The Server-Side main Program in Python

This section discussed how to initialize and finalize the server-side run time.

On this page:

Initializing and Finalizing the Server-Side Run Time in Python

The main entry point to the Ice run time is represented by the local interface `Ice::Communicator`. As for the client side, you must initialize the Ice run time by calling `Ice.initialize` before you can do anything else in your server. `Ice.initialize` returns a reference to an instance of `Ice.Communicator`:

Python

```python
import sys, traceback, Ice

status = 0
ic = None
try:
    ic = Ice.initialize(sys.argv)
    # ...
except:
    traceback.print_exc()
    status = 1

# ...
```

`Ice.initialize` accepts the argument list that is passed to the program by the operating system. The function scans the argument list for any command-line options that are relevant to the Ice run time; any such options are removed from the argument list so, when `Ice.initialize` returns, the only options and arguments remaining are those that concern your application. If anything goes wrong during initialization, `initialize` throws an exception.

You can pass a second argument of type `InitializationData` to `Ice.initialize`. `InitializationData` is defined as follows:

Python

```python
class InitializationData(object):
    def __init__(self):
        self.properties = None
        self.logger = None
        self.threadHook = None
```

You can pass in an instance of this class to set properties for the communicator, establish a logger, and to establish a thread notification hook.

Before leaving your program, you *must* call `Communicator.destroy`. The `destroy` operation is responsible for finalizing the Ice run time. In particular, `destroy` waits for any operation implementations that are still executing in the server to complete. In addition, `destroy` ensures that any outstanding threads are joined with and reclaims a number of operating system resources, such as file descriptors and memory. Never allow your program to terminate without calling `destroy` first; doing so has undefined behavior.

The general shape of our server-side program is therefore as follows:

Python

```python
import sys, traceback, Ice

status = 0
ic = None
try:
    ic = Ice.initialize(sys.argv)
    # ...
except:
    traceback.print_exc()
    status = 1

if ic:
    try:
        ic.destroy()
    except:
        traceback.print_exc()
        status = 1

sys.exit(status)
```

Note that the code places the call to `Ice.initialize` into a `try` block and takes care to return the correct exit status to the operating system. Also note that an attempt to destroy the communicator is made only if the initialization succeeded.

The `Ice.Application` Class in Python

The preceding program structure is so common that Ice offers a class, `Ice.Application`, that encapsulates all the correct initialization and finalization activities. The synopsis of the class is as follows (with some detail omitted for now):

Python

```python
class Application(object):

    def __init__(self, signalPolicy=0):

    def main(self, args, configFile=None, initData=None):

    def run(self, args):

    def appName():
        # ...
    appName = staticmethod(appName)

    def communicator():
        # ...
    communicator = staticmethod(communicator)
```

The intent of this class is that you specialize `Ice.Application` and implement the abstract `run` method in your derived class. Whatever code you would normally place in your main program goes into `run` instead. Using `Ice.Application`, our program looks as follows:

Python

```python
import sys, Ice

class Server(Ice.Application):
    def run(self, args):
        # Server code here...
        return 0

app = Server()
status = app.main(sys.argv)
sys.exit(status)
```

You also can call `main` with an optional file name or an `InitializationData` structure. If you pass a configuration file name to `main`, the property settings in this file are overridden by settings in a file identified by the `ICE_CONFIG` environment variable (if defined). Property settings supplied on the command line take precedence over all other settings.

The `Application.main` function does the following:

1. It installs an exception handler. If your code fails to handle an exception, `Application.main` prints the exception information before returning with a non-zero return value.
2. It initializes (by calling `Ice.initialize`) and finalizes (by calling `Communicator.destroy`) a communicator. You can get access to the communicator for your server by calling the static `communicator` accessor.
3. It scans the argument list for options that are relevant to the Ice run time and removes any such options. The argument list that is passed to your `run` method therefore is free of Ice-related options and only contains options and arguments that are specific to your application.
4. It provides the name of your application via the static `appName` member function. The return value from this call is the first element of the argument vector passed to `Application.main`, so you can get at this name from anywhere in your code by calling `Ice.Application.appName` (which is often necessary for error messages).
5. It installs a signal handler that properly shuts down the communicator.
6. It installs a per-process logger if the application has not already configured one. The per-process logger uses the value of the `Ice.ProgramName` property as a prefix for its messages and sends its output to the standard error channel. An application can also specify an alternate logger.

