


AMI in C-Sharp with AsyncResult



 The AMI mapping using the AsyncResult API is deprecated and provided only for backward compatibility. New applications should use the [Task API](#).

On this page:

- [Basic Asynchronous API in C#](#)
 - [Asynchronous Proxy Methods in C#](#)
 - [Asynchronous Exception Semantics in C#](#)
- [AsyncResult Interface in C#](#)
- [Polling for Completion in C#](#)
- [Generic Completion Callbacks in C#](#)
 - [Using Cookies for Generic Completion Callbacks in C#](#)
- [Type-Safe Completion Callbacks in C#](#)
 - [Using Cookies for Type-Safe Completion Callbacks in C#](#)
- [Asynchronous Oneway Invocations in C#](#)
- [Flow Control in C#](#)
- [Asynchronous Batch Requests in C#](#)
- [Concurrency Semantics for AMI in C#](#)

Basic Asynchronous API in C#

Consider the following simple Slice definition:

Slice

```
module Demo
{
    interface Employees
    {
        string getName(int number);
    }
}
```

[Back to Top ^](#)

Asynchronous Proxy Methods in C#


Besides the synchronous proxy methods, `slice2cs` generates the following asynchronous proxy methods:

C#

```
public interface EmployeesPrx : Ice.ObjectPrx
{
    Ice.AsyncResult<Demo.Callback_Employees_getName>
    begin_getName(int number);

    Ice.AsyncResult<Demo.Callback_Employees_getName>
    begin_getName(int number,
        _System.Collections.Generic.Dictionary<string, string> ctx__);

    string end_getName(Ice.AsyncResult r__);
}
```

 Two additional overloads of `begin_getName` are generated for use with [generic completion callbacks](#).

As you can see, the single `getName` operation results in `begin_getName` and `end_getName` methods. (The `begin_` method is overloaded so you can pass a [per-invocation context](#).)

- The `begin_getName` method sends (or queues) an invocation of `getName`. This method does not block the calling thread.
- The `end_getName` method collects the result of the asynchronous invocation. If, at the time the calling thread calls `end_getName`, the result is not yet available, the calling thread blocks until the invocation completes. Otherwise, if the invocation completed some time before the call to `end_getName`, the method returns immediately with the result.

A client could call these methods as follows:

C#

```
EmployeesPrx e = ...;
Ice.AsyncResult r = e.begin_getName(99);

// Continue to do other things here...

string name = e.end_getName(r);
```

Because `begin_getName` does not block, the calling thread can do other things while the operation is in progress.

Note that `begin_getName` returns a value of type `Ice.AsyncResult`. (The class derives from `System.IAsyncResult`.) This value contains the state that the Ice run time requires to keep track of the asynchronous invocation. You must pass the `AsyncResult` that is returned by the `begin_` method to the corresponding `end_` method.

The `begin_` method has one parameter for each in-parameter of the corresponding Slice operation. Similarly, the `end_` method has one out-parameter for each out-parameter of the corresponding Slice operation (plus the `AsyncResult` parameter). For example, consider the following operation:

Slice

```
double op(int inp1, string inp2, out bool outp1, out long outp2);
```

The `begin_op` and `end_op` methods have the following signature:

C#

```
Ice.AsyncResult<Demo.Callback_Employees_op> begin_op(int inp1, string inp2);

double end_op(out bool outp1, out long outp2, Ice.AsyncResult r__);
```

[Back to Top](#) ^

Asynchronous Exception Semantics in C#

If an invocation raises an exception, the exception is thrown by the `end_` method, even if the actual error condition for the exception was encountered during the `begin_` method ("on the way out"). The advantage of this behavior is that all exception handling is located with the code that calls the `end_` method (instead of being present twice, once where the `begin_` method is called, and again where the `end_` method is called).

There is one exception to the above rule: if you destroy the communicator and then make an asynchronous invocation, the `begin_` method throws `CommunicatorDestroyedException`. This is necessary because, once the run time is finalized, it can no longer throw an exception from the `end_` method.

The only other exception that is thrown by the `begin_` and `end_` methods is `System.ArgumentException`. This exception indicates that you have used the API incorrectly. For example, the `begin_` method throws this exception if you call an operation that has a return value or out-parameters on a oneway proxy. Similarly, the `end_` method throws this exception if you use a different proxy to call the `end_` method than the proxy you used to call the `begin_` method, or if the `AsyncResult` you pass to the `end_` method was obtained by calling the `begin_` method for a different operation.

