user data structure specified by the user_data parameter to the T-Server information that is related to the call specified by the conn_id parameter.
RequestClearCall	Deletes all parties, that is, all telephony objects, from the call specified by conn_id and disconnects the call.
RequestCompleteConference	Completes a previously-initiated conference by merging the held call specified by the held_conn_id parameter with the active consultation call specified by the current_conn_id parameter on behalf of the telephony object specified by the DN. Assigns the held_conn_id to the resulting conference call. Clears the consultation call specified by the current_conn_id parameter.
RequestCompleteTransfer	On behalf of the telephony object specified by the DN parameter, completes a previously initiated two-step transfer by merging the held call specified by the conn_id parameter with the active consultation call specified by the current_conn_id parameter. Assigns held_conn_id to the resulting call. Releases the telephony object specified by the DN parameter from both calls and clears the consultation call specified by the current_conn_id parameter.
RequestDeleteFromConference	A telephony object specified by DN deletes the telephony object specified by dn_to_drop from the conference call specified by conn_id. The client that invokes this service must be a party on the call in question.
RequestDeletePair	On behalf of the telephony object specified by the DN parameter, deletes the key-value pair specified by the key parameter from the user data attached to the call specified by the conn_id parameter.
RequestDeleteUserData	On behalf of the telephony object specified by the DN parameter, deletes all of the user data attached to the call specified by the conn_id parameter.
RequestHoldCall	On behalf of the telephony object specified by the DN parameter, places the call specified by the conn_id parameter on hold.
RequestInitiateConference	On behalf of the telephony object specified by the DN parameter, places the existing call specified by the conn_id parameter on hold and originates a consultation call from the same telephony object to the called party, which is specified by the

Switch Function	Description
	destination parameter with the purpose of a conference call.
RequestInitiateTransfer	On behalf of the telephony object specified by the DN parameter, places the existing call specified by the conn_id parameter on hold and originates a consultation call from the same telephony object to the called party, which is specified by the destination parameter for the purpose of a two-step transfer.
RequestListenDisconnect	On an existing conference call, sets Deaf mode for the party specified by the listener_dn parameter. For example, if two agents wish to consult privately, the subscriber may temporarily be placed in Deaf mode.
RequestListenReconnect	On an existing conference call, cancels Deaf mode for the party defined by the listener_dn parameter.
RequestMakeCall	Originates a regular call from the telephony object specified by the DN parameter to the called party specified by the Destination parameter.
RequestMakePredictiveCall	Makes a predictive call from the thisDN DN to the otherDN called party. A predictive call occurs before any agent-subscriber interaction is created. For example, if a fax machine answers the call, no agent connection occurs. The agent connection occurs only if there is an actual subscriber available on line.
RequestMergeCalls	On behalf of the telephony object specified by the DN parameter, merges the held call specified by the held_conn_id parameter with the active call specified by the current_conn_id parameter in a manner specified by the merge_type parameter. The resulting call will have the same conn_id as the held call.
RequestMuteTransfer	Initiates a transfer of the call specified by the conn_id parameter from the telephony object specified by the DN parameter to the party specified by the destination parameter; completes the transfer without waiting for the destination party to pick it up. Releases the telephony object specified by the DN parameter from the call.
RequestQueryCall	Requests the information specified by info_type about the telephony object specified by conn_id. If the query type is supported, the requested information will be returned in EventPartyInfo.
RequestReconnectCall	Releases the telephony object specified by the DN parameter from the active call specified by the current_conn_id parameter and retrieves the previously held call, specified by the held_conn_id parameter, to the same object. This function is commonly used to clear an active call and to return to a held call, or to cancel a consult call (due to

Switch Function	Description
	lack of an answer, because the device is busy, and so on) and then to return to a held call.
RequestRedirectCall	Requests that the call be redirected, without an answer, from the party specified by the DN parameter to the party specified by the dest_dn parameter.
RequestRegisterAddress	Registers for a DN. Your application must register the DN before sending the RequestAgentLogin.
RequestReleaseCall	Releases the telephony object specified by the DN parameter from the call specified by the conn_id parameter.
RequestRetrieveCall	Connects the held call specified by the conn_id parameter to the telephony object specified by the DN parameter.
RequestSendDtmf	On behalf of the telephony object specified by the DN parameter, sends the digits that are expected by an interactive voice response system.
RequestSetCallInfo	Changes the call attributes. Warning: Improper use of this function may result in unpredictable behavior on the part of the T-Server and the Genesys Framework. If you have any doubt on how to use it, please consult with Genesys.
RequestSetMessageWaitingOff	Sets the Message Waiting indication to off for the telephony object specified by the DN parameter.
RequestSetMessageWaitingOn	Sets the Message Waiting indication to on for the telephony object specified by the DN parameter.
RequestSetMuteOff	On an existing conference call, cancels the Mute mode for the party specified by the DN parameter.
RequestSetMuteOn	On an existing conference call, sets Mute mode for the party specified by the DN parameter.
RequestSingleStepConference	Adds a new party to an existing call and creates a conference.
RequestSingleStepTransfer	Transfers the call from a specified directory number DN that is currently engaged in the call specified by the conn_id parameter to a destination DN that is specified by the destination parameter.
RequestUnregisterAddress	Unregisters a DN.
RequestUpdateUserData	On behalf of the telephony object specified by the DN parameter, updates the user data that is attached to the call specified by the conn_id parameter with the data specified by the user_data parameter.

Legacy Topics

Topics in this section are no longer applicable for new development, but are maintained here for backwards compatibility.

- Using the Message Broker Application Block
- Event Handling Using the Message Broker Application Block
- Using the Protocol Manager Application Block
- Connecting to a Server Using the Protocol Manager Application Block
- Explicitly Choosing a Netty or Mina Connection Layer

Using the Message Broker Application Block

> **Important**
> This application block is considered a legacy product starting with release 8.1.1. Documentation is provided for backwards compatibility, but new development should consider using the improved method of message handling.

The Message Broker Application Block is a reusable production-quality component that makes it easy for your applications to handle events in an efficient way. It has been designed using industry best practices and provided with source code so it can be used "as is," extended, or tailored if you need to. Please see the License Agreement for details.

For information on the other application blocks that ship with the Genesys SDKs, consult Introducing the Platform SDK.

Java

Installing the Message Broker Application Block

Software Requirements

To work with the Message Broker Application Block, you must ensure that your system meets the software requirements established in the Genesys Supported Operating Environment Reference Guide, as well as meeting the following minimum software requirements:

- JDK 1.6 or higher

Building the Message Broker Application Block

To build the Message Broker Application Block:

1. Open the `<Platform SDK Folder>\applicationblocks\messagebroker` folder.
2. Run either `build.bat` or `build.sh`, depending on your platform.

This will create the `messagebrokerappblock.jar` file, located within the `<Platform SDK Folder>\applicationblocks\messagebroker\dist\lib` directory.

Working with the Message Broker Application Block

You can find basic information on how to use the Message Broker Application Block in the article on Event Handling Using the Message Broker Application Block.

Configuring Message Broker

When you first work with Message Broker, you will probably use a single instance of `EventBrokerService`. This means that all messages coming into your application will first pass through this single instance, as shown in below. Note that configuration diagrams used here do not show the Protocol Manager Application Block, in order to focus on the architecture of Message Broker.

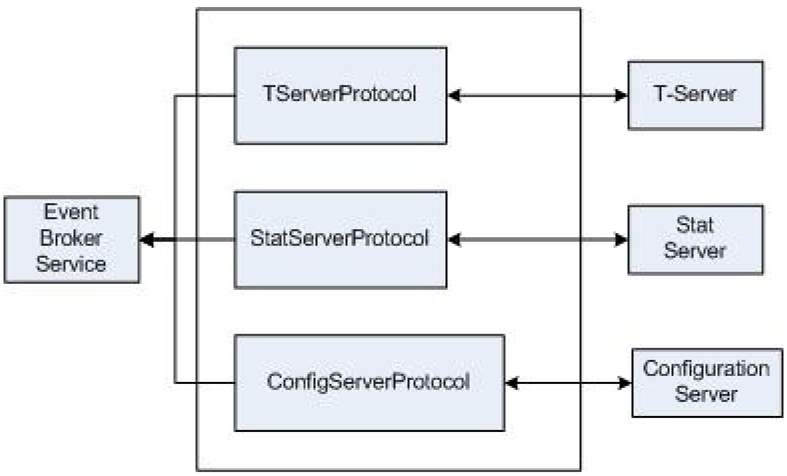

However, there may be high-traffic scenarios that require multiple instances of Message Broker. This might happen if you have one or more servers whose events use so much of Message Broker's processing time that events from other servers must wait for an unacceptable amount of time. In that case, you could dedicate an instance of `EventBrokerService` to the appropriate server.

For example, you may have a scenario in which you frequently receive large volumes of statistics. To handle that situation, you could dedicate an `EventBrokerService` instance to Stat Server. In other situations, you might regularly receive large amounts of Configuration Layer data from Configuration Server. You could handle this in a similar way by giving Configuration Server its own instance of `EventBrokerService`, as shown here:

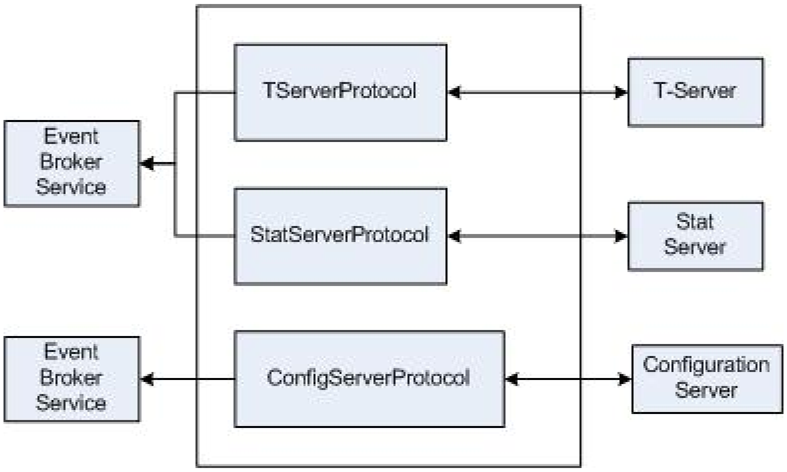

Sometimes you may have large message volumes for each server, in which case you could use a separate instance of `EventBrokerService` for each server, as shown here.

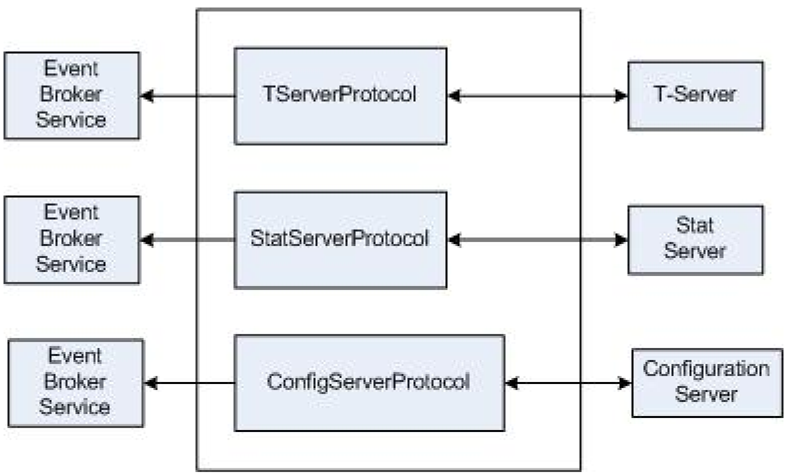

Using Message Filters

Message Broker comes with several types of message filters. You can filter on individual messages using `MessageIdFilter` or `MessageNameFilter`. In most cases you will want to use `MessageIdFilter`, as it is more efficient than `MessageNameFilter`. You can also use a `MessageRangeFilter` to filter on several messages at a time.

As shown in the article on Event Handling Using the Message Broker Application Block, you can specify these filters when you register an event handler with the Event Broker Service. Here is a sample of how to set up a `MessageIdFilter`:

```
[Java]

eventBrokerService.register(new StatPackageOpenedHandler(),
        packageEvents);
```

There may be times when you want to process several events in the same event handler. In such cases, you can use a MessageRangeFilter, which will direct all of these events to that handler. Here is a sample of how to set up the filter:

```
[Java]

int[] messageRange = new int[] {EventPackageOpened.ID, EventPackageClosed.ID};
MessageRangeFilter packageStatusEvents = new MessageRangeFilter
        (messageRange);
eventBrokerService.register(new StatPackageStatusChangedHandler(),
        packageStatusEvents);
```

Your event handler might look something like this:

```
[Java]

class StatPackageStatusChangedHandler implements Action {

        public void handle(Message obj) {
                // Common processing goes here...
                if (obj.messageId() == EventPackageOpened.ID) {
                        // EventPackageOpened processing goes here...
                } else {
                        // EventPackageClosed processing goes here...
                }
        }
}
```

Some servers use events that have the same name as events used by another server. One example is EventError, which is used by just about every server except Stat Server. The Event Handling Using the Message Broker Application Block article shows how to use a Protocol Description object to filter events by server type in order to avoid confusion when handling these events.

There also may be times when you have several instances of a given server in your environment and you want to filter by a specific one. To do this, first specify an Endpoint for that server, using a name for the server in the Endpoint constructor:

```
[Java]

String statServer1EndpointName = "StatServer1";
Endpoint statServer1Endpoint =
        new Endpoint(statServer1EndpointName, statServer1Uri);
```

Now create the filter:

```
[Java]

MessageIdFilter statServer1EndpointFilter =
        new MessageIdFilter(EventPackageOpened.ID);
```

And set the EndpointName in the filter:

```
[Java]

statServer1EndpointFilter.setEndpointName(statServer1EndpointName);
```

When you register this filter, the handler you specify will only receive messages that were sent from the instance you mentioned above:

```
[Java]

eventBrokerService.register(new StatPackageOpenedHandler_StatServer1(),
        statServer1EndpointFilter);
```

Architecture and Design

The Message Broker Application Block is designed to make it easy for your applications to handle events in an efficient way.

