



Magic Internet Kit

Programmer's Guide



General Magic

Copyright © 1996 General Magic, Inc.

All rights reserved.

License

Your use of the software discussed in this document shall be permitted only pursuant to the terms in a software license between you and General Magic.

Trademarks

The General Magic logo, the Magic Cap logo, the Telescript logo, Magic Cap, Telescript, and the rabbit-from-a-hat logo are trademarks of General Magic, and may be registered in certain jurisdictions.

All other trademarks and service marks are the property of their respective owners.

Limit of Liability/Disclaimer of Warranty

THIS BOOK IS SOLD "AS IS." Even though General Magic. has reviewed this book in detail, GENERAL MAGIC MAKES NO REPRESENTATION OR WARRANTIES, EITHER EXPRESS OR IMPLIED, WITH RESPECT TO THE ACCURACY OR COMPLETENESS OF THE CONTENTS OF THIS BOOK. GENERAL MAGIC SPECIFICALLY DISCLAIMS ANY IMPLIED WARRANTIES OR MERCHANTABILITY OR FITNESS FOR ANY PARTICULAR PURPOSE AND SHALL IN NO EVENT BE LIABLE FOR ANY LOSS OF PROFIT OR ANY OTHER COMMERCIAL DAMAGE, INCLUDING BUT NOT LIMITED TO SPECIAL, INCIDENTAL, CONSEQUENTIAL, OR OTHER DAMAGES, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Some states do not allow for the exclusion or limitation of implied warranties or incidental or consequential damage. So, the exclusions in this paragraph might not apply to you. This warranty gives you specific legal rights. You may also have other rights which vary from state to state.

Josh and Ed's Excellent Internet Kit

Lyrics and electric banjo by Josh Carter. Drums and synths by Ed Satterthwaite. Backing vocals by Zarko Draganic, Dean Yu, and C.J. Silverio. Tour management by Mark "The Red" Harlan.

Get out the banjo and let's boogie to this

Restricted Rights. For defense agencies: Use, duplication, or disclosure is subject to the restrictions set forth in subparagraph (c)(1)(ii) of DFAR section 252.227-7013 and its successors. For civilian agencies: Use, duplication, or disclosure is subject to the restrictions set forth in subparagraphs (a) through (d) of FAR section 52.227-19 and its successors. Unpublished—rights reserved under the copyright laws of the United States.

General Magic

420 North Mary Avenue	Telephone: 408 774 4000
Sunnyvale	Fax: 408 774 4010
California 94086 USA	E-mail: dev-info@genmagic.com
	URL: http://www.genmagic.com/

Patent Pending

Portions of the Magic Cap software and the Telescript software are patent pending in the United States and other countries.

Table of Contents

Introduction	5
Real connectivity made real easy	5
About the Workbench 1996 release	6
What you should already know	6
About this document	7
Creating a New Package	9
Templates	9
Templates provided in the kit	9
Cloning a template	10
Where to go from here	10
Connections	11
The Connection class	11
Methods of Connection	12
CreateStream error handling	13
Connection subclasses in the Magic Internet Kit	14
InternetConnection	14
I'm connected. Now what?	16
Streams	17
Essential methods of Stream	17
CountReadPending	17
Read	17
Write	18
Other useful methods	18
ReadUntil	18
WriteLiteral	19
WriteTextAsASCII, WriteTextAsUnicode	19
Synchronous stream issues	19
Multithreading with Actors	21
Actor concepts	21
Creating an actor	21
Using an actor	22

Destroying an actor 23
Moving between actors 23
Actors in the Magic Internet Kit 24

Introduction

Real connectivity made real easy

The Magic Internet Kit is a complete development kit for creating Magic Cap communicating applications. The heart of the Magic Internet Kit is a powerful yet easy-to-use object framework which provides:

- TCP/IP support for writing full-featured Internet/Intranet applications.
- Supporting protocols for TCP/IP such as PPP, with PAP and CHAP authentication, and DNS for resolving host names.
- Serial communications over a communicator's built-in modem and MagicBus port.
- Ability to add additional TCP/IP hardware drivers without needing to change or rebuild client applications

From the first steps of creating a communicating application to maintaining the code later, the Magic Internet Kit makes your work easy and your development cycle fast. From the very beginning, you can use one of several templates provided with the kit to get you started right where you want to be. These templates range from an empty package with very little user interface to a ready-to-go terminal package.

While the handy templates get you started quickly, the real power of the Magic Internet Kit lies in its robust and flexible programming interfaces. This framework is arranged to provide you with a single set of methods that you can call to create any type of communications stream supported by the kit.