[Back to Top](#) ^

AsyncResult Interface in C#

The `AsyncResult` that is returned by the `begin_` method encapsulates the state of the asynchronous invocation:

C#

```
public interface AsyncResult : System.IAsyncResult
{
    void cancel();

    Ice.Communicator getCommunicator();
    Ice.Connection getConnection();
    ObjectPrx getProxy();
    string getOperation();
    object AsyncState { get; }

    bool IsCompleted { get; }
    void waitForCompleted();

    bool isSent();
    void waitForSent();

    void throwLocalException();

    bool sentSynchronously();

    AsyncResult whenSent(Ice.AsyncCallback cb);
    AsyncResult whenSent(Ice.SentCallback cb);
    AsyncResult whenCompleted(Ice.ExceptionCallback ex);
}

public interface AsyncResult<T> : AsyncResult
{
    AsyncResult<T> whenCompleted(T cb, Ice.ExceptionCallback excb);

    new AsyncResult<T> whenCompleted(Ice.ExceptionCallback excb);
    new AsyncResult<T> whenSent(Ice.SentCallback cb);
}
```

The methods and properties have the following semantics:

- `void cancel()`
This method prevents a queued invocation from being sent or, if the invocation has already been sent, ignores a reply if the server sends one. `cancel` is a local operation and has no effect on the server. A canceled invocation is considered to be completed, meaning `IsCompleted` returns true, and the result of the invocation is an `Ice.InvocationCanceledException`.
- `Communicator getCommunicator()`
This method returns the communicator that sent the invocation.
- `Connection getConnection()`
This method returns the connection that was used for the invocation. Note that, for typical asynchronous proxy invocations, this method returns a nil value because the possibility of automatic retries means the connection that is currently in use could change unexpectedly. The `getConnection` method only returns a non-nil value when the `AsyncResult` object is obtained by calling `begin_flushBatchRequests` on a `Connection` object.
- `ObjectPrx getProxy()`
This method returns the proxy that was used to call the `begin_` method, or nil if the `AsyncResult` object was not obtained via an asynchronous proxy invocation.
- `string getOperation()`
This method returns the name of the operation.
- `object AsyncState { get; }`
This property stores an object that you can use to pass shared state from the `begin_` to the `end_` method.
- `bool IsCompleted { get; }`
This property is true if, at the time it is called, the result of an invocation is available, indicating that a call to the `end_` method will not block the caller. Otherwise, if the result is not yet available, the method returns false.
- `void waitForCompleted()`
This method blocks the caller until the result of an invocation becomes available.
- `bool isSent()`
When you call the `begin_` method, the Ice run time attempts to write the corresponding request to the client-side transport. If the transport cannot

accept the request, the Ice run time queues the request for later transmission. `isSent` returns true if, at the time it is called, the request has been written to the local transport (whether it was initially queued or not). Otherwise, if the request is still queued or an exception occurred before the request could be sent, `isSent` returns false.

- `void waitForSent()`
This method blocks the calling thread until a request has been written to the client-side transport, or an exception occurs. After `waitForSent` returns, `isSent` returns true if the request was successfully written to the client-side transport, or false if an exception occurred. In the case of a failure, you can call the corresponding `end_` method or `throwLocalException` to obtain the exception.
- `void throwLocalException()`
This method throws the local exception that caused the invocation to fail. If no exception has occurred yet, `throwLocalException` does nothing.
- `bool sentSynchronously()`
This method returns true if a request was written to the client-side transport without first being queued. If the request was initially queued, `sentSynchronously` returns false (independent of whether the request is still in the queue or has since been written to the client-side transport).
- `AsyncResult whenSent(Ice.SentCallback cb)`
`AsyncResult<T> whenSent(Ice.SentCallback cb)`
`AsyncResult whenCompleted(Ice.ExceptionCallback ex)`
`AsyncResult<T> whenCompleted(T cb, Ice.ExceptionCallback excb)`
`AsyncResult<T> whenCompleted(Ice.ExceptionCallback excb)`
These methods allow you to specify callback methods that are called by the Ice run time. The `whenSent` methods set a callback that triggers when an asynchronous invocation is written to the client-side transport. The `whenCompleted` methods set a callback that triggers when an asynchronous invocation completes (also see [Generic Completion Callbacks in C#](#)).