Message Broker allows you to set up individual classes to handle specific events coming from Genesys servers. It receives all of the events from the servers you specify, and sends each one to the appropriate handler class. Message Broker is a high-performance way to hide the complexity of event-driven programming — so you can focus on other areas of your application.

> **Tip**
> Message Broker has been designed for use with the Protocol Manager Application Block. Protocol Manager is another high-performance component that makes it easy for your applications to connect to Genesys servers. You can find basic information on how to use the Protocol Manager Application Block in the article on Connecting to a Server.

The Message Broker Application Block Architecture

The Message Broker Application Block uses a service-based API that enables you to write individual methods that handle one or more events.

For example, you might want to handle every occurrence of `EventAgentLogin` with a specific dedicated method, while there might be other events that you wish to send to a common event-handling method. Message Broker allows you write these methods and register them with an event broker that manages them for you.

Message Filters

Message Broker uses *message filters* to identify specific messages, assign them to specified methods, and route them accordingly.

Design Patterns

This section gives an overview of the design patterns used in the Message Broker Application Block.

Publish/Subscribe Pattern

There are many occasions when one class (the subscriber) needs to be notified when something changes in another class (the publisher). The Message Broker Application Block use the Publish/Subscribe pattern to inform the client application when events arrive from the server.

Factory Method Pattern

It is common practice for a class to include constructors that enable clients of the class instantiate it. There are times, however, when a client may need to instantiate one of several different classes. In some of these situations, the client should not need to decide which class is being created. In this case, a Factory Method pattern is used. The Factory Method pattern lets a class developer define the interface for creating an object, while retaining control of which class to instantiate.

How To Properly Manage the EventBrokerService Lifecycle

Unfortunately, a commonly encountered problem is that users create `EventBrokerService` but do not dispose of it properly. `EventBrokerService` exclusively uses an invoker thread to run an infinite cycle with `MessageReceiver.receive()` and incoming messages handling logic. `EventBroker` is created by user code, so it should be disposed by user code as well. Useful methods are `MessageBrokerService.deactivate()` and `MessageBrokerService.dispose()`.

In PSDK 8.1 this class is deprecated and a new one is added to resolve the problem with thread waiting: `EventReceivingBrokerService`. This new class implements the `MessageReceiver` interface and may be used as external receiver for Platform SDK protocols. In this case, we have no intermediate redundant queue and incoming messages are delivered from protocol(s) to handler(s) directly. This class still requires async invoker to execute messages handling, but in this case the invoker is called once per incoming message, so it's thread is not blocked during the `.receive()` operation.

So, `EventReceivingBrokerService` does not need `.dispose()` and is GC friendly.

> **Tip**
> A similar change has been made to `RequestBrokerService`.

Also note that the `Invoker` instance still represents a "costly" resource (thread) and is managed by user code, so proper attention (allocation/deallocation) is required.

Q: Does it matter if the event broker service is created by the `BrokerServiceFactory` or not?

A: Actually, `BrokerServiceFactory` just creates and activates the corresponding broker instance. So if a broker is created by a call to the factory, it must be disposed of by user code in accordance to its usage there.

.NET

Installing the Message Broker Application Block

Before you install the Message Broker Application Block, it is important to review the software requirements and the structure of the software distribution.

Software Requirements

To work with the Message Broker Application Block, you must ensure that your system meets the software requirements established in the Genesys Supported Operating Environment Reference Guide.

Building the Message Broker Application Block

The Platform SDK distribution includes a `Genesyslab.Platform.ApplicationBlocks.Commons.Broker.dll` file that you can use as is. This file is located in the `bin` directory at the root level of the Platform SDK directory. To build your own copy of this application block, follow the instructions below:

1. Open the `<Platform SDK Folder>\ApplicationBlocks\MessageBroker` folder.
2. Double-click `MessageBroker.sln`.
3. Build the solution.

Working with the Message Broker Application Block

You can find basic information on how to use the Message Broker Application Block in the article on Event Handling Using the Message Broker Application Block.

Configuring Message Broker

When you first work with Message Broker, you will probably use a single instance of `EventBrokerService`. This means that all messages coming into your application will first pass through this single instance, as shown in the figure below. Note that the following configuration diagrams do not show the Protocol Manager Application Block, in order to focus on the architecture of Message Broker.

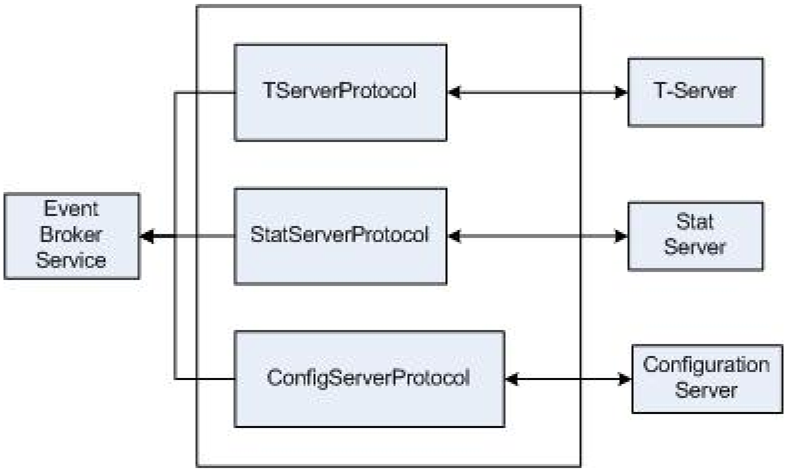

However, there may be high-traffic scenarios that require multiple instances of Message Broker. This might happen if you have one or more servers whose events use so much of Message Broker's processing time that events from other servers must wait for an unacceptable amount of time. In that case, you could dedicate an instance of `EventBrokerService` to the appropriate server.

For example, you may have a scenario in which you frequently receive large volumes of statistics. To handle that situation, you could dedicate an `EventBrokerService` instance to Stat Server. In other situations, you might regularly receive large amounts of Configuration Layer data from Configuration Server. You could handle this in a similar way by giving Configuration Server its own instance of `EventBrokerService`, as shown in the following figure:

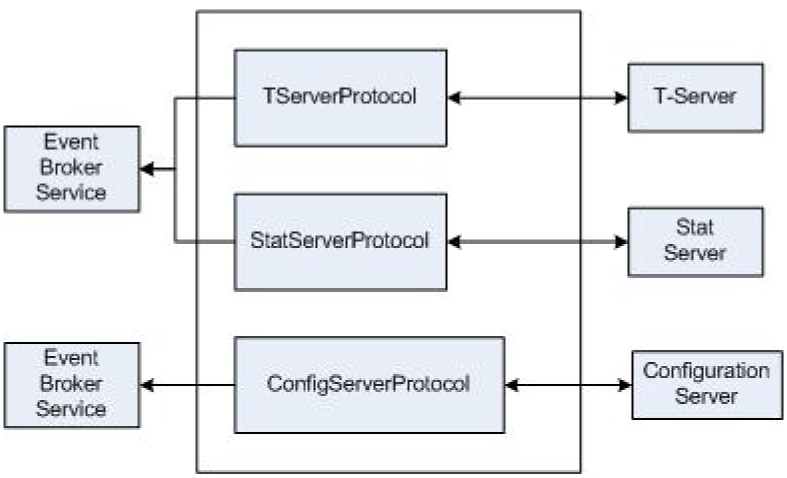

Sometimes you may have large message volumes for each server, in which case you could use a separate instance of `EventBrokerService` for each server, as shown here.

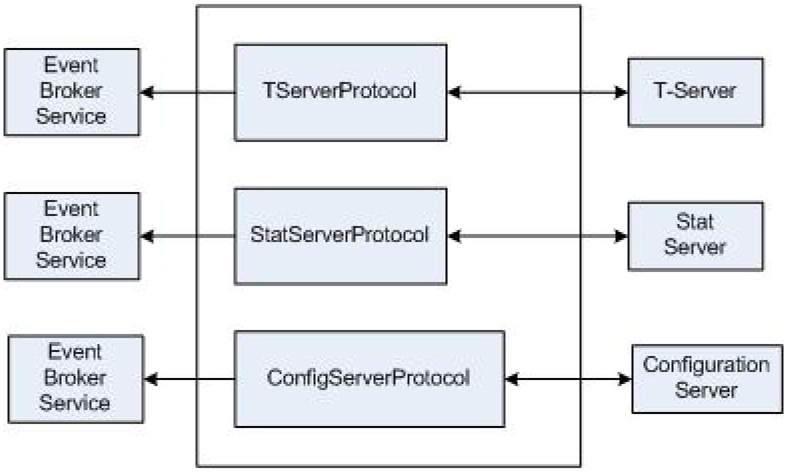

Using Message Filters

Message Broker comes with several types of message filters. You can filter on individual messages using `MessageIdFilter` or `MessageNameFilter`. In most cases you will want to use `MessageIdFilter`, as it is more efficient than `MessageNameFilter`. You can also use a `MessageRangeFilter` to filter on several messages at a time.

As shown in the article on Event Handling Using the Message Broker Application Block in the beginning of this guide, you can specify these filters when you register an event handler with the Event Broker Service. Here is a sample of how to set up a `MessageIdFilter`:

```
[C#]

eventBrokerService.Register(this.OnEventPackageClosed,
        new MessageIdFilter(EventPackageClosed.MessageId));
```

There may be times when you want to process several events in the same event handler. In such cases, you can use a `MessageRangeFilter`, which will direct all of these events to that handler. Here is a sample of how to set up the filter:

```
[C#]

eventBrokerService.Register(this.OnEventPackageStatusChanged, new MessageRangeFilter(new
int[]   {
        EventPackageOpened.MessageId, EventPackageClosed.MessageId}));
```

Your event handler might look something like this:

```
[C#]

private void OnEventPackageStatusChanged(IMessage theMessage)
{
        // Common processing goes here...
        if (theMessage.Id == EventPackageOpened.MessageId)
        {
                // EventPackageOpened processing goes here...
        }
        else
```

```
        {
                // EventPackageClosed processing goes here...
        }
}
```

Some servers use events that have the same name as events used by another server. One example is `EventError`, which is used by just about every server except Stat Server. The Event Handling Using the Message Broker Application Block article shows how to use a Protocol Description object to filter events by server type in order to avoid confusion when handling these events.

There also may be times when you have several instances of a given server in your environment and you want to filter by a specific one. To do this, first specify an `Endpoint` for that server, using a name for the server in the `Endpoint` constructor:

```
[C#]

string statServer1EndpointName = "StatServer1";
Endpoint statServer1Endpoint =
        new Endpoint(statServer1EndpointName, statServer1Uri);
```

Now create the filter:

```
[C#]

MessageIdFilter statServer1EndpointFilter =
        new MessageIdFilter(EventPackageOpened.MessageId);
```

And set the `EndpointName` property of the filter:

```
[C#]

statServer1EndpointFilter.EndpointName = statServer1EndpointName;
```

When you register this filter, the handler you specify will only receive messages that were sent from the instance you mentioned above:

```
[C#]

eventBrokerService.Register(
        this.OnEventPackageOpened_StatServer1, statServer1EndpointFilter);
```

Architecture and Design

The Message Broker Application Block is designed to make it easy for your applications to handle events in an efficient way.

Message Broker allows you to set up individual classes to handle specific events coming from Genesys servers. It receives all of the events from the servers you specify, and sends each one to the appropriate handler class. Message Broker is a high-performance way to hide the complexity of event-driven programming — so you can focus on other areas of your application.

Tip

Message Broker has been designed for use with the Protocol Manager Application Block. Protocol Manager is another high-performance component that makes it easy for your applications to connect to Genesys servers. You can find basic information on how to use the Protocol Manager Application Block in the article on Connecting to a Server Using the Protocol Manager Application Block.

The Message Broker Application Block Architecture

The Message Broker Application Block uses a service-based API that enables you to write individual methods that handle one or more events.

For example, you might want to handle every occurrence of `EventAgentLogin` with a specific dedicated method, while there might be other events that you wish to send to a common event-handling method. Message Broker allows you write these methods and register them with an event broker that manages them for you.

Message Filters

Message Broker uses *message filters* to identify specific messages, assign them to specified methods, and route them accordingly. These message filters are shown in greater detail in the figure below.

Design Patterns

This section gives an overview of the design patterns used in the Message Broker Application Block.

Publish/Subscribe Pattern

There are many occasions when one class (the subscriber) needs to be notified when something changes in another class (the publisher). Message Broker uses the Publish/Subscribe pattern to inform the client application when events arrive from the server.

Factory Method Pattern

It is common practice for a class to include constructors that enable clients of the class instantiate it. There are times, however, when a client may need to instantiate one of several different classes. In some of these situations, the client should not need to decide which class is being created. In this case, a Factory Method pattern is used. The Factory Method pattern lets a class developer define the interface for creating an object, while retaining control of which class to instantiate.

Event Handling Using the Message Broker Application Block

Important

The Message Broker Application Block is considered a legacy product as of release 8.1.1 due to changes to the default event-receiving mechanism. Documentation related to this application block is retained for backwards compatibility. For information about event handling *without* use of the deprecated Message Broker Application Block, refer to the Event Handling article.

Once you have connected to a server using the Protocol Manager Application Block, much of the work of your application will be to send messages to that server and then handle the events you receive from it.

Genesys recommends that you use the **Message Broker Application Block** for most of your event handling needs. This article shows how to send and receive simple synchronous events without using Message Broker and then discusses how to use Message Broker for asynchronous event handling.

Tip

It is important to determine whether your application needs to use synchronous or asynchronous messages. In general, you will probably use only one or the other type in your application. If you decide to use synchronous messages, you must make sure that your code handles all of the messages you receive from your servers. For example, if you send a `RequestReadObjects` message to Configuration Server, you will receive several `EventObjectsRead` messages, followed by an `EventObjectsSent` message. If your application does not handle all of these messages, it will not work properly.