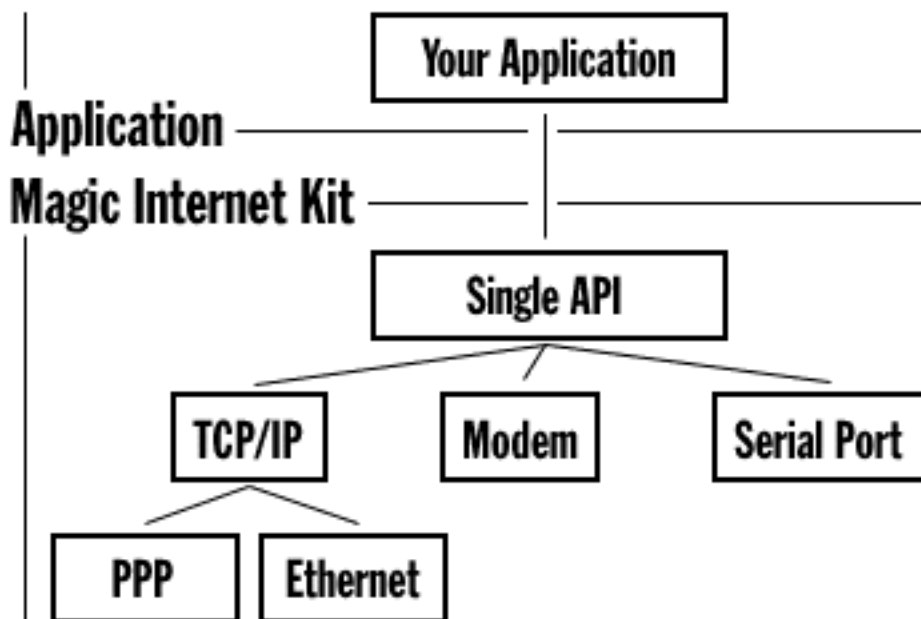


Figure 1 Magic Internet Kit layout

As illustrated in the above figure, many means of communication are available on the bottom end, and the application can access all of them with one API. The Magic Internet Kit API will be discussed in the Connections chapter of this guide.

About the Workbench 1996 release

The Magic Internet Kit released at the Magic Cap Workbench 1996 is an in-progress version designed for Rosemary development. This version includes support for TCP/IP connectivity and classes for taking advantage of the Rosemary Internet Center. Unfortunately, it does not include serial support, the full suite of template applications, or automated package creation tools. These are in development and will be released shortly. Please also excuse errors or omissions in this document; it is still in progress as well.

To get the latest copy of the Magic Internet Kit and the *Programmer's Guide*, see the kit's web site, located at:

<http://www.genmagic.com/Develop/MagicCap/Internet/>

What you should already know

The *Magic Internet Kit Programmer's Guide* assumes that you are familiar with the basics of TCP/IP and related protocols. If you need an introduction, we recommend referring to *Internetworking with TCP/IP* by Douglas Comer or *TCP Illustrated* by

Richard Stevens. If you don't want to write TCP/IP applications, then you don't need to worry; the Magic Internet Kit also works directly using the modem or serial port.

This document also assumes that you are somewhat familiar with Magic Cap development. If you need to find out more about the basics of writing a Magic Cap application, see the *Magic Cap Package Development Guide*. You can find this document and other information at General Magic's *Magic Cap Development Overview* web page located at:

<http://www.genmagic.com/Develop/MagicCap/Overview/index.html>

About this document

The *Magic Internet Kit Programmer's Guide* starts with a discussion of the templates available for creating a package, and then proceeds on to discussing the high-level APIs for communications work. Later chapters dive into detail about TCP/IP and multithreaded communications functionality. This section will briefly describe the topics of each chapter.

Creating a New Package

This chapter describes the process of creating a new communicating package. Included are discussions of the templates applications provided by the Magic Internet Kit and how to clone a template for your own application.

Connections

This chapter discusses the centerpiece of the Magic Internet Kit: the **Connection** class. First it describes the API provided by the Connection class, and then discusses specific subclasses of Connection, including **InternetConnection**.

Streams

This chapter describes usage of communication streams within Magic Cap. The **Stream** class is used for reading and writing data, so topics include the methods used to read, write, and catch errors with the Stream class.

Multithreading with Actors

Multithreading, or running tasks in the background, is essential when writing a communicating application, so this chapter describes how to do that. The topics discussed in this chapter are illustrated in code with the CujoTerm template.