[Back to Top ^](#)

Polling for Completion in C#

The `AsyncResult` methods allow you to poll for call completion. Polling is useful in a variety of cases. As an example, consider the following simple interface to transfer files from client to server:

Slice

```
interface FileTransfer
{
    void send(int offset, ByteSeq bytes);
};
```

The client repeatedly calls `send` to send a chunk of the file, indicating at which offset in the file the chunk belongs. A naïve way to transmit a file would be along the following lines:

C#

```
FileHandle file = open(...);
FileTransferPrx ft = ...;
const int chunkSize = ...;
int offset = 0;
while (!file.eof()) {
    byte[] bs;
    bs = file.read(chunkSize); // Read a chunk
    ft.send(offset, bs);        // Send the chunk
    offset += bs.Length;
}
```

This works, but not very well: because the client makes synchronous calls, it writes each chunk on the wire and then waits for the server to receive the data, process it, and return a reply before writing the next chunk. This means that both client and server spend much of their time doing nothing — the client does nothing while the server processes the data, and the server does nothing while it waits for the client to send the next chunk.

Using asynchronous calls, we can improve on this considerably:

C#

```
FileHandle file = open(...);
FileTransferPrx ft = ...;
const int chunkSize = ...;
int offset = 0;

LinkedList<Ice.AsyncResult> results = new LinkedList<Ice.AsyncResult>();
const int numRequests = 5;

while(!file.eof())
{
    byte[] bs;
    bs = file.read(chunkSize);

    // Send up to numRequests + 1 chunks asynchronously.
    Ice.AsyncResult r = ft.begin_send(offset, bs);
    offset += bs.Length;

    // Wait until this request has been passed to the transport.
    r.waitForSent();
    results.AddLast(r);

    // Once there are more than numRequests, wait for the least
    // recent one to complete.
    while(results.Count > numRequests)
    {
        Ice.AsyncResult r = results.First;
        results.RemoveFirst();
        r.waitForCompleted();
    }
}

// Wait for any remaining requests to complete.
while(results.Count > 0)
{
    Ice.AsyncResult r = results.First;
    results.RemoveFirst();
    r.waitForCompleted();
}
```

With this code, the client sends up to `numRequests + 1` chunks before it waits for the least recent one of these requests to complete. In other words, the client sends the next request without waiting for the preceding request to complete, up to the limit set by `numRequests`. In effect, this allows the client to "keep the pipe to the server full of data": the client keeps sending data, so both client and server continuously do work.

Obviously, the correct chunk size and value of `numRequests` depend on the bandwidth of the network as well as the amount of time taken by the server to process each request. However, with a little testing, you can quickly zoom in on the point where making the requests larger or queuing more requests no longer improves performance. With this technique, you can realize the full bandwidth of the link to within a percent or two of the theoretical bandwidth limit of a native socket connection.

[Back to Top ^](#)

Generic Completion Callbacks in C#

The `begin_` method is overloaded to allow you to provide completion callbacks. Here are the corresponding methods for the `getName` operation:

C#

```
Ice.AsyncResult begin_getName(  
    int number,  
    Ice.AsyncCallback cb__,  
    object cookie__);  
  
Ice.AsyncResult begin_getName(  
    int number,  
    _System.Collections.Generic.Dictionary<string, string> ctx__,  
    Ice.AsyncCallback cb__,  
    object cookie__);
```

The second version of `begin_getName` lets you override the default context. (We discuss the purpose of the `cookie` parameter in the next section.) Following the in-parameters, the `begin_` method accepts a parameter of type `Ice.AsyncCallback`, which is a delegate for a callback method. The Ice run time invokes the callback method when an asynchronous operation completes. Your callback method must have `void` return type and accept a single parameter of type `AsyncResult`, for example:

C#

```
private class MyCallback  
{  
    public void finished(Ice.AsyncResult r)  
    {  
        EmployeesPrx e = (EmployeesPrx)r.getProxy();  
        try  
        {  
            string name = e.end_getName(r);  
            System.Console.WriteLine("Name is: " + name);  
        }  
        catch(Ice.Exception ex)  
        {  
            System.Console.Err.WriteLine("Exception is: " + ex);  
        }  
    }  
}
```

The implementation of your callback method must call the `end_` method. The proxy for the call is available via the `getProxy` method on the `AsyncResult` that is passed by the Ice run time. The return type of `getProxy` is `Ice.ObjectPrx`, so you must down-cast the proxy to its correct type.