The messages you send to a server are in the form of requests. For example, you may send a request to log in an agent or to gather statistics. You might also send a request to update a configuration object, or to shut down an application.

In each of these cases, the server will respond with an event message, as shown below.

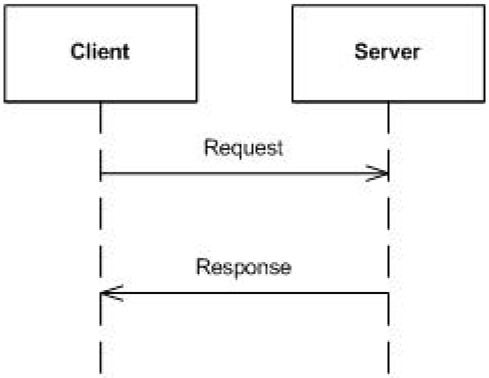

Some of the requests you send may best be handled with a synchronous response, while others may best be handled asynchronously. Let’s talk about synchronous requests first.

Java

Synchronous Requests

Sometimes you might want a synchronous response to your request. For example, if you are using the Open Media Platform SDK, you may want to log in an agent. To do this, you need to let the server know that you want to log in. And then you need to wait for confirmation that your login was successful.

The first thing you need to do is to create a login request, as shown here:

```
[Java]

RequestAgentLogin requestAgentLogin =
        RequestAgentLogin.create(
                tenantId,
                placeId,
                reason);
```

This version of `RequestAgentLogin.Create` specifies most of the information you will need in order to perform the login, but there is one more piece of data required. Here is how to add it:

```
[Java]

requestAgentLogin.setMediaList(mediaList);
```

Once you have created the request and set all required properties, you can make a synchronous request by using the request method of your ProtocolManagementService object, like this:

```
[Java]

Message response = null;
response = protocolManagementServiceImpl.getProtocol("Interaction_Server_App")
        .request(requestAgentLogin);
```

> **Tip**
> For information on how to use the ProtocolManagementServiceImpl class of the Protocol Manager Block to communicate with a Genesys server, see the article on Connecting to a Server.

There are two important things to understand when you use the `request` method:

- When you execute this method call, the calling thread will be blocked until it has received a response from the server.
- This method call will only return one message from the server. If the server returns subsequent messages in response to this request, you must process them separately. This can happen in the example of sending a `RequestReadObjects` message to Configuration Server, as mentioned at the beginning of this article.

The response from the server will come in the form of a `Message`. This is the interface implemented by all events in the Platform SDK. Some types of requests will be answered by an event that is specific to the request, while others may receive a more generic response of `EventAck`, which simply acknowledges that your request was successful. If a request fails, the server will send an `EventError`.

A successful `RequestAgentLogin` will receive an `EventAck`, while an unsuccessful one will receive an `EventError`. You can use a switch statement to test which response you received, as outlined here:

```
[Java]

switch(response.messageId())
{
        case EventAck.ID:
                OnEventAck(response);
        case EventError.ID:
                OnEventError(response);
        ...
}
```

Using Message Broker to Handle Asynchronous Requests

There are times when you need to receive asynchronous responses from a server.

First of all, some requests to a server can result in multiple events. For example, if you send a `RequestReadObjects` message, which is used to read objects from the Genesys Configuration Layer, Configuration Server may send more than one `EventObjectsRead` messages in response, depending on whether there is too much data to be handled by a single `EventObjectsRead`.

In other cases, events may be unsolicited. To continue with our example, once you have received all of the `EventObjectsRead` messages, Configuration Server will also send an `EventObjectsSent`, which confirms that it has completed your request.

To make an asynchronous request, you would use the send method of your `ProtocolManagementServiceImpl` class. For example, you might need to fetch information about some objects in the Genesys Configuration Layer. Here is how to set up a RequestReadObjects, followed by the send:

```
[Java]

KeyValueCollection filterKey = new KeyValueCollection();
filterKey.addObject("switch_dbid", 113);
filterKey.addObject("dn_type", CfgDNType.CFGExtension.asInteger());
RequestReadObjects requestReadObjects = RequestReadObjects.create(
        CfgObjectType.CfgDN.asInteger(), filterKey);
protocolManagementServiceImpl.getProtocol("Config_Server_App")
        .send(requestReadObjects);
```

This snippet is searching for all DNs that have a type of Extension and are associated with the switch that has a database ID of 113.

There are several ways to handle the response from the server, but Genesys recommends that you use the Message Broker Application Block, which is included with the Platform SDK. Message Broker allows you to set up individual classes to handle specific events. It receives the events from the servers you are working with, and sends them to the appropriate handler class. Message Broker is a high-performance way to hide the complexity of event-driven programming — so you can focus on other areas of your application.

To use the Message Broker Application Block, add the following .jar file to the classpath for your application:

- messagebrokerappblock.jar

This .jar file was precompiled using the default Application Block code, and can be located at: <Platform SDK Folder>\lib.

> ### Tip
> You can also view or modify the Message Broker Application Block source code. To do this, open the Message Broker Java source files that were installed with the Platform SDK. The Java source files for this project are located at: <Platform SDK Folder>\applicationblocks\messagebroker\src\java. If you make any changes to the project, you will have to run Ant (or use the build.bat file for this Application Block) to rebuild the .jar archive listed above. After you run Ant, add the resulting .jar to your classpath.

Now you can add the appropriate import statements to your source code. For example:

```
[Java]

import com.genesyslab.platform.applicationblocks.commons.broker.*;
```

In order to use the Message Broker Application Block, you need to create an EventBrokerService object to handle the events your application receives. Since you are using the Protocol Manager Application Block to connect to your servers, as shown in the section on Connecting to a Server, you should specify the ProtocolManagementServiceImpl object in the EventBrokerService constructor:

```
[Java]
```

```
EventBrokerService mEventBrokerService = new EventBrokerService(
        (MessageReceiverSupport) protocolManagementServiceImpl
                .getReceiver());
```

You also need to set up the appropriate filters for your event handlers and register the handlers with the `EventBrokerService`. This allows that service to determine which classes will be used for event-handling. Note that you should register these classes before you open the connection to the server. Otherwise, the server might send events before you are ready to handle them. The sample below shows how to filter on Message ID, which is an integer associated with a particular message:

```
[Java]

mEventBrokerService.register(new ConfObjectsReadHandler(),
        new MessageIdFilter(EventObjectsRead.ID));
mEventBrokerService.register(new ConfObjectsSentHandler(),
        new MessageIdFilter(EventObjectsSent.ID));
mEventBrokerService.register(new StatPackageInfoHandler(),
        new MessageIdFilter(EventPackageInfo.ID));
```

Once you have registered your event-handling classes, you can activate the `EventBrokerService` and open the connection to your server. In the following snippet, connections are being opened to both Configuration Server and Stat Server:

```
[Java]

mEventBrokerService.activate();

protocolManagementServiceImpl.getProtocol("Config_Server_App")
        .open();
protocolManagementServiceImpl.getProtocol("Stat_Server_App").open();
```

At this point, you are ready to set up classes to handle the events you have received from the server. Here is a simple class that handles the `EventObjectsRead` messages:

```
[Java]

class ConfObjectsReadHandler implements Action {

        public void handle(Message obj) {
                EventObjectsRead objectsRead = (EventObjectsRead) obj;
                // Add processing here...
        }
}
```

As mentioned earlier, once Configuration Server has sent all of the information you requested, it will let you know it has finished by sending an `EventObjectsSent` message. Note that this handler has a structure that is similar to the one for `EventObjectsRead`:

```
[Java]

class ConfObjectsSentHandler implements Action {

        public void handle(Message obj) {
                EventObjectsSent objectsSent = (EventObjectsSent) obj;
                // Add processing here...
        }
}
```

`Message Broker only routes non-null messages of the type you specify to your message`

```
handlers. For example, if you send a RequestReadObjects and no objects in the
Configuration Layer meet your filtering criteria, you will not receive an
EventObjectsRead. In that case, you will only receive an EventObjectsSent. Therefore,
you do not need to check for a null message in your EventObjectsRead handler.|2
```

The `EventPackageInfo` handler also has a similar structure, but in this case, we show how to print information about the statistics contained in the requested package:

```
[Java]

class StatPackageInfoHandler implements Action {

  public void handle(Message obj) {
    EventPackageInfo eventPackageInfo = (EventPackageInfo) obj;
    if (eventPackageInfo != null)
    {
      int statisticsCount = eventPackageInfo.getStatistics().getCount();
      StatisticsCollection statisticsCollection = eventPackageInfo.getStatistics();

      for (int i = 0; i < statisticsCount; i++)
      {
        Statistic statistic = statisticsCollection.getStatistic(i);

        System.out.println("\nStatistic Metric is: " +
             statistic.getMetric().toString());
        System.out.println("Statistic Object is: " +
            statistic.getObject());
        System.out.println("Statistic IntValue is: " +
            statistic.getIntValue());
        System.out.println("Statistic StringValue is: " +
            statistic.getStringValue());
        System.out.println("Statistic ObjectValue is: " +
            statistic.getObjectValue());
        System.out.println("Statistic ExtendedValue is: " +
            statistic.getExtendedValue());
         System.out.println("Statistic Tenant is: " +
            statistic.getObject().getTenant());
        System.out.println("Statistic Type is: " +
            statistic.getObject().getType());
        System.out.println("Statistic Id is: " +
            statistic.getObject().getId());
        System.out.println("Statistic TimeProfile is: " +
            statistic.getMetric().getTimeProfile());
         System.out.println("Statistic StatisticType is: " +
            statistic.getMetric().getStatisticType());
        System.out.println("Statistic TimeRange is: " +
            statistic.getMetric().getTimeRange());
        }
      }
    }
  }
```

Filtering Messages by Server

Each server in the Genesys environment makes use of a particular set of events that corresponds to the tasks of that server. For example, Configuration Server sends `EventObjectsRead` and `EventObjectsSent` messages, among others, while Stat Server's events include `EventPackageInfo` and `EventPackageOpened`. Although your applications can identify each of these events by name, it is more efficient to use the ID field associated with an event, which you specify as an `int`. You can do this by using a `MessageIdFilter`, as shown here:

```
[Java]

mEventBrokerService.register(new ConfEventErrorHandler(),
        new MessageIdFilter(EventError.ID));
```

However, the integer used for the Message ID of, say, a Configuration Server message, could be same as the integer used for a completely different message on another server. This could lead to problems if your application works with messages from more than one server. For example, if a multi-server application includes a handler that processes a specific type of message from the first server and that message has an ID of 12, any messages from the other servers that also have a Message ID of 12 will be sent by your `MessageIdFilter` to the same handler.

Fortunately, the Platform SDK allows you to filter messages on a server-by-server basis in addition to filtering on `MessageId`. Here is how to set up a Protocol Description object that allows you to specify that you want some of your handlers to work only with events that are coming from Configuration Server:

```
[Java]

ConfServerProtocol confServerProtocol = (ConfServerProtocol)
        protocolManagementServiceImpl.getProtocol("Config_Server_App");
ProtocolDescription configProtocolDescription = null;
if (confServerProtocol != null)
{
        configProtocolDescription =
                confServerProtocol.getProtocolDescription();
}
```

Once you have set up this Protocol Description, you can use it to indicate that you only want to process events associated with that server, in addition to specifying which event or events you want each handler to process:

```
[Java]

mEventBrokerService.register(new ConfEventErrorHandler(),
        new MessageIdFilter(configProtocolDescription, EventError.ID));
```

You are now ready to open the connection to Configuration Server:

```
[Java]

protocolManagementServiceImpl.
        getProtocol("Config_Server_App").open();
```

Using One Handler for Multiple Events

There may be times when you would like to use a single event handler for more than one event. In that case, you can create the handler and then register the appropriate events with it. For example, you might create a handler for both `EventObjectsRead` and `EventObjectsSent`:

```
[Java]

        class ConfEventHandler implements Action {
                ...
        }
```

You might use a `case` statement inside the handler, in order to process each event appropriately. In any case, once you have set up this handler, all you need to do is register both events with it, as shown here:

```
[Java]

mEventBrokerService.register(new ConfEventHandler(),
        new MessageIdFilter(configProtocolDescription, EventObjectsRead.ID));
mEventBrokerService.register(new ConfEventHandler(),
        new MessageIdFilter(configProtocolDescription, EventObjectsSent.ID));
```

These are the basics of how to use the Message Broker Application Block. For more information, see the Using the Message Broker Application Block article.

.NET

Synchronous Requests

Sometimes you might want a synchronous response to your request. For example, if you are using the Open Media Platform SDK, you may want to log in an agent. To do this, you need to let the server know that you want to log in. And then you need to wait for confirmation that your login was successful.

The first thing you need to do is to create a login request, as shown here:

```
[C#]

RequestAgentLogin requestAgentLogin =
        RequestAgentLogin.Create(
                tenantId,
                placeId,
                reason);
```

This version of `RequestAgentLogin.Create` specifies most of the information you will need in order to perform the login, but there is one more piece of data required. Here is how to add it:

```
[C#]

requestAgentLogin.MediaList = mediaList;
```

Once you have created the request and set all required properties, you can make a synchronous request by using the `Request` method of your `ProtocolManagementService` object, like this:

```
[C#]

IMessage response =
        protocolManagementService["InteractionServer"].
            Request(requestAgentLogin);
```

> **Tip**

> For information on how to use the `ProtocolManagementService` class of the Protocol Manager Application Block to communicate with a Genesys server, see the article on Connecting to a Server Using the Protocol Manager Application Block.

There are two important things to understand when you use the Request method:

- When you execute this method call, the calling thread will be blocked until it has received a response from the server.
- This method call will only return one message from the server. If the server returns subsequent messages in response to this request, you must process them separately. This can happen in the example of sending a `RequestReadObjects` message to Configuration Server, as mentioned at the beginning of this article.