Using `Ice.Application` ensures that your program properly finalizes the Ice run time, whether your server terminates normally or in response to an exception or signal. We recommend that all your programs use this class; doing so makes your life easier. In addition, `Ice.Application` also provides features for signal handling and configuration that you do not have to implement yourself when you use this class.

Using `Ice.Application` on the Client Side in Python

You can use `Ice.Application` for your clients as well: simply implement a class that derives from `Ice.Application` and place the client code into its `run` method. The advantage of this approach is the same as for the server side: `Ice.Application` ensures that the communicator is destroyed correctly even in the presence of exceptions or signals.

Catching Signals in Python

The simple server we developed in Hello World Application had no way to shut down cleanly: we simply interrupted the server from the command line to force it to exit. Terminating a server in this fashion is unacceptable for many real-life server applications: typically, the server has to perform some cleanup work before terminating, such as flushing database buffers or closing network connections. This is particularly important on receipt of a signal or keyboard interrupt to prevent possible corruption of database files or other persistent data.

To make it easier to deal with signals, `Ice.Application` encapsulates Python's signal handling capabilities, allowing you to cleanly shut down on receipt of a signal:

Python

```
class Application(object):
    # ...
    def destroyOnInterrupt():
        # ...
    destroyOnInterrupt = classmethod(destroyOnInterrupt)

    def shutdownOnInterrupt():
        # ...
    shutdownOnInterrupt = classmethod(shutdownOnInterrupt)

    def ignoreInterrupt():
        # ...
    ignoreInterrupt = classmethod(ignoreInterrupt)

    def callbackOnInterrupt():
        # ...
    callbackOnInterrupt = classmethod(callbackOnInterrupt)

    def holdInterrupt():
        # ...
    holdInterrupt = classmethod(holdInterrupt)

    def releaseInterrupt():
        # ...
    releaseInterrupt = classmethod(releaseInterrupt)

    def interrupted():
        # ...
    interrupted = classmethod(interrupted)

    def interruptCallback(self, sig):
        # Default implementation does nothing.
        pass
```

The methods behave as follows:

- destroyOnInterrupt
  This method installs a signal handler that destroys the communicator if it is interrupted. This is the default behavior.

- shutdownOnInterrupt
  This method installs a signal handler that shuts down the communicator if it is interrupted.

- ignoreInterrupt
  This method causes signals to be ignored.

- callbackOnInterrupt
  This method configures Ice.Application to invoke interruptCallback when a signal occurs, thereby giving the subclass responsibility for handling the signal.

- holdInterrupt
  This method temporarily blocks signal delivery.

- releaseInterrupt
  This method restores signal delivery to the previous disposition. Any signal that arrives after holdInterrupt was called is delivered when you call releaseInterrupt.

- interrupted
  This method returns True if a signal caused the communicator to shut down, False otherwise. This allows us to distinguish intentional shutdown from a forced shutdown that was caused by a signal. This is useful, for example, for logging purposes.

- interruptCallback
  A subclass overrides this method to respond to signals. The function may be called concurrently with any other thread and must not raise exceptions.

By default, `Ice.Application` behaves as if `destroyOnInterrupt` was invoked, therefore our server program requires no change to ensure that the program terminates cleanly on receipt of a signal. (You can disable the signal-handling functionality of `Ice.Application` by passing `1` to the constructor. In that case, signals retain their default behavior, that is, terminate the process.) However, we add a diagnostic to report the occurrence of a signal, so our program now looks like:

Python

```python
import sys, Ice

class MyApplication(Ice.Application):
    def run(self, args):

        # Server code here...

        if self.interrupted():
            print self.appName() + ": terminating"

        return 0

app = MyApplication()
status = app.main(sys.argv)
sys.exit(status)
```

Note that, if your server is interrupted by a signal, the Ice run time waits for all currently executing operations to finish. This means that an operation that updates persistent state cannot be interrupted in the middle of what it was doing and cause partial update problems.

`Ice.Application` and Properties in Python

Apart from the functionality shown in this section, `Ice.Application` also takes care of initializing the Ice run time with property values. Properties allow you to configure the run time in various ways. For example, you can use properties to control things such as the thread pool size or port number for a server. The `main` method of `Ice.Application` accepts an optional second parameter allowing you to specify the name of a configuration file that will be processed during initialization.

Limitations of `Ice.Application` in Python

`Ice.Application` is a singleton class that creates a single communicator. If you are using multiple communicators, you cannot use `Ice.Application`. Instead, you must structure your code as we saw in Hello World Application (taking care to always destroy the communicator).

See Also

- Hello World Application
- Properties and Configuration
- Communicator Initialization
- Logger Facility
- Thread Safety