Your callback method should catch and handle any exceptions that may be thrown by the `end_` method. If you allow an exception to escape from the callback method, the Ice run time produces a log entry by default and ignores the exception. (You can disable the log message by setting the property `Ice.Warn.AMICallback` to zero.)

To inform the Ice run time that you want to receive a callback for the completion of the asynchronous call, you pass a delegate for your callback method to the `begin_` method:

C#

```
EmployeesPrx e = ...;  
  
MyCallback cb = new MyCallback();  
Ice.AsyncCallback del = new Ice.AsyncCallback(cb.finished);  
  
e.begin_getName(99, del, null);
```

The trailing `null` argument specifies a cookie, which we will discuss shortly.

You can avoid explicit instantiation of the delegate and, more tersely, write:

C#

```
EmployeesPrx e = ...;

MyCallback cb = new MyCallback();
e.begin_getName(99, cb.finished, null);
```

In fact, there's often no need to define a callback class at all. An anonymous method is one alternative:

C#

```
EmployeesPrx e = ...;

e.begin_getName(99,
    delegate(Ice.AsyncResult r)
    {
        EmployeesPrx e = (EmployeesPrx)r.getProxy();
        try
        {
            string name = e.end_getName(r);
            System.Console.WriteLine("Name is: " + name);
        }
        catch(Ice.Exception ex)
        {
            System.Console.Err.WriteLine("Exception is: " + ex);
        }
    }, null);
```

A lambda is another alternative:

C#

```
EmployeesPrx e = ...;

e.begin_getName(99,
    r => {
        EmployeesPrx e = (EmployeesPrx)r.getProxy();
        try
        {
            string name = e.end_getName(r);
            System.Console.WriteLine("Name is: " + name);
        }
        catch(Ice.Exception ex)
        {
            System.Console.Err.WriteLine("Exception is: " + ex);
        }
    }, null);
```

[Back to Top ^](#)

Using Cookies for Generic Completion Callbacks in C#

It is common for the `end_` method to require access to some state that is established by the code that calls the `begin_` method. As an example, consider an application that asynchronously starts a number of operations and, as each operation completes, needs to update different user interface elements with the results. In this case, the `begin_` method knows which user interface element should receive the update, and the `end_` method needs access to that element.

The API allows you to pass such state by providing a cookie. A cookie is any class instance; the class can contain whatever data you want to pass, as well as any methods you may want to add to manipulate that data.

Here is an example implementation that stores a `Widget`. (We assume that this class provides whatever methods are needed by the `end_` method to update the display.) When you call the `begin_` method, you pass the appropriate cookie instance to inform the `end_` method how to update the display:

C#

```
// Invoke the getName operation with different widget cookies.
MyCallback cb = ...;
e.begin_getName(99, cb.finished, widget1);
e.begin_getName(24, cb.finished, widget2);
```

The `end_` method can retrieve the cookie from the `AsyncResult` by reading the `AsyncState` property. For this example, we assume that widgets have a `writeString` method that updates the relevant UI element:

C#

```
public void finished(Ice.AsyncResult r)
{
    EmployeesPrx e = (EmployeesPrx)r.getProxy();
    Widget widget = (Widget)r.AsyncState;
    try
    {
        string name = e.end_getName(r);
        widget.writeString(name);
    }
    catch(Ice.Exception ex)
    {
        handleException(ex);
    }
}
```

The cookie provides a simple and effective way for you to pass state between the point where an operation is invoked and the point where its results are processed. Moreover, if you have a number of operations that share common state, you can pass the same cookie instance to multiple invocations.

[Back to Top ^](#)

Type-Safe Completion Callbacks in C#

The [generic callback API](#) is not entirely type-safe:

- You must down-cast the return value of `getProxy` to the correct proxy type before you can call the `end_` method.
- You must call the correct `end_` method to match the operation called by the `begin_` method.
- You must remember to catch exceptions when you call the `end_` method; if you forget to do this, you will not know that the operation failed.