The response from the server will come in the form of an `IMessage`. This is the interface implemented by all events in the Platform SDK. Some types of requests will be answered by an event that is specific to the request, while others may receive a more generic response of `EventAck`, which simply acknowledges that your request was successful. If a request fails, the server will send an `EventError`.

A successful `RequestAgentLogin` will receive an `EventAck`, while an unsuccessful one will receive an `EventError`. You can use a switch statement to test which response you received, as outlined here:

```
[C#]

switch(response.Id)
{
        case EventAck.MessageId:
                OnEventAck(response);
        case EventError.MessageId:
                OnEventError(response);
        ...
}
```

Using Message Broker to Handle Asynchronous Requests

There are times when you need to receive asynchronous responses from a server.

First of all, some requests to a server can result in multiple events. For example, if you send a `RequestReadObjects` message, which is used to read objects from the Genesys Configuration Layer, Configuration Server may send more than one `EventObjectsRead` messages in response, depending on whether there is too much data to be handled by a single `EventObjectsRead`.

In other cases, events may be unsolicited. To continue with our example, once you have received all of the `EventObjectsRead` messages, Configuration Server will also send an `EventObjectsSent`, which confirms that it has completed your request.

To make an asynchronous request, you would use the Send method of your `ProtocolManagementService` class. Here is how to set up a `RequestReadObjects`, followed by the `Send`:

```
[C#]

KeyValueCollection filterKey = new KeyValueCollection();
filterKey.Add("switch_dbid", 113);
filterKey.Add("dn_type", (int) CfgDNType.Extension);
RequestReadObjects requestReadObjects =
        RequestReadObjects.Create(
                (int) CfgObjectType.CFGDN,
                filterKey);
protocolManagementService["ConfigServer"].Send(requestReadObjects);
```

This snippet is searching for all DNs that have a type of `Extension` and are associated with the switch that has a database ID of 113.

There are several ways to handle the response from the server, but Genesys recommends that you use the Message Broker Application Block, which is included with the Platform SDK. Message Broker allows you to set up individual handlers for specific events. It receives the events from the servers you are working with, and sends them to the appropriate handler. Message Broker is a high-performance way to hide the complexity of event-driven programming — so you can focus on other areas of your application.

To use the Message Broker Application Block, open the Solution Explorer for your application project and add a reference to the following file:

- Genesyslab.Platform.ApplicationBlocks.Commons.Broker.dll

This dll file is precompiled using the default Application Block code, and can be located at: `<Platform SDK Folder>\Bin`.

> ### Tip
> You can also view or modify the Message Broker Application Block source code. To do this, open the Message Broker Visual Studio project that was installed with the Platform SDK. The solution file for this project is located at: `<Platform SDK Folder>\ApplicationBlocks\MessageBroker`. If you make any changes to the project, you will have to rebuild the .dll file listed above.

Once you have added the reference, you can add a using statement to your source code:

```
[C#]

using Genesyslab.Platform.ApplicationBlocks.Commons.Broker;
```

In order to use the Message Broker Application Block, you need to create an `EventBrokerService` object to handle the events your application receives. Declare this object with your other fields:

```
[C#]

EventBrokerService eventBrokerService;
```

Then you can set up the `EventBrokerService` to receive events from the Protocol Manager Application Block's `ProtocolManagementService` class, which you are using to connect to your servers, as shown in the section on Connecting to a Server:

```
[C#]

eventBrokerService = new EventBrokerService(protocolManagementService.Receiver);
```

Now you are ready to set up your event handlers.

Note that there are two ways to do this. In 7.5, when Message Broker was introduced, you needed to use attributes to filter the events you wanted processed by a particular handler. Starting in 7.6, you can still do it that way, but you can also set up your filters in the statement that registers an event handler with the Event Broker service, rather than using attributes that are associated with the handler itself. This new method may perform better than the old way, but we will show you how to use both.

Using Event Handlers Without Attributes

Let us start by setting up a couple of event handlers. First, here is a simple handler for the `EventError` message:

```
[C#]

private void OnConfEventError(IMessage theMessage)
{
        EventError eventError = theMessage as EventError;
        /// Add processing here...
}
```

And here is one for the `EventObjectsRead` message:

```
[C#]

private void OnConfEventObjectsRead(IMessage theMessage)
{
        EventObjectsRead objectsRead = theMessage as EventObjectsRead;
        /// Add processing here...
}
```

As mentioned earlier, once Configuration Server has sent all of the information you requested, it will let you know it has finished by sending an `EventObjectsSent` message. Here is a handler for that:

```
[C#]

private void OnConfEventObjectsSent(IMessage theMessage)
{
        EventObjectsSent objectsSent = theMessage as EventObjectsSent;
        /// Add processing here...
}
```

Now you can set up the appropriate filters for your event handlers and register the handlers with the `EventBrokerService`. This allows that service to determine which classes will be used for event-handling. Note that you should register these handlers before you open the connection to the server. Otherwise, the server might send events before you are ready to handle them. The sample below shows how to filter on Message ID, which is an integer associated with a particular message:

```
[C#]

eventBrokerService.Register(
        this.OnConfEventError,
        new MessageIdFilter(EventError.MessageId));
```

```
eventBrokerService.Register(
        this.OnConfEventObjectsRead,
        new MessageIdFilter(EventObjectsRead.MessageId));
eventBrokerService.Register(
        this.OnConfEventObjectsSent,
        new MessageIdFilter(EventObjectsSent.MessageId));
```

```
Message Broker only routes non-null messages of the type you specify to your message
handlers. For example, if you send a RequestReadObjects and no objects in the
Configuration Layer meet your filtering criteria, you will not receive an
EventObjectsRead. In that case, you will only receive an EventObjectsSent. Therefore,
you do not need to check for a null message in your EventObjectsRead handler.
```

Filtering Messages by Server

Each server in the Genesys environment makes use of a particular set of events that corresponds to the tasks of that server. For example, Configuration Server sends `EventObjectsRead` and `EventObjectsSent` messages, among others, while Stat Server's events include `EventPackageInfo` and `EventPackageOpened`. Although your applications can identify each of these events by name, it is more efficient to use the ID field associated with an event, which you specify as an int. You can do this by using a `MessageIdFilter`, as shown here:

```
[C#]

eventBrokerService.Register(this.OnConfEventError);
```

However, the integer used for the Message ID of, say, a Configuration Server message, could be same as the integer used for a completely different message on another server. This could lead to problems if your application works with messages from more than one server. For example, if a multi-server application includes a handler that processes a specific type of message from the first server and that message has an ID of 12, any messages from the other servers that also have a Message ID of 12 will be sent by your `MessageIdFilter` to the same handler.

Fortunately, the Platform SDK allows you to filter messages on a server-by-server basis in addition to filtering on `MessageId`. Here is how to set up a Protocol Description object that allows you to specify that you want some of your handlers to work only with events that are coming from Configuration Server:

```
[C#]

ConfServerProtocol confServerProtocol =
        protocolManagementService["Config_Server_App"]
                as ConfServerProtocol;
ProtocolDescription configProtocolDescription = null;
if (confServerProtocol != null)
{
        configProtocolDescription =
                confServerProtocol.ProtocolDescription;
}
```

Once you have set up this Protocol Description, you can use it to indicate that you only want to process events associated with that server, in addition to specifying which event or events you want each handler to process:

```
[C#]

eventBrokerService.Register(
```

```
        this.OnConfEventError,
                new MessageIdFilter(
                        configProtocolDescription,
                        EventError.MessageId));
eventBrokerService.Register(
        this.OnConfEventObjectsRead,
                new MessageIdFilter(
                         configProtocolDescription,
                        EventObjectsRead.MessageId));
eventBrokerService.Register(
        this.OnConfEventObjectsSent,
                new MessageIdFilter(
                        configProtocolDescription,
                        EventObjectsSent.MessageId));
```

You are now ready to open the connection to Configuration Server:

```
[C#]

protocolManagementService["Config_Server_App"].Open();
```

Using One Handler for Multiple Events

There may be times when you would like to use a single event handler for more than one event. In that case, you can create the handler and then register the appropriate events with it. For example, you might create a handler for both `EventObjectsRead` and `EventObjectsSent`:

```
[C#]

private void OnConfEvents (IMessage theMessage) {
        ...
}
```

You might use a case statement inside the handler, in order to process each event appropriately. In any case, once you have set up this handler, all you need to do is register both events with it, as shown here:

```
[C#]

eventBrokerService.Register(
        this.OnConfEvents,
                new MessageIdFilter(
                        configProtocolDescription,
                        EventObjectsRead.MessageId));
eventBrokerService.Register(
        this.OnConfEvents,
                new MessageIdFilter(
                        configProtocolDescription,
                        EventObjectsSent.MessageId));
```

Using Attributes with Your Event Handlers

As mentioned above, you can also use attributes to filter your event handlers. It is important to note that this may not perform as well as the method outlined above, but in case you would like to use attributes in your application, here is how to proceed.

When you use attributes, you have to specify the name of the protocol object you are using, and the name of the SDK it is part of, as shown here:

```
[C#]

private const string protocolName = "ConfServer";
private const string sdkName = "Configuration";
```

These values can be determined by accessing the `ProtocolDescription.ProtocolName` and `ProtocolDescription.SdkName` properties of your protocol object. They are also provided in the following table.

SDK	SdkName	Protocol Object	ProtocolName
Configuration Platform SDK	Configuration	ConfServerProtocol	ConfServer
Contacts Platform SDK	Contacts	UniversalContactServerProtocol	ContactServer
Management Platform SDK	Management	• LocalControlAgentProtocol • MessageServerProtocol • SolutionControlServerProtocol	• LocalControlAgent • MessageServer • SolutionControlServer
Open Media Platform SDK	OpenMedia	• InteractionServerProtocol • ExternalServiceProtocol	• InteractionServer • ExternalService
Outbound Contact Platform SDK	Outbound	OutboundServerProtocol	OutboundServer
Routing Platform SDK	Routing	• RoutingServerProtocol • UrsCustomProtocol	• RoutingServer • CustomServer
Statistics Platform SDK	Reporting	StatServerProtocol	StatServer
Voice Platform SDK	Voice	TServerProtocol	TServer
Web Media Platform SDK	WebMedia	• BasicChatProtocol • FlexChatProtocol • EmailProtocol • EspEmailProtocol • CallbackProtocol	• BasicChat • FlexChat • Email • EspEmail • Callback

Table 1: Platform SDK SdkName and ProtocolName Values

You also need to register the methods you will handle your events with. This allows the `EventBrokerService` to determine which methods will be used for event-handling. When registering for event handlers that use attributes, you only specify the name of the event-handling method. In this case, you need to handle three different events. Note that you should register these methods before you open the connection to the server, as shown here. Otherwise, the server might send events before you are ready to handle them:

```
[C#]
```

```
eventBrokerService.Register(this.OnConfEventObjectsRead);
eventBrokerService.Register(this.OnConfEventObjectsSent);
eventBrokerService.Register(this.OnConfEventError);
protocolManagementService["Config_Server_App"].Open();
```

At this point, you are ready to set up methods to handle the events you have received from the server. Here is a simple method that handles the `EventError` message:

```
[C#]

[MessageIdFilter(EventError.MessageId, ProtocolName = "ConfServer", SdkName =
"Configuration")]
private void OnConfEventError(IMessage theMessage)
{
        EventError eventError = theMessage as EventError;
        /// Add processing here...
}
```

Notice that there is a `MessageIdFilter` attribute right before the method body. This attribute indicates that all `EventError` messages for the Configuration Platform SDK's Configuration protocol will be handled by this method.

The attributes and methods for `EventObjectsRead` have a similar structure:

```
[C#]

[MessageIdFilter(EventObjectsRead.MessageId, ProtocolName = "ConfServer", SdkName =
"Configuration")]
private void OnConfEventObjectsRead(IMessage theMessage)
{
        EventObjectsRead objectsRead = theMessage as EventObjectsRead;
        /// Add processing here...
}
```

And so do the attributes and methods for `EventObjectsSent`:

```
 [C#]

[MessageIdFilter(EventObjectsSent.MessageId, ProtocolName = "ConfServer", SdkName =
"Configuration")]
private void OnConfEventObjectsSent(IMessage theMessage)
{
        //protocolManagementService["Config_Server_App"].Close();
        EventObjectsSent objectsSent = theMessage as EventObjectsSent;
        /// Add processing here...

}
```

If you want to process more than one event with a single handler, you can set up multiple attributes for that handler, like this:

```
[C#]

[MessageIdFilter(EventObjectsRead.MessageId, ProtocolName = "ConfServer", SdkName =
"Configuration")]
[MessageIdFilter(EventObjectsSent.MessageId, ProtocolName = "ConfServer", SdkName =
"Configuration")]
```

```
private void OnConfEvents (IMessage theMessage) {
        ...
}
```

These are the basics of how to use the Message Broker Application Block. For more information, see the Using the Message Broker Application Block article.

Using the Protocol Manager Application Block

> **Important**
> This application block is considered a legacy product starting with release 8.1.1. Documentation is provided for backwards compatibility, but new development should consider using the improved method of connecting to servers.

One of the two main functions of the Platform SDK is to enable your applications to establish and maintain connections with Genesys servers.

The Protocol Manager Application Block is a reusable production-quality component that provides unified management of server protocol objects. It takes care of opening and closing connections to many different servers, as well as reconfiguration of high availability connections. It has been designed using industry best practices and provided with source code so it can be used "as is," extended, or tailored if you need to. Please see the License Agreement for details.

For information on the other application blocks that ship with the Genesys SDKs, consult Introducing the Platform SDK.

Java

Installing the Protocol Manager Application Block

Before you install the Protocol Manager Application Block, it is important to review the software requirements.

Software Requirements

To work with the Protocol Manager Application Block, you must ensure that your system meets the software requirements established in the Genesys Supported Operating Environment Reference Guide, as well as meeting the following minimum software requirements:

- JDK 1.6 or higher

Building the Protocol Manager Application Block

To build the Protocol Manager Application Block:

1. Open the <Platform SDK Folder>\applicationblocks\protocolmanager folder.
2. Run either build.bat or build.sh, depending on your platform.