Creating a New Package

Templates

The first step in creating a new package is to clone an existing template. These templates are functional Magic Cap packages in themselves, but are designed to be starting points for your own application development. This chapter describes the templates provided in the Magic Internet Kit and how to use them.

Templates provided in the kit

Note: The Magic Internet Kit is still under development and does not include its full suite of templates. This section describes the existing templates, but more will be included in future releases.

Templates are provided to give you a good starting point for a new application. There are several interface elements that many communicating packages may all want to use, for example a choice box that lets the user pick an Internet provider, so the Magic Internet Kit lets you start with varying levels of functionality built in. This section will describe each of the included templates.

Terminal (a.k.a. CujoTerm)

This is the classic TTY terminal package; CujoTerm gives you a simple terminal that you can connect to a remote host over whatever connections means the kit supports. CujoTerm's **TTYTerminal** class is a **TextField** subclass which displays text from the communication stream, accepts typing and sends data to the stream, and even lets the user drop text coupons that will be sent to the stream.

CujoTerm also uses a **ProviderChoiceBox**, a subclass of **ChoiceBox**, which allows the user to choose an Internet service provider (ISP) for the connection. **ProviderChoiceBox** is defined by the core Internet Kit framework since it's very handy, so you can use this class even if you don't start from the CujoTerm template.

Internet providers are set up in Rosemary's Internet Center, so your package does not have to include this functionality; you can leverage the information using `ProviderChoiceBox` or your own code.

The key component of `CujoTerm` is its architecture. This terminal, while basic in user-level features, has an advanced multi-threaded structure behind it that can support multiple simultaneous TCP/IP connections. The two key classes that implement this functionality are the **`CommsActor`** and **`CommsManager`**. The code that executes on its own thread is part of the `CommsActor`, and the `CommsManager` is used to easily create and destroy `CommsActors`. The concepts behind these classes will be discussed in greater detail in the Multithreaded Communications chapter of this document.

`CujoTerm` is a good template to use for testing out code or investigating protocols. For example, if you wanted to write a package that queried and responded to a database, you could use `CujoTerm` to make sure that the commands you are sending to the server are actually doing the right thing. Additionally, since `CujoTerm` works "out of the box," you can use it as a good sample package to learn from.

Finger

The `Finger` template package is a basic finger client as defined by RFC 742. Its user interface is very sparse, but like `CujoTerm` it includes the handy `ProviderChoiceBox` class for picking an Internet provider. `Finger` is a good clean slate to start a new package from if you don't anticipate using all of `CujoTerm`'s goodies.

Cloning a template

Cloning a template is just like cloning any other package. First, set MPW's active directory to the directory of your preferred template. The templates are located in the "Magic Internet Kit:Templates" folder inside your "MagicDeveloper" folder.

Once you have set MPW's current directory, select "Clone Package" menu item from the "Build" menu, and follow the script's prompts for choosing your new package's name and location.

Where to go from here

Now that you're all set up with a clone of your preferred template, it's time to make it do some real work. The next chapter, `Connections`, shows you how to connect to a remote host, use the connection, and then destroy it when you're done.

Connections

The heart of the Magic Internet Kit is the easy-to-use object framework for connecting your application to the outside world. This act can take several forms, from serial communications to TCP/IP over a variety of data links. This chapter describes the interfaces that your packages use to create, use, and destroy links to remote hosts.

The Connection class

The **Connection** class defines the central programming interface for the Magic Internet Kit's framework. Subclasses of Connection are defined for each type of communications means supported by the kit; for example **InternetConnection** is used for TCP/IP-based communications.

Connection subclasses serve two purposes: they contain the data needed to establish their particular type of data stream, and also implement the API defined to by Connection to create and destroy their streams. Like other Magic Cap objects, connection data is stored in fields that are accessed via attributes. The exact fields are dependent on the type of connection, so templates like **CujoTerm** use **AttributeText** objects and variants thereof to provide a user interface for filling in data. This model allows the package to have an appropriate user interface for the connection type, for example host name and port fields for a TCP/IP connection, but the interface is totally separate from the mainstream application code.

This section will discuss the API defined by the Connection class, and following sections will discuss details of each of the Connection subclasses.

Source Code Note: The Connection class definition can be found in the Connection:Connection.cdef file.

Methods of Connection

There are three methods defined by Connection: CanCreateStream, CreateStream, and DestroyStream. The methods do what they sound like, and we'll cover what each one does in detail.

CanCreateStream

```
operation CanCreateStream(): Boolean;
```

This method is called by client code to determine if the connection could potentially create a stream. For example, if this Connection subclass communicates using the modem, CanCreateStream could check to see if the phone line is plugged in, make sure the user has set up a dialing location, and also check that the modem is not already in use.