`slice2cs` generates an additional type-safe API that takes care of these chores for you. To use type-safe callbacks, you supply delegates for two callback methods:

- a success callback that is called if the operation succeeds
- a failure callback that is called if the operation raises an exception

Here is a callback class for an invocation of the `getName` operation:

C#

```
public class MyCallback
{
    public void getNameCB(string name)
    {
        System.Console.WriteLine("Name is: " + name);
    }

    public void failureCB(Ice.Exception ex)
    {
        System.Console.Err.WriteLine("Exception is: " + ex);
    }
}
```


The callback methods can have any name you prefer and must have `void` return type. The failure callback always has a single parameter of type `Ice.Exception`. The success callback parameters depend on the operation signature. If the operation has non-`void` return type, the first parameter of the success callback is the return value. The return value (if any) is followed by a parameter for each out-parameter of the corresponding Slice operation, in the order of declaration.

At the calling end, you call the `begin_` method as follows:

C#

```
MyCallback cb = new MyCallback();

e.begin_getName(99).whenCompleted(cb.getNameCB, cb.failureCB);
```

Note the `whenCompleted` method on the `AsyncResult` that is returned by the `begin_` method. This method establishes the link between the `begin_` method and the callbacks that are called by the Ice run time by setting the delegates for the success and failure methods.

It is legal to pass a null delegate for the success or failure methods. For the success callback, this is legal only for operations that have `void` return type and no out-parameters. This is useful if you do not care when the operation completes but want to know if the call failed. If you pass a null exception delegate, the Ice run time will ignore any exception that is raised by the invocation.

Defining a callback class as we've shown above is only necessary in practice when the callback has additional state to manage. In many cases, there's no need for a callback class and the delegates can be defined inline:

C#

```
e.begin_getName(99).whenCompleted(
    delegate(string name)
    {
        System.Console.WriteLine("Name is: " + name);
    },
    delegate(Ice.Exception ex)
    {
        System.Console.Err.WriteLine("Exception is: " + ex);
    });
```

Using lambda functions makes the callbacks even more compact:

C#

```
e.begin_getName(99).whenCompleted(
    name => {
        System.Console.WriteLine("Name is: " + name);
    },
    ex => {
        System.Console.Err.WriteLine("Exception is: " + ex);
    });
```

[Back to Top ^](#)

Using Cookies for Type-Safe Completion Callbacks in C#

The type-safe API does not support cookies. If you want to pass state from the `begin_` method to the `end_` method, you must use the [generic API](#) or, alternatively, place the state into the callback class containing the callback methods. Here is a simple implementation of a callback class that stores a widget that can be retrieved by the `end_` method:

C#

```
public class MyCallback
{
    public MyCallback(Widget w)
    {
        _w = w;
    }

    private Widget _w;

    public void getNameCB(string name)
    {
        _w.writeString(name);
    }

    public void failureCB(Ice.Exception ex)
    {
        _w.writeError(ex);
    }
}
```

When you call the `begin_` method, you pass the appropriate callback instance to inform the `end_` method how to update the display:

C#

```
EmployeesPrx e = ...;
Widget widget1 = ...;
Widget widget2 = ...;

MyCallback cb1 = new MyCallback(widget1);
MyCallback cb2 = new MyCallback(widget2);

// Invoke the getName operation with different widget callbacks.

e.begin_getName(99).whenCompleted(cb1.getNameCB, cb1.failureCB);
e.begin_getName(24).whenCompleted(cb2.getNameCB, cb2.failureCB);
```

[Back to Top ^](#)

Asynchronous Oneway Invocations in C#

You can invoke operations via oneway proxies asynchronously, provided the operation has `void` return type, does not have any out-parameters, and does not raise user exceptions. If you call the `begin_` method on a oneway proxy for an operation that returns values or raises a user exception, the `begin_` method throws a `System.ArgumentException`.

For the generic API, the callback method looks exactly as for a twoway invocation. However, for oneway invocations, the Ice run time does not call the callback method unless the invocation raised an exception during the `begin_` method ("on the way out").