This will create the protocolmanagerappblock.jar file, located within the <Platform SDK Folder>\applicationblocks\protocolmanager\dist\lib directory.

Working with the Protocol Manager Application Block

You can find basic information on how to use the Protocol Manager Application Block in the article on Connecting to a Server Using the Protocol Manager Application Block.

Configuring ADDP

To enable ADDP, set the UseAddp property of your Configuration object to true. You can also set server and client timeout intervals, as shown here:

```
[Java]

statServerConfiguration.setUseAddp(true);
statServerConfiguration.setAddpServerTimeout(10);
statServerConfiguration.setAddpClientTimeout(10);
```

> **Tip**
> To avoid connection exceptions in the scenario where a client has configured ADDP but the server has not, "ADDP" is included as a default value for the "protocol" key in the configure() method of the ServerChannel class.

Configuring Warm Standby

Enable warm standby in your application by setting your Configuration object's FaultTolerance property to FaultToleranceMode.WarmStandby, as shown here. You can also configure the backup server's URI, the timeout interval, and the number of times your application will attempt to contact the primary server before switching to the backup:

```
[Java]

statServerConfiguration
                .setFaultTolerance(FaultToleranceMode.WarmStandby);
statServerConfiguration.setWarmStandbyTimeout(10);
statServerConfiguration.setWarmStandbyAttempts((short) 5);
try {
        statServerConfiguration.setWarmStandbyUri(new URI("tcp://"
                + statServerBackupHost
                + ":"
                + statServerBackupPort));
} catch (URISyntaxException e) {
```

```
        e.printStackTrace();
}
```

High-Performance Message Parsing

The Platform SDK exposes the protocols of supported Genesys servers as an API. This means you can write .NET and Java applications that communicate with these servers in their native protocols.

Every message you receive from a Genesys server is formatted in some way. Most Genesys servers use binary protocols, while some use XML-based protocols. When your application receives one of these messages, it parses the message and places it in the message queue for the appropriate protocol.

By default, the Platform SDK uses a single thread for all of this message parsing. Since this parsing can be time-consuming in certain cases, some applications may face serious performance issues. For example, some applications may receive lots of large binary-format messages, such as some of the statistics messages generated by Stat Server, while others might need to parse messages in non-binary formats, such as the XML format used to communicate with Genesys Multimedia (or e-Services) servers.

This section gives an example of how you can modify Protocol Manager to selectively enable multi-threaded parsing of incoming messages, in order to work around these kinds of performance issues. It is important to stress that you must take a careful look at which kind of multi-threading options to pursue in your applications, since your needs are specific to your situation.

Tip

Your application may also face other performance bottlenecks. For example, you may need more than one instance of the Message Broker Application Block if you handle large numbers of messages. For more information on how to configure Message Broker for high-performance situations, see the Message Broker Application Block Guide.

This example shows how to call `com.genesyslab.platform.commons.threading.DefaultInvoker`, which uses `SingleThreadInvoker` behind the scenes. As mentioned, you need to determine whether this is the right solution for your application.

The main thing to take from this example is how to set up an invoker interface, so that you can use another invoker if `DefaultInvoker` doesn't meet your needs. For example, Genesys also supplies `com.genesyslab.platform.commons.threading.SingleThreadInvoker`, which assigns a single dedicated thread to each protocol that enables it in your application. This may be useful in some cases where you have to parse XML messages.

The enhancement shown here will only require small changes to two of the classes in Protocol Manager, namely `ProtocolConfiguration` and `ProtocolFacility`.

To get started, let's declare a new multi-threaded parsing property in the `ProtocolConfiguration` class. In this example, the property is called `useMultiThreadedMessageParsing`. It is declared right after some ADDP and Warm Standby declarations:

```
[Java]

private boolean useAddp;
private FaultToleranceMode faultTolerance;
private Boolean useMultiThreadedMessageParsing;
```

Now you can code the getter and setter methods for the property itself, as shown here:

```
[Java]

public Boolean getUseMultiThreadedMessageParsing()
{
        return useMultiThreadedMessageParsing;
}

public void setUseMultiThreadedMessageParsing(Boolean value)
{
        useMultiThreadedMessageParsing = value;
}
```

Once you have made these changes, add an if statement to the `ApplyChannelConfiguration` method of the `ProtocolFacility` class so that your applications can selectively enable this property:

```
[Java]

private void applyChannelConfiguration(
        ProtocolConfiguration conf, ProtocolInstance instance)
{
        if (conf.getUri() != null)
        {
                instance.getProtocol().setEndpoint(
                        new Endpoint(conf.getName(), conf.getUri()));
        }

        if (conf.getUseMultiThreadedMessageParsing() != null &&
                        conf.getUseMultiThreadedMessageParsing().booleanValue())
        {
                instance.getProtocol().
                        setConnectionInvoker(DefaultInvoker.getSingletonInstance());
        }
...
```

Enabling `UseMultiThreadedMessageParsing` now calls `DefaultInvoker`.

To enable multi-threaded parsing, set the `useMultiThreadedMessageParsing` property of your Configuration object to true. Here is how to enable the new property for Stat Server messages:

```
[Java]

statServerConfiguration.setUseMultiThreadedMessageParsing(true);
```

Receiving Copies of Synchronous Server Messages

Most of the time, when you send a synchronous message to a server, you are satisfied to receive the response synchronously. But there can be situations where you want to receive a copy of the response asynchronously, as well. This section shows how to do that.

As in the previous section, this enhancement will only require small changes to the `ProtocolConfiguration` and `ProtocolFacility` classes.

To get started, let's declare a new `copyResponse` property in the `ProtocolConfiguration` class. You can put this declaration right after the `useMultiThreadedMessageParsing` declaration we created in the previous section:

```
[Java]

private boolean useAddp;
private FaultToleranceMode faultTolerance;
private Boolean useMultiThreadedMessageParsing;
private Boolean copyResponse;
```

Now you can code the getter and setter methods for the property itself, as shown here:

```
[Java]

public Boolean getCopyResponse()
{
        return copyResponse;
}

public void setCopyResponse(Boolean value)
{
        copyResponse = value;
}
```

It might be a good idea to let anyone using Protocol Manager know whether this property is enabled. One way to do this is to add it to the `toString` method in this class:

```
[Java]

public String toString()
{
        StringBuilder sb = new StringBuilder();
        .
        .
        .
        sb.append(MessageFormat.format(
                    "AddpClientTimeout: {0}\n", addpClientTimeout));
        sb.append(MessageFormat.format(
                    "AddpServerTimeout: {0}\n", addpServerTimeout));
        sb.append(MessageFormat.format(
                "CopyResponse: {0}\n", copyResponse));
        ...
```

Once you have made these changes, add an if statement to the `applyChannelConfiguration` method of the `ProtocolFacility` class so that your applications can selectively enable this property:

```
[Java]

private void applyChannelConfiguration(
        ProtocolConfiguration conf, ProtocolInstance instance)
{
        if (conf.getUri() != null)
        {
                instance.getProtocol().setEndpoint(
                        new Endpoint(conf.getName(), conf.getUri()));
        }
```

```
        if (conf.getCopyResponse() != null)
        {
                instance.getProtocol().setCopyResponse(
                        conf.getCopyResponse());
        }
...
```

To receive a copy of synchronous server messages, set the `CopyResponse` property of your Configuration object to true. Here is how to enable the new property for Stat Server messages:

[Java]

```
statServerConfiguration.setCopyResponse(true);
```

Supporting New Protocols

When the Platform SDK was first developed, it supported many, but not all, of the servers in the Genesys environment. As the SDK has matured, support has been added for more servers. As you might expect, a given version of the Protocol Manager Application Block only supports those servers that were supported by the Platform SDK at the time of its release. Since you may want to work with a server that is not currently supported by Protocol Manager, it can be helpful to know how add support for that server.

This section shows how the Protocol Manager Application Block supports the Stat Server Protocol. You can use it as a guide if you need to add support for other servers or protocols.

Adding support for the Stat Server Protocol involved three basic steps:

1. Create a new subclass of `ProtocolConfiguration` called `StatServerConfiguration`.
2. Create a new subclass of `ProtocolFacility` called `StatServerFacility`.
3. Add a statement to the `initialize` method of `ProtocolManagementServiceImpl` that associates `StatServerFacility` with `StatServerProtocol`.

The StatServerConfiguration Class

Here is the code for `StatServerConfiguration`:

[Java]

```
package com.genesyslab.platform.applicationblocks.commons.protocols;
import com.genesyslab.platform.reporting.protocol.StatServerProtocol;
import java.text.MessageFormat;

public final class StatServerConfiguration extends ProtocolConfiguration
{

    private String clientName;
    private Integer clientId;

    public StatServerConfiguration(String name)
    {
        super(name, StatServerProtocol.class);
    }
```

```
     public Integer getClientId()
    {
        return clientId;
    }

    public void setClientId(Integer clientId)
    {
        this.clientId = clientId;
    }

    public String getClientName()
    {
        return clientName;
    }

    public void setClientName(String clientName)
    {
        this.clientName = clientName;
    }

    public String toString()
    {
        StringBuilder sb = new StringBuilder(super.toString());

        sb.append(MessageFormat.format("ClientName: {0}\n", clientName));
        sb.append(MessageFormat.format("ClientId: {0}\n", this.clientId));

        return sb.toString();
    }
}
```

As you can see, this class imports the protocol object, but you will also need to use `MessageFormat` when we create the `toString()` method, so there must be an import statement for that class, as well:

```
[Java]

import com.genesyslab.platform.reporting.protocol.StatServerProtocol;
import java.text.MessageFormat;
```

Here are the class declaration and the field and constructor declarations. Stat Server requires client name and ID, so these must both be present in `StatServerConfiguration`:

```
[Java]

public final class StatServerConfiguration extends ProtocolConfiguration
{

    private String clientName;
    private Integer clientId;

    public StatServerConfiguration(String name)
    {
        super(name, StatServerProtocol.class);
    }
```

Here are the getter and setter methods for the client name and ID:

```
[Java]

public Integer getClientId()
```

```
{
        return clientId;
}

public void setClientId(Integer clientId)
{
        this.clientId = clientId;
}

public String getClientName()
{
        return clientName;
}

public void setClientName(String clientName)
{
        this.clientName = clientName;
}
```

And finally, the `toString()` method:

```
[Java]

public String toString()
{
        StringBuilder sb = new StringBuilder(super.toString());

        sb.append(MessageFormat.format("ClientName: {0}\n", clientName));
        sb.append(MessageFormat.format("ClientId: {0}\n", this.clientId));

        return sb.toString();
}
```

The StatServerFacility Class

Now we can take a look at the `StatServerFacility` class. Once again, we will start with the code for the entire class:

```
[Java]

package com.genesyslab.platform.applicationblocks.commons.protocols;

import com.genesyslab.platform.commons.protocol.Endpoint;
import com.genesyslab.platform.commons.protocol.Protocol;
import com.genesyslab.platform.reporting.protocol.StatServerProtocol;
import java.net.URI;

public final class StatServerFacility extends ProtocolFacility
{

    public void applyConfiguration(
                ProtocolInstance instance, ProtocolConfiguration conf)
    {
        super.applyConfiguration(instance, conf);
        StatServerConfiguration statConf = (StatServerConfiguration)conf;
        StatServerProtocol statProtocol =
                    (StatServerProtocol) instance.getProtocol();

/*
        if (statConf.getClientName() != null)
        {
```

```
            statProtocol.setClientName(statConf.getClientName());
        }
*/
        if (statConf.getClientId() != null)
        {
            statProtocol.setClientId(statConf.getClientId());
        }
    }

    public Protocol createProtocol(String name, URI uri)
    {
        return new StatServerProtocol(new Endpoint(name, uri));
    }
}
```

This class needs the following import statements:

```
[Java]

import com.genesyslab.platform.commons.protocol.Endpoint;
import com.genesyslab.platform.commons.protocol.Protocol;
import com.genesyslab.platform.reporting.protocol.StatServerProtocol;
import java.net.URI;
```

Here is how to declare the class:

```
[Java]

public final class StatServerFacility extends ProtocolFacility
```

There are two methods in this class. The first is `applyConfiguration`:

```
[Java]

public void applyConfiguration(
                        ProtocolInstance instance, ProtocolConfiguration conf)
{
        super.applyConfiguration(instance, conf);
        StatServerConfiguration statConf = (StatServerConfiguration)conf;
        StatServerProtocol statProtocol =
                                (StatServerProtocol) instance.getProtocol();

/*
        if (statConf.getClientName() != null)
        {
                statProtocol.setClientName(statConf.getClientName());
        }
*/
        if (statConf.getClientId() != null)
        {
                statProtocol.setClientId(statConf.getClientId());
        }
}
```

The second method is `createProtocol`:

```
[Java]

public Protocol createProtocol(String name, URI uri)
{
        return new StatServerProtocol(new Endpoint(name, uri));
}
```

Updating ProtocolManagementServiceImpl

To complete this enhancement, a single line of code was added to the initialize method of `ProtocolManagementServiceImpl`:

```
[Java]

private void Initialize()
{
        this.facilities.Add(typeof(ConfServerProtocol), new ConfServerFacility());
        this.facilities.Add(typeof(TServerProtocol), new TServerFacility());
        this.facilities.Add(typeof(InteractionServerProtocol), new
InteractionServerFacility());
        this.facilities.Add(typeof(StatServerProtocol), new StatServerFacility());
        this.facilities.Add(typeof(OutboundServerProtocol), new OutboundServerFacility());
        this.facilities.Add(typeof(LocalControlAgentProtocol), new LcaFacility());
        this.facilities.Add(typeof(SolutionControlServerProtocol), new ScsFacility());
        this.facilities.Add(typeof(MessageServerProtocol), new MessageServerFacility());
}
```

Architecture and Design

The Protocol Manager Application Block uses a service-based API. You can use this API to open and close your connection with Genesys servers and to dynamically reconfigure the parameters for a given protocol. Protocol Manager also includes built-in warm standby capabilities.