The application calling CanCreateStream should be aware that the boolean return value is not a guarantee that a connection attempt will succeed. There are many factors that can effect the success of a connection attempt, and many of these cannot be judged until the attempt is actually made. For example, if the object communicates using the modem, and the remote server is not answering, the object won't know that until after dialing. CanCreateStream therefore should only be counted on as a "sanity check" before doing any real work to try connecting.

CreateStream

```
operation CreateStream(): Stream;
```

CreateStream is the method that tells a Connection subclass to create a communication stream to its remote host. If the attempt succeeds, this method will return a reference to the stream it created. The stream returned from CreateStream is always a subclass of **CommunicationStream**, for example the **Modem** class for serial modem connections, or the **TCPStream** class for TCP/IP connections.

If the connection attempt fails, CreateStream will throw a **CommsException** heavyweight exception. Error handling with CommsExceptions is slightly different than other exceptions in Magic Cap, so this topic will be discussed shortly in the following CreateStream error handling section.

Note: With the InternetConnection class, you can call CreateStream repeatedly to get multiple streams over the same data link.

DestroyStream

```
operation DestroyStream(stream: Stream);
```

DestroyStream is the inverse of CreateStream; whenever a stream is no longer needed, this method is used to destroy it in whatever manner is appropriate for the specific Connection subclass. This method will also destroy any objects or buffers that were allocated by CreateStream. For example, if CreateStream creates a TCPStream object, DestroyStream will be sure to destroy that object and its local

buffers. Never call Magic Cap's **Destroy** method on a stream returned by `CreateStream`; always use the kit's `DestroyStream` method since it knows how to deal with each type of stream appropriately.

CreateStream error handling

As a general rule with communications, connecting to a remote host requires quite a few things to all work together with each other, and therefore there are many potential places for things to go wrong. Errors in Magic Cap are typically handled with **Exceptions** that get thrown when the error occurs and are then caught by application code designed to specifically handle that error. The exception itself is just an indexical value that serves as a unique identifier for a type of error. For more information on handling these types of exceptions, see the *Magic Cap Package Development Guide*.

With the multitude of possible errors that could occur when creating a communication stream, an application using the typical Magic Cap exception model would have to catch almost a dozen distinct types of exceptions. The code for setting up the error handling could be pages long! The Magic Internet Kit takes a slightly different but much cleaner approach by using **heavyweight exceptions**.

A heavyweight exception is similar to any other exception in how it is used, but in this case the exception is not an error code but rather a real, live object. In the case of `CreateStream`, the object that is thrown is a **CommsException** or one of its subclasses. Code that is calling `CreateStream`, therefore, can catch all of the possible `CommsExceptions` by catching them by their class. The magical method that does this is called, quite reasonably, **CatchByClass**. Here is an example of how to use `CatchByClass` to catch all `CommsException` errors:

```
Reference exception = CatchByClass(CommsException_);

if (exception != nilObject)
{
    /* we caught a CommsException object */
    return false; /* your error handling code goes here */
}

stream = CreateStream(iMyConnection);

Commit(); /* CommsException_ */
```

Once you have caught a `CommsException` object, you can find out which exception it is by comparing it to the ones listed in the `Connection:CommsException.odef` file included in your Magic Internet Kit folder. For example:

```
if ((exception = CatchByClass(CommsException_)) != nilObject)
{
    /* we caught a CommsException object */
    if (exception == ieCannotConnectHardware)
    {
        /* can't connect the hardware to the remote host */
        return false;
    }
    else if /* etc... */
}
}
```

The `CujoTerm` template uses this approach and announces error messages to the user for each exception.

Additionally, errors are grouped into subclasses of `CommsException`, for example all three of the DNS-related errors are of the `DNSCommsException` class. This gives you an extra level of granularity if you want it. For the case of DNS, you can check to see if an exception was any of the DNS errors with the following code:

```
if ((exception = CatchByClass(CommsException_)) != nilObject)
{
    /* we caught a CommsException object */
    if (Implements(exception, DNSCommsException_))
    {
        /* a DNS-related error occurred */
        return false;
    }
    else if /* etc... */
}
```

Furthermore, this feature means that you can `CatchByClass` on particular subclasses of `CommsException` if you want to have different handlers for them. In practice this last technique might not be useful since all `CommsException` errors are fatal for a given connection attempt, but it might come in handy.

1.x to Rosemary Note: With the Rosemary Internet Kit, you no longer have to destroy the exception objects. The objects thrown are indexicals, but they all point to a single `CommsException` instance. As such, only one `CommsException` exists but the different indexical values make each one unique.