For the type-safe API, you only specify a delegate for the failure method. For example, here is how you could call `ice_ping` asynchronously:

C#

```
ObjectPrx p = ...;
MyCallback cb = new MyCallback();
p.begin_ice_ping().whenCompleted(cb.failureCB);
```

[Back to Top ^](#)

Flow Control in C#

Asynchronous method invocations never block the thread that calls the `begin_` method: the Ice run time checks to see whether it can write the request to the local transport. If it can, it does so immediately in the caller's thread. (In that case, `AsyncResult.sentSynchronously` returns true.) Alternatively, if the local transport does not have sufficient buffer space to accept the request, the Ice run time queues the request internally for later transmission in the background. (In that case, `AsyncResult.sentSynchronously` returns false.)

This creates a potential problem: if a client sends many asynchronous requests at the time the server is too busy to keep up with them, the requests pile up in the client-side run time until, eventually, the client runs out of memory.

The API provides a way for you to implement flow control by counting the number of requests that are queued so, if that number exceeds some threshold, the client stops invoking more operations until some of the queued operations have drained out of the local transport.

For the [generic API](#), you can create an additional callback method:

C#

```
public class MyCallback
{
    public void finished(Ice.AsyncResult r)
    {
        // ...
    }

    public void sent(Ice.AsyncResult r)
    {
        // ...
    }
}
```

As with any other callback method, you are free to choose any name you like. For this example, the name of the callback method is `sent`. You inform the Ice run time that you want to be informed when a call has been passed to the local transport by calling `whenSent`:

C#

```
MyCallback cb = new MyCallback();

e.begin_getName(99).whenCompleted(cb.getNameCB, cb.failureCB).whenSent(cb.sent);
```

If the Ice run time can immediately pass the request to the local transport, it does so and invokes the `sent` method from the thread that calls the `begin_` method. On the other hand, if the run time has to queue the request, it calls the `sent` method from a different thread once it has written the request to the local transport. In addition, you can find out from the `AsyncResult` that is returned by the `begin_` method whether the request was sent synchronously or was queued, by calling `sentSynchronously`.

For the [generic API](#), the `sent` method has the following signature:

C#

```
void sent(Ice.AsyncResult r);
```

For the [type-safe API](#), the signature is:

C#

```
void sent(bool sentSynchronously);
```

For the generic API, you can find out whether the request was sent synchronously by calling `sentSynchronously` on the `AsyncResult`. For the type-safe API, the boolean `sentSynchronously` parameter provides the same information.

The `sent` methods allow you to limit the number of queued requests by counting the number of requests that are queued and decrementing the count when the Ice run time passes a request to the local transport.

[Back to Top ^](#)

Asynchronous Batch Requests in C#

You can invoke operations via batch oneway proxies asynchronously, provided the operation has `void` return type, does not have any out-parameters, and does not raise user exceptions. If you call the `begin_` method on a oneway proxy for an operation that returns values or raises a user exception, the `begin_` method throws a `System.ArgumentException`.

A batch oneway invocation never calls the generic or type-safe callbacks unless an error occurs before the request is queued. The returned `Ice.AsyncResult` for a batch oneway invocation is always completed and indicates the successful queuing of the batch invocation. The returned result can also be marked completed if an error occurs before the request is queued.

Applications that send [batched requests](#) can either flush a batch explicitly or allow the Ice run time to flush automatically. The proxy method `ice_flushBatchRequests` performs an immediate flush using the synchronous invocation model and may block the calling thread until the entire message can be sent. Ice also provides asynchronous versions of this method so you can flush batch requests asynchronously.

`begin_ice_flushBatchRequests` and `end_ice_flushBatchRequests` are proxy methods that flush any batch requests queued by that proxy.

In addition, similar methods are available on the communicator and the `Connection` object that is returned by `AsyncResult.getConnection`. These methods flush batch requests sent via the same communicator and via the same connection, respectively.

[Back to Top ^](#)

Concurrency Semantics for AMI in C#

The Ice run time always invokes your callback methods from a separate thread, with one exception: it calls the `sent` callback from the thread calling the `begin_` method if the request could be sent synchronously. In the `sent` callback, you know which thread is calling the callback by looking at the `sentSynchronously` member or parameter.

[Back to Top ^](#)

See Also

- [Request Contexts](#)
- [Batched Invocations](#)
- [Collocated Invocation and Dispatch](#)