Protocol Manager uses a `ServerConfiguration` object to describe each server it manages.

.NET

Installing the Protocol Manager Application Block

Before you install the Protocol Manager Application Block, it is important to review the software requirements and the structure of the software distribution.

Software Requirements

To work with the Protocol Manager Application Block, you must ensure that your system meets the software requirements established in the Genesys Genesys Supported Operating Environment Reference Guide.

Building the Protocol Manager Application Block

The Platform SDK distribution includes a Genesyslab.Platform.ApplicationBlocks.Commons.Protocols.dll file that you can use as is. This file is located in the bin directory at the root level of the Platform SDK directory. To build your own copy of this application block, follow the instructions below:

1. Open the <Platform SDK Folder>\ApplicationBlocks\ProtocolManager folder.
2. Double-click ProtocolManager.sln.
3. Build the solution.

Working with the Protocol Manager Application Block

You can find basic information on how to use the Protocol Manager Application Block in the article on Connecting to a Server Using the Protocol Manager Application Block at the beginning of this guide.

Configuring ADDP

To enable ADDP, set the `UseAddp` property of your Configuration object to true. You can also set server and client timeout intervals, as shown here:

```
[C#]

statServerConfiguration.UseAddp = true;
statServerConfiguration.AddpServerTimeout = 10;
statServerConfiguration.AddpClientTimeout = 10;
```

Configuring Warm Standby

Hot standby is not designed to handle situations where both the primary and backup servers are down. It is also not designed to connect to your backup server if the primary server was down when you initiated your connection. However, in cases like these, warm standby will attempt to connect. In fact, warm standby will keep trying one server and then the other, until it does connect. Because of this, you will probably want to enable warm standby in your applications, even if you are already using hot standby.

You can enable warm standby in your application by setting your Configuration object's `FaultTolerance` property to `FaultToleranceMode.WarmStandby`, as shown here. You can also configure the backup server's URI, the timeout interval, and the number of times your application will attempt to contact the primary server before switching to the backup:

```
[C#]

statServerConfiguration.FaultTolerance = FaultToleranceMode.WarmStandby;
statServerConfiguration.WarmStandbyTimeout = 5000;
statServerConfiguration.WarmStandbyAttempts = 5;
statServerConfiguration.WarmStandbyUri = statServerBackupUri;
```

High-Performance Message Parsing

The Platform SDK exposes the protocols of supported Genesys servers as an API. This means you can write .NET and Java applications that communicate with these servers in their native protocols.

Every message you receive from a Genesys server is formatted in some way. Most Genesys servers use binary protocols, while some use XML-based protocols. When your application receives one of these messages, it parses the message and places it in the message queue for the appropriate protocol.

By default, the Platform SDK uses a single thread for all of this message parsing. Since this parsing can be time-consuming in certain cases, some applications may face serious performance issues. For example, some applications may receive lots of large binary-format messages, such as some of the statistics messages generated by Stat Server, while others might need to parse messages in non-binary formats, such as the XML format used to communicate with Genesys Multimedia (or e-Services) servers.

This section gives an example of how you can modify Protocol Manager to selectively enable multi-threaded parsing of incoming messages, in order to work around these kinds of performance issues. It is important to stress that you must take a careful look at which kind of multi-threading options to pursue in your applications, since your needs are specific to your situation.

> **Tip**
>
> Your application may also face other performance bottlenecks. For example, you may need more than one instance of the Message Broker Application Block if you handle large numbers of messages. For more information on how to configure Message Broker for high-performance situations, see the Using the Message Broker Application Block.

This example shows how to call `Genesyslab.Platform.Commons.Threading.DefaultInvoker`, which uses the .NET thread pool for your message parsing needs. As mentioned, you need to determine whether this is the right solution for your application, since, for example, the .NET thread pool may be heavily used for other tasks.

The main thing to take from this example is how to set up an invoker interface, so that you can use another invoker if `DefaultInvoker` doesn't meet your needs. For example, Genesys also supplies `Genesyslab.Platform.Commons.Threading.SingleThreadInvoker`, which assigns a single dedicated thread to each protocol that enables it in your application. This may be useful in some cases where you have to parse XML messages.

The enhancement shown here will only require small changes to two of the classes in Protocol Manager, namely `ProtocolConfiguration` and `ProtocolFacility`.

To get started, let's declare a new multi-threaded parsing property in the `ProtocolConfiguration` class. In this example, the property is called `useMultiThreadedMessageParsing`. It is nullable and is declared right after some ADDP and Warm Standby declarations:

```
[C#]

private bool? useAddp;
```

```
private FaultToleranceMode? faultTolerance;
private string addpTrace;
private bool? useMultiThreadedMessageParsing;
```

Now you can code the property itself, as shown here:

```
[C#]

public bool? UseMultiThreadedMessageParsing
{
        get { return this.useMultiThreadedMessageParsing; }
        set { this.useMultiThreadedMessageParsing = value; }
}
```

Once you have made these changes, add an if statement to the `ApplyChannelConfiguration` method of the `ProtocolFacility` class so that your applications can selectively enable this property:

```
[C#]

private void ApplyChannelConfiguration(ProtocolInstance entry, ProtocolConfiguration conf)
{
        if( conf.Uri != null )
        {
                entry.Protocol.Endpoint = new Endpoint(conf.Name, conf.Uri);
        }

        if (conf.UseMultiThreadedMessageParsing != null &&
conf.UseMultiThreadedMessageParsing.Value)
                {
                        entry.Protocol.SetConnectionInvoker(DefaultInvoker.InvokerSingleton);
        }
...
```

Enabling `UseMultiThreadedMessageParsing` now calls `DefaultInvoker`, which uses the .NET thread pool, as mentioned above.

To enable multi-threaded parsing, set the `UseMultiThreadedMessageParsing` property of your Configuration object to true. Here is how to enable the new property for Stat Server messages:

```
[C#]

statServerConfiguration.UseMultiThreadedMessageParsing = true;
```

Receiving Copies of Synchronous Server Messages

Most of the time, when you send a synchronous message to a server, you are satisfied to receive the response synchronously. But there can be situations where you want to receive a copy of the response asynchronously, as well. This section shows how to do that.

As in the previous section, this enhancement will only require small changes to the `ProtocolConfiguration` and `ProtocolFacility` classes.

To get started, let's declare a new `copyResponse` property in the `ProtocolConfiguration` class. You can put this declaration right after the `useMultiThreadedMessageParsing` declaration we created in the previous section:

```
[C#]

private bool? useAddp;
private FaultToleranceMode? faultTolerance;
private string addpTrace;
private bool? useMultiThreadedMessageParsing;
private bool? copyResponse;
```

Now you can code the property itself, as shown here:

```
[C#]

public bool? CopyResponse
{
        get { return this.copyResponse; }
        set { this.copyResponse = value; }
}
```

It might be a good idea to let anyone using Protocol Manager know whether this property is enabled. One way to do this is to add it to the `ToString` method overrides in this class:

```
[C#]

public override string ToString()
{
        StringBuilder sb = new StringBuilder();
        .
        .
        .
        sb.AppendFormat("AddpClientTimeout: {0}\n", this.addpClientTimeout.ToString());
        sb.AppendFormat("AddpServerTimeout: {0}\n", this.addpServerTimeout.ToString());
        sb.AppendFormat("CopyResponse: {0}\n", this.copyResponse.ToString());
        ...
```

Once you have made these changes, add an if statement to the `ApplyChannelConfiguration` method of the `ProtocolFacility` class so that your applications can selectively enable this property:

```
[C#]

private void ApplyChannelConfiguration(ProtocolInstance entry, ProtocolConfiguration conf)
{
        if( conf.Uri != null )
        {
                entry.Protocol.Endpoint = new Endpoint(conf.Name, conf.Uri);
        }

        if (conf.CopyResponse != null)
{
        entry.Protocol.CopyResponse = conf.CopyResponse.Value;
}
...
```

To receive a copy of synchronous server messages, set the `CopyResponse` property of your Configuration object to true. Here is how to enable the new property for Stat Server messages:

```
[C#]

statServerConfiguration.CopyResponse = true;
```

Supporting New Protocols

When the Platform SDK was first developed, it supported many, but not all, of the servers in the Genesys environment. As the SDK has matured, support has been added for more servers. As you might expect, a given version of the Protocol Manager Application Block only supports those servers that were supported by the Platform SDK at the time of its release. Since you may want to work with a server that is not currently supported by Protocol Manager, it can be helpful to know how add support for that server.

For example, early versions of Protocol Manager were developed before the Platform SDK supported Universal Contact Server (UCS). This section shows how to add UCS support to the Protocol Manager Application Block. You can also use these instructions as a guide if you need to add support for other servers.

This enhancement involves three basic steps:

- Create a new subclass of `ProtocolConfiguration`. We will call this class `ContactServerConfiguration`.
- Create a new subclass of `ProtocolFacility` called `ContactServerFacility`.
- Add a statement to the `Initialize` method of `ProtocolManagementService` that associates the new `ContactServerFacility` class with `UniversalContactServerProtocol`.

Creating a ContactServerConfiguration Class

We will use the `StatServerConfiguration` class as a template for the new `ContactServerConfiguration` class. Here is the code for `StatServerConfiguration`:

```
[C#]

using System;
using System.Text;

using Genesyslab.Platform.Reporting.Protocols;

namespace Genesyslab.Platform.ApplicationBlocks.Commons.Protocols
{
    public sealed class StatServerConfiguration : ProtocolConfiguration
    {
        #region Fields

        private string clientName;
        private int? clientId;

        #endregion Fields

        public StatServerConfiguration(string name)
            : base(name, typeof(StatServerProtocol))
        {
        }

        #region Properties

                public string ClientName
                {
                        get { return this.clientName; }
```

```
                        set { this.clientName = value; }
                }

                public int? ClientId
                {
                        get { return this.clientId; }
                        set { this.clientId = value; }
                }

                #endregion Properties

        public override string ToString()
        {
            StringBuilder sb = new StringBuilder();
            sb.Append(base.ToString());

            sb.AppendFormat("ClientName: {0}\n", this.clientName);
            sb.AppendFormat("ClientId: {0}\n", this.clientId.ToString());

            return sb.ToString();
        }
    }
}
```

To get started, make a copy of `StatServerConfiguration.cs` and call it `ContactServerConfiguration.cs`. Rename the Platform SDK using statement and the class name, as shown here:

```
[C#]

using System;
using System.Text;
using Genesyslab.Platform.Contacts.Protocols;

namespace Genesyslab.Platform.ApplicationBlocks.Commons.Protocols
{
    public sealed class ContactServerConfiguration : ProtocolConfiguration
    {
    ...
```

The connection parameters required by Stat Server are different from those used by UCS. Instead of `clientName` and `clientId`, UCS requires `applicationName`. Like `clientName`, `applicationName` is of type string. One fairly simple way to modify this class is to delete all references to `clientId` and rename the references to `clientName` to `applicationName`. Make sure to retain the capitalization in the property name, which should become `ApplicationName`.

```
[C#]

using System;
using System.Text;
using Genesyslab.Platform.Contacts.Protocols;

namespace Genesyslab.Platform.ApplicationBlocks.Commons.Protocols
{
    public sealed class ContactServerConfiguration : ProtocolConfiguration
    {
        #region Fields

        private string applicationName;
        private int? clientId;
```

```
        #endregion Fields

        ...

        #region Properties

                public string ApplicationName
                {
                        get { return this.applicationName; }
                        set { this.applicationName = value; }
                }

                public int? ClientId
                {
                        get { return this.clientId; }
                        set { this.clientId = value; }
                }

                #endregion Properties

        public override string ToString()
        {
            StringBuilder sb = new StringBuilder();
            sb.Append(base.ToString());

            sb.AppendFormat("applicationName: {0}\n", this.applicationName);
            sb.AppendFormat("ClientId: {0}\n", this.clientId.ToString());

            return sb.ToString();
        }
    }
}
```

The constructor also needs to be renamed. This code:

```
[C#]

public StatServerConfiguration(string name)
        : base(name, typeof(StatServerProtocol))
{
}
```

should be replaced with this:

```
[C#]

public ContactServerConfiguration(string name)
        : base(name, typeof(UniversalContactServerProtocol))
{
}
```

When you have made all of these changes, your new class should look like this:

```
[C#]

using System;
using System.Text;
using Genesyslab.Platform.Contacts.Protocols;

namespace Genesyslab.Platform.ApplicationBlocks.Commons.Protocols
{
    public sealed class ContactServerConfiguration : ProtocolConfiguration
```

```
    {
        #region Fields

        private string applicationName;

        #endregion Fields

        public ContactServerConfiguration(string name)
            : base(name, typeof(UniversalContactServerProtocol))
        {
        }
        #region Properties

        public string ApplicationName
        {
            get { return this.applicationName; }
            set { this.applicationName = value; }
        }

        #endregion Properties

        public override string ToString()
        {
            StringBuilder sb = new StringBuilder();
            sb.Append(base.ToString());

            sb.AppendFormat("ApplicationName: {0}\n", this.applicationName);

            return sb.ToString();
        }

    }
}
```

Creating a ContactServerFacility Class

The next step is to create a copy of `StatServerFacility.cs` and name it `ContactServerFacility.cs`. Here is what the `StatServerFacility` class looks like:

```
[C#]

using System;
using System.Text;

using Genesyslab.Platform.Commons.Collections;
using Genesyslab.Platform.Commons.Protocols;
using Genesyslab.Platform.Reporting.Protocols;
using Genesyslab.Platform.Commons.Logging;

namespace Genesyslab.Platform.ApplicationBlocks.Commons.Protocols
{
    internal sealed class StatServerFacility : ProtocolFacility
    {
        public override void ApplyConfiguration(ProtocolInstance entry, ProtocolConfiguration
conf, ILogger logger)
        {
            base.ApplyConfiguration(entry, conf, logger);

            StatServerConfiguration statConf = (StatServerConfiguration)conf;
            StatServerProtocol statProtocol = (StatServerProtocol)entry.Protocol;
```

```
            if (statConf.ClientName != null)
            {
                statProtocol.ClientName = statProtocol.ClientName;
            }
            if (statConf.ClientId != null)
            {
                statProtocol.ClientId = statConf.ClientId.Value;
            }
        }

        public override ClientChannel CreateProtocol(string name, Uri uri)
        {
            return new StatServerProtocol(new Endpoint(name, uri));
        }
    }
}
```

Start by renaming the `using` statement and the class name:

```
[C#]

using System;
using Genesyslab.Platform.Commons.Logging;
using Genesyslab.Platform.Commons.Protocols;
using Genesyslab.Platform.Contacts.Protocols;

namespace Genesyslab.Platform.ApplicationBlocks.Commons.Protocols
{
    internal sealed class ContactServerFacility : ProtocolFacility
    ...
```

Rename `statConf` and `statProtocol`, giving them the correct configuration and protocol types:

```
[C#]

ContactServerConfiguration ucsConf = (ContactServerConfiguration)conf;
UniversalContactServerProtocol ucsProtocol =
        (UniversalContactServerProtocol)entry.Protocol;
```

And delete the references to `ClientId`:

```
[C#]

if (statConf.ClientId != null)
{
        statProtocol.ClientId = statConf.ClientId.Value;
```

Now you can rename `ClientName` to `ApplicationName`:

```
[C#]

if (ucsConf.ApplicationName != null)
{
        ucsProtocol.ApplicationName = ucsConf.ApplicationName;
}
```

When you are finished, you will have a new class that looks like this:

```
[C#]

using System;
```

```
using Genesyslab.Platform.Commons.Logging;
using Genesyslab.Platform.Commons.Protocols;
using Genesyslab.Platform.Contacts.Protocols;

namespace Genesyslab.Platform.ApplicationBlocks.Commons.Protocols
{
    internal sealed class ContactServerFacility : ProtocolFacility
    {
        public override void ApplyConfiguration(ProtocolInstance entry, ProtocolConfiguration
conf, ILogger logger)
        {
            base.ApplyConfiguration(entry, conf, logger);

            ContactServerConfiguration ucsConf = (ContactServerConfiguration)conf;
            UniversalContactServerProtocol ucsProtocol =
(UniversalContactServerProtocol)entry.Protocol;

            if (ucsConf.ApplicationName != null)
            {
                ucsProtocol.ApplicationName = ucsConf.ApplicationName;
            }
        }

        public override ClientChannel CreateProtocol(string name, Uri uri)
        {
            return new UniversalContactServerProtocol(new Endpoint(name, uri));
        }
    }
}
```

Updating ProtocolManagementService

To complete this enhancement, add a single line of code to the `Initialize` method of `ProtocolManagementService`:

```
[C#]

private void Initialize()
{
        this.facilities.Add(typeof(ConfServerProtocol), new ConfServerFacility());
        this.facilities.Add(typeof(TServerProtocol), new TServerFacility());
        this.facilities.Add(typeof(InteractionServerProtocol), new
InteractionServerFacility());
        this.facilities.Add(typeof(StatServerProtocol), new StatServerFacility());
        this.facilities.Add(typeof(OutboundServerProtocol), new OutboundServerFacility());
        this.facilities.Add(typeof(LocalControlAgentProtocol), new LcaFacility());
        this.facilities.Add(typeof(SolutionControlServerProtocol), new ScsFacility());
        this.facilities.Add(typeof(MessageServerProtocol), new MessageServerFacility());
        this.facilities.Add(typeof(UniversalContactServerProtocol), new
ContactServerFacility());
}
```

Your copy of Protocol Manager now works with Universal Contact Server!

Architecture and Design

The Protocol Manager Application Block uses a service-based API. You can use this API to open and close your connection with Genesys servers and to dynamically reconfigure the parameters for a given protocol. Protocol Manager also includes built-in warm standby capabilities.

Protocol Manager uses a `ServerConfiguration` object to describe each server it manages. The figure below gives examples of the structure of some of these objects.

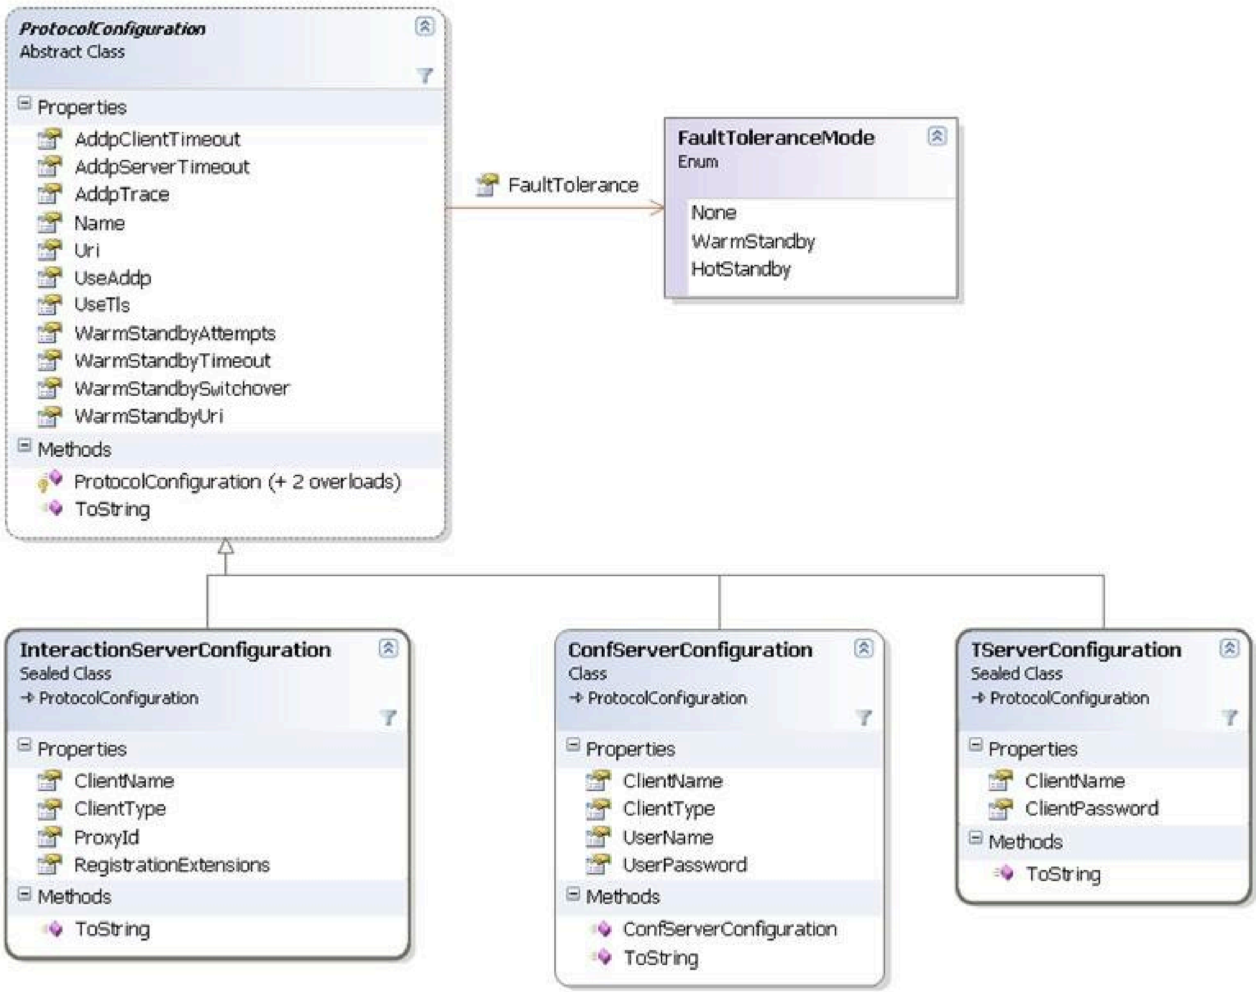

Tip

Any protocol can be reconfigured dynamically.

Connecting to a Server Using the Protocol Manager Application Block

Important

The Protocol Manager Application Block is considered a legacy product as of release 8.1.1 due to improvements in the configuration of core protocol classes. Documentation related to this application block is retained for backwards compatibility. For information about connecting to Genesys servers *without* use of the Protocol Manager Application Block, refer to the Connecting to a Server article.

The applications you write with the Platform SDK will need to communicate with one or more Genesys servers. So the first thing you need to do is create a connection with these servers. Genesys recommends that you use the **Protocol Manager Application Block** to do this. Protocol Manager is designed for high-performance communication with Genesys servers. It also includes built-in support for warm standby.

Once you have connected to a server, you will be sending and receiving messages to and from this server. The next article shows how to use the Message Broker Application Block for efficient event handling using the Message Broker Application Block.

Tip

Protocol Manager may not support all of the servers you need to use in your application. For information about how to update Protocol Manager to communicate with these servers, see the Using the Protocol Manager Application Block article.

Java

To use the Protocol Manager Application Block, add the following file to your classpath:

- `protocolmanagerappblock.jar`

This jar file was precompiled using the default Application Block code, and can be located at: `<Platform SDK Folder>\lib`.

Tip

> You can also view or modify the Protocol Manager Application Block source code. To do this, open the Protocol Manager Java source files that were installed with the Platform SDK. The Java source files for this project are located at: `<Platform SDK Folder>\applicationblocks\protocolmanager\src\java`. If you make any changes to the project, you will have to run Ant (or use the `build.bat` file for this Application Block) to rebuild the jar archive listed above. After you run Ant, add the resulting jar to your classpath.

Now you can add `import` statements to your source code. For example:

```
[Java]

import com.genesyslab.platform.applicationblocks.commons.protocols.*;
import com.genesyslab.platform.applicationblocks.warmstandby.*;
```

You will also have to add additional JAR archives to your classpath and add `import` statements to your project for each specific protocol you are working with. The steps are not explicitly described here because the archives and classes required will vary depending on which SDKs and protocols you plan to use.

In order to use the Protocol Manager, you need to create a `ProtocolManagementServiceImpl` object. This object manages all of your server connections. Declare it with your other fields:

```
[Java]

ProtocolManagementServiceImpl protocolManagementServiceImpl;
```

Then you can initialize the service object inside the appropriate method body:

```
[Java]

protocolManagementServiceImpl =
        new ProtocolManagementServiceImpl();
```

You are now ready to create an object that will be used to specify how to communicate with the server. For example, if you are working with Configuration Server, you will set up a `ConfServerConfiguration` object:

[Java]

ConfServerConfiguration confServerConfiguration = new
ConfServerConfiguration("Config_Server_App");

Note that you have to provide a string when you create the `ConfServerConfiguration` object. This string should be unique for each protocol used in your application. It might be a good idea to use the name of the server's application object from the configuration layer, which guarantees uniqueness as well as clearly identifying which server you are communicating with.

After setting up the `ConfServerConfiguration` object, you need to specify the URI of the Configuration Server you want to communicate with, as well as a few other necessary pieces of information:

```
[Java]
```

```
try {
        confServerConfiguration.setUri(
                new URI("tcp://" + confServerHost + ":" + confServerPort));
} catch (URISyntaxException e) {
        e.printStackTrace();
}
confServerConfiguration.setClientApplicationType(CfgAppType.CFGSCE);
confServerConfiguration.setClientName(clientName);
confServerConfiguration.setUserName(userName);
confServerConfiguration.setUserPassword(password);
```

At this point, you can register your `ConfServerConfiguration` object with Protocol Manager:

```
[Java]

protocolManagementServiceImpl.register(confServerConfiguration);
```

Now you can tell Protocol Manager to open the connection to your server:

```
[Java]

try {
        protocolManagementServiceImpl.getProtocol("Config_Server_App").open();
} catch (ProtocolException e) {
        e.printStackTrace();
} catch (IllegalStateException e) {
        e.printStackTrace();
} catch (InterruptedException e) {
        e.printStackTrace();
}
```

You may want to set up a connection to more than one server. To do that, you could repeat the steps outlined above. Here is an example of how you might do that in order to add a connection to Stat Server:

```
[Java]

StatServerConfiguration statServerConfiguration = new StatServerConfiguration(
                "Stat_Server_App");
try {
        statServerConfiguration.setUri(new URI(statServerUri));
} catch (URISyntaxException e) {
        // TODO Auto-generated catch block
        e.printStackTrace();
}
protocolManagementServiceImpl.register(statServerConfiguration);
.
.
.
// Add this line to the try block for the Configuration Server open()
protocolManagementServiceImpl.getProtocol("Stat_Server_App").open();
```

In some cases, you may want to use the `beginOpen()` method instead of using the `open()` method. `beginOpen()` will open all of your connections with a single method call. However, unlike `open()`, `beginOpen()` is asynchronous. This means you will need to make sure you have received the `onChannelOpened` event before you send any messages. Otherwise, you might be trying to use a connection that does not yet exist.

In order to use `beginOpen()`, you need to implement the `ChannelListener` interface:

```
[Java]

import com.genesyslab.platform.commons.protocol.ChannelListener;
.
.
.
public class YourApplication
        implements ChannelListener, ...
```

You will also need to add a channel listener after you register your `ServerConfiguration` objects:

```
[Java]

protocolManagementServiceImpl.register(confServerConfiguration);
protocolManagementServiceImpl.register(statServerConfiguration);
.
.
.
protocolManagementServiceImpl.addChannelListener(this);
```

Now you can add a method to handle the `OnChannelOpened` event:

```
[Java]

public void onChannelOpened(EventObject event) {
        if ( event.getSource() instanceof ClientChannel )        {
                ClientChannel channel = (ClientChannel)event.getSource();

                if ( channel instanceof ConfServerProtocol ) {
                        // Work with Configuration Server messages...
                }
                else if ( channel instanceof StatServerProtocol ) {
                        // Work with Stat Server messages...
                }
        }
}
```

Having done that, you can remove these lines from the `try` block:

```
[Java]

protocolManagementServiceImpl.getProtocol("Config_Server_App").open();
protocolManagementServiceImpl.getProtocol("Stat_Server_App").open();
```

And replace them with this one:

```
[Java]

protocolManagementServiceImpl.beginOpen();
```

However, if you want to issue an asynchronous open for a specific protocol, you can invoke `beginOpen` for that protocol, like this:

```
[Java]

protocolManagementServiceImpl.getProtocol("Config_Server_App").beginOpen();
protocolManagementServiceImpl.getProtocol("Stat_Server_App").beginOpen();
```

Tip

When using the `beginOpen()` method, make sure that your code waits for the `onChannelOpened` event to fire before attempting to send or receive messages.

Once you have opened your connection, you can send and receive messages, as shown in the article on Event Handling. But before getting to that, please note that when you have finished communicating with your servers, you should close the connection, like this:

```
[Java]

protocolManagementServiceImpl.beginClose();
```

Or like this:

```
[Java]

protocolManagementServiceImpl.getProtocol("Config_Server_App")
        .close();
protocolManagementServiceImpl.getProtocol("Stat_Server_App")
        .close();
```

Or like this:

```
[Java]

protocolManagementServiceImpl.getProtocol("Config_Server_App")
                .beginClose();
protocolManagementServiceImpl.getProtocol("Stat_Server_App")
                .beginClose();
```

This introduction has only covered the most basic features of the Protocol Manager Application Block. Consult the Protocol Manager Application Block Guide for more information on how to use Protocol Manager, including the following topics:

- Configuring ADDP
- Configuring Warm Standby
- High-Performance Message Parsing
- Supporting New Protocols

To learn how to send and receive messages, go to the article on Event Handling Using the Message Broker Application Block.

.NET

To use the Protocol Manager Application Block, open the Solution Explorer for your application project and add references to the following files:

- `Genesyslab.Platform.ApplicationBlocks.Commons.Protocols.dll`

- `Genesyslab.Platform.ApplicationBlocks.WarmStandby.dll`

These dll files are precompiled using the default Application Block code, and can be located at: `<Platform SDK Folder>\Bin`.

> **Tip**
> You can also view or modify the Protocol Manager Application Block source code. To do this, open the Protocol Manager Visual Studio project that was installed with the Platform SDK. The solution file for this project is located at: `<Platform SDK Folder>\ApplicationBlocks\ProtocolManager`. If you make any changes to the project, you will have to rebuild the two .dll files listed above.

Once you have added the references, you can add using statements to your source code:

```
[C#]

using Genesyslab.Platform.ApplicationBlocks.Commons.Protocols;
using Genesyslab.Platform.ApplicationBlocks.WarmStandby;
```

You will also have to reference additional libraries and add using statements to your project for each specific protocol you are working with. The steps are not explicitly described here because the files and namespaces required will vary depending on which SDKs and protocols you plan to use.

In order to use the Protocol Manager, you now need to create a `ProtocolManagementService` object. This object manages all of your server connections. Declare it with your other fields:

```
[C#]

ProtocolManagementService protocolManagementService;
```

Then you can initialize the service object inside the appropriate method body:

```
[C#]

protocolManagementService =
        new ProtocolManagementService();
```

You are now ready to create an object that will be used to specify how to communicate with the server. For example, if you are working with Configuration Server, you will set up a `ConfServerConfiguration` object:

```
[C#]

ConfServerConfiguration confServerConfiguration =
        new ConfServerConfiguration("Config_Server_App");
```

Note that you have to provide a string when you create the `ConfServerConfiguration` object. This string should be unique for each protocol used in your application. It might be a good idea to use the name of the server's application object from the configuration layer, which guarantees uniqueness as well as clearly identifying which server you are communicating with.

After setting up the `ConfServerConfiguration` object, you need to specify the URI of the

Configuration Server you want to communicate with, as well as a few other necessary pieces of information:

```
[C#]

confServerConfiguration.Uri =
        new Uri("tcp://" + confServerHost + ":" + confServerPort);
confServerConfiguration.ClientApplicationType = CfgAppType.CFGSCE;
confServerConfiguration.ClientName = clientName;
confServerConfiguration.UserName = userName;
confServerConfiguration.UserPassword = password;
```

At this point, you can register your `ConfServerConfiguration` object with Protocol Manager:

```
[C#]

protocolManagementService.Register(confServerConfiguration);
```

Now you can tell Protocol Manager to open the connection to your server:

```
[C#]

protocolManagementService["Config_Server_App"].Open();
```

You may want to set up a connection to more than one server. To do that, you could repeat the steps outlined above. Here is an example of how you might do that in order to add a connection to Stat Server:

```
[C#]

StatServerConfiguration statServerConfiguration = new
StatServerConfiguration("Stat_Server_App");
statServerConfiguration.Uri = statServerUri;
protocolManagementService.Register(statServerConfiguration);
.
.
.
protocolManagementService["Stat_Server_App"].Open();
```

In some cases, you may want to use the `BeginOpen()` method instead of using the `Open()` method. `BeginOpen()` will open all of your connections with a single method call. However, unlike `Open()`, `BeginOpen()` is asynchronous. This means you will need to make sure you have received the `Opened` event before you send any messages. Otherwise, you might be trying to use a connection that does not yet exist.

Once you have set up an event handler for the `Opened` event, you can remove these lines from your code:

```
[C#]

protocolManagementService["Config_Server_App"].Open();
protocolManagementService["Stat_Server_App"].Open();
```

And replace them with this one:

```
[C#]

protocolManagementService.BeginOpen();
```

However, if you want to issue an asynchronous open for a specific protocol, you can invoke BeginOpen for that protocol, like this:

```
[C#]

protocolManagementService["Config_Server_App"].BeginOpen();
protocolManagementService["Stat_Server_App"].BeginOpen();
```

Tip

When using the BeginOpen() method, make sure that your code waits for the Opened event to fire before attempting to send or receive messages.

Once you have opened your connection, you can send and receive messages, as shown in the article on Event Handling Using the Message Broker Application Block. But before getting to that, please note that when you have finished communicating with your servers, you should close the connection, like this:

```
[C#]

protocolManagementService.BeginClose();
```

Or like this:

```
[C#]

protocolManagementService["Config_Server_App"].Close();
protocolManagementService["Stat_Server_App"].Close();
```

Or like this:

```
[C#]

protocolManagementService["Config_Server_App"].BeginClose();
protocolManagementService["Stat_Server_App"].BeginClose();
```

This introduction has only covered the most basic features of the Protocol Manager Application Block. Consult Using the Protocol Manager Application Block for more information on how to use Protocol Manager, including the following topics:

- Configuring ADDP
- Configuring Warm Standby
- High-Performance Message Parsing
- Supporting New Protocols

To learn how to send and receive messages, go to the article on Event Handling Using the Message Broker Application Block.

Transport Layer Security (TLS) Support

Platform SDK now supports Transport Layer Security (TLS). This section contains two sample configurations, but it is important to understand your environment and its unique requirements before using this new support. You should refer to the appropriate server manual to configure TLS on your server. You should also refer to Part 3 of the Genesys 8.0 Security Deployment Guide, "Communications Integrity—Transport Layer Security".

The first sample configuration shows a situation where the client application specifies the name of a server-based certificate:

```
[C#]

SolutionControlServerProtocol scsProtocol
        = new SolutionControlServerProtocol(myEndpoint);
KeyValueCollection kvCollection = new KeyValueCollection();
kvCollection[CommonConnection.TlsKey] = 1;
kvCollection[CommonConnection.CertificateNameKey] = "name";
KeyValueConfiguration kvConfig = new KeyValueConfiguration(kvCollection);
scsProtocol.Configure(kvConfig);
```

In this sample configuration, "name" is the name of the certificate, which is located in the certificate store on the server and used in the TLS configuration of the port/application/server in the Genesys Configuration Layer.

The second sample configuration shows a client application using its own client certificate to authenticate on the server:

```
[C#]

SolutionControlServerProtocol scsProtocol
        = new SolutionControlServerProtocol(myEndpoint);
KeyValueCollection kvCollection = new KeyValueCollection();
kvCollection[CommonConnection.TlsKey] = 1;
kvCollection[CommonConnection.CertificateKey] = @"c:\directory\certificate.p12";
kvCollection[CommonConnection.CertificatePwdKey] = "password";
KeyValueConfiguration kvConfig = new KeyValueConfiguration(kvCollection);
scsProtocol.Configure(kvConfig);
```

In this sample configuration, `CommonConnection.CertificateKey` is the path to the certificate file located on the client machine, while `CommonConnection.CertificatePwdKey` is the password which will be used to open the certificate file, if it is password protected.

Explicitly Choosing a Netty or Mina Connection Layer

Problem: A user wants to explicitly specify which connection layer Platform SDK should use, instead of accepting the default one.

Resolution: Set the `com.genesyslab.platform.commons.connection.factory.class` system property (defined by the `ConnectionManager.CONN_FACTORY_KEY` constant) to the fully qualified name of connection factory class to use:

- `com.genesyslab.platform.commons.connection.impl.netty.NettyConnectionFactory` (or `NettyConnectionFactory.class.getName()`)
- `com.genesyslab.platform.commons.connection.impl.mina.MinaConnectionFactory` (or `MinaConnectionFactory.class.getName()`)

```
$> java -Dcom.genesyslab.platform.commons.connection.factory.class=
      com.genesyslab.platform.commons.connection.impl.netty.NettyConnectionFactory
      -jar your_application.jar
```

This property can be set either from command line using `-D` switch or from code, as shown in the sample below.

Please note that Platform SDK looks up connection factory each time a connection is created, allowing the user to use different connection layers for different connections.

```
package test;

import com.genesyslab.platform.commons.collections.KeyValueCollection;
import com.genesyslab.platform.commons.connection.Connection;
import com.genesyslab.platform.commons.connection.ConnectionManager;
import com.genesyslab.platform.commons.connection.configuration.KeyValueConfiguration;
import com.genesyslab.platform.commons.connection.impl.netty.NettyConnectionFactory;
import com.genesyslab.platform.commons.connection.impl.mina.MinaConnectionFactory;
import com.genesyslab.platform.commons.protocol.Endpoint;
import com.genesyslab.platform.commons.protocol.ServerChannel;
import com.genesyslab.platform.management.protocol.solutioncontrolserver.runtime.channel.
   SolutionControlServerInternalProtocolFactory;

public class TlsServer {
  private ServerChannel channel = null;

  public void start() throws ProtocolException, InterruptedException {
    // System.setProperty(ConnectionManager.CONN_FACTORY_KEY,
    // NettyConnectionFactory.class.getName());
    System.setProperty(ConnectionManager.CONN_FACTORY_KEY,
    MinaConnectionFactory.class.getName());
    KeyValueCollection kvc = new KeyValueCollection();
    kvc.addString(Connection.TLS_KEY, "1"); // Turn on TLS
    // Java keystore format, use "keytool" utility from JDK to generate one
    kvc.addString(Connection.SSL_KEYSTORE_PATH_KEY, "D:\\Test\\keystore");
    kvc.addString(Connection.SSL_KEYSTORE_PASS, "password");
    Endpoint endpoint = new Endpoint("TlsServer", "localhost", 5500);
    channel = new ServerChannel(endpoint, new SolutionControlServerInternalProtocolFactory());
```

```
    channel.configure(new KeyValueConfiguration(kvc));
    channel.open();
  }
  // And so on...
}
```

Platform SDK Resources

This page describes additional resources located on this site and the Genesys Support site.

Related Documentation

Depending on what type of development you are doing with the Platform SDKs, the following resources may be useful for providing background information about your Genesys environment.

Genesys Events and Models Reference Manual

Use with: *T-Server, Interaction Server*

Download: Genesys Events and Models Reference Manual

If you are working with T-Server or Interaction Server, you should download and start reading the *Genesys Events and Models Reference Manual* right away. This document provides you with a large collection of two different types of important information, organized into two separate sections.

- **Part 1: Genesys Events** is the events portion of this document. The information in this part is wide-ranging, and includes everything from the names and descriptions of events, to the attributes that go with these events, to the definitions of event sub-states.
- **Part 2: Genesys Interaction Models** is the models portion of this document. It contains a selected list of call and interaction models. This information is also wide ranging. Based on the history of how this information has been presented in the past in various documents, model details may differ from chapter to chapter.

In both parts of this document, chapters are organized according to the type of event or model being described. So, for example, both parts one and two have specific chapters on voice-based issues that center on T-Library's generation of events and how calls are routed in a contact center.

Framework Stat Server User's Guide

Use with: *Stat Server*

See the Stat Server page to download the User's Guide by version.

Reporting Technical Reference

Use with: *Stat Server*

Download: Reporting Technical Reference 8.0 Overview

Download: Reporting Technical Reference Guide for the Genesys 7.2 Release

Multimedia Open Media Interaction Models Reference Manual

Use with: *Interaction Server*

Download: Multimedia 7.5 Open Media Interaction Models Reference Manual

Code Samples

The documentation for the Platform SDK includes a number of code samples. These samples are for illustrative purpose only:

- Complex Platform SDK 7.6 .NET Code Sample
- Configuration Platform SDK 7.6 Java (with Message Broker) Code Sample
- Configuration Platform SDK 7.6 .NET Code Sample
- Open Media Platform SDK 7.6 Java (Client) Code Sample
- Open Media Platform SDK 7.6 .NET (Client) Code Sample
- Open Media Platform SDK 7.6 Java (Server) Code Sample
- Open Media Platform SDK 7.6 .NET (Server) Code Sample
- Statistics Platform SDK 7.6 Java Code Sample
- Statistics Platform SDK 7.6 .NET Code Sample
- Voice Platform SDK 7.6 Java Code Sample
- Voice Platform SDK 7.6 .NET Code Sample