Connection subclasses in the Magic Internet Kit

In the Workbench 1996 release of the Magic Internet Kit, the only subclass of `Connection` is `InternetConnection`, which is used for TCP/IP based communications. Other subclasses are in development for serial modem and serial port use, and these will be available shortly. The following section will describe the details of the `InternetConnection` class, and in the future will discuss additional `Connection` subclasses.

InternetConnection

`InternetConnection` is used to create and destroy TCP/IP streams. This is usually a fairly daunting task, but `InternetConnection` takes care of all the details for you; your application only has to deal with the above three methods defined by `Connection`. Thanks to the Internet Center in Rosemary, filling in the fields of an `InternetConnection` object is also a snap. The user fills in all of the required information for their Internet Service Provider (ISP) in the Internet Center, and then the application can easily use that information. First we'll discuss filling in `InternetConnection`'s fields, and then how that data is used.

InternetConnection attributes

```
attribute ServiceChoice:  InternetServiceChoice;
attribute ServiceInfo:   InternetServiceInfo;
attribute HostName:      Text;
attribute HostPort:     Unsigned;
```

InternetConnection defines two key attributes: ServiceChoice and ServiceInfo. These are filled in with new classes defined in Rosemary: **InternetServiceChoice** and **InternetServiceInfo**. The first class, InternetServiceChoice, defines a service provider and contains the relevant data needed for a connection. This includes a **Means** object which holds information about the hardware driver, the ISP's name, and other moderately interesting things. Fortunately, your application does not have to create or set up these objects; they are created by the Internet Center when the user sets up an ISP. The easiest way to fill in this field is to use the Magic Internet Kit's **ProviderChoiceBox** class or design a variant thereof. This class is used by both the Finger and CujoTerm templates, and it provides a list of the ISPs that the user has set up in the Internet Center. It then sets its iCurrentServiceChoice indexical to the matching InternetServiceChoice object. Simply put an instance of ProviderChoiceBox in your package and set the ServiceChoice attribute of your InternetConnection objects to iCurrentServiceChoice.

The second class, InternetServiceInfo, specifies the remote host name and port that you want to connect to. Setting a destination host name and port are very common tasks, so InternetConnection helps out by defining its own **HostName** and **HostPort** attributes. Getting and setting these attributes with an InternetConnection object will actually get and set fields within its InternetServiceInfo object, but your application does not have to notice the difference. Additionally, creating a new InternetConnection object will automatically create a matching InternetServiceInfo object for it to use.

Note: The HostName attribute of InternetConnection can be either a symbolic host name (e.g. www.genmagic.com) or an IP number (e.g. 192.216.22.142). If the host name is symbolic, InternetConnection will know to use Domain Name Resolution and look the name up automatically.

The easiest way to set up an InternetConnection is to use **AttributeText** and **UnsignedAttributeText** objects that point to the HostName and HostPort attributes. The first class, AttributeText, serves as the data store for a text field and fills in a text attribute of its target object with the contents of the field. The second class, UnsignedAttributeText, is a subclass of AttributeText defined by the Magic Internet Kit for targeting an Unsigned attribute. See the CujoTerm sample for usage of these two classes.

Source Code Note: Examples of AttributeText and UnsignedAttributeText can be found in the CujoTerm:InternetSetup.odef file.

Details of InternetConnection

InternetConnection is quite intelligent about its implementation of the core Connection methods. CanCreateStream checks to make sure that its ServiceChoice and ServiceInfo objects have enough data for the connection, and also makes sure the hardware is not already in use by a non-TCP/IP application. One very useful

feature in Rosemary is that TCP/IP is implemented in the Magic Cap itself, so multiple applications can share a data link to a service provider. `InternetConnection` is aware of this feature and supports it accordingly.

`CreateStream` also fully utilizes Magic Cap's resources for multiple applications sharing a data link. If a data link to a given service provider does not exist, it will create one automatically. If a data link is already up and running, even if it was created by another application, `CreateStream` will use the existing link. Also note that `CreateStream` will automatically look up symbolic host names if needed, so your application does not have to worry about DNS.

`DestroyStream` is just as intelligent as `CreateStream`. If `CreateStream` starts up a new data link, `DestroyStream` will take it down as long as there are no other users. In fact, if another application starts up a stream on the data link created by your application, `DestroyStream` will leave the link in place for the other application.

Internally, `CreateStream` and `DestroyStream` keep a list of streams and another object, called a **TransferTicket**, which is used for the lower-level communication APIs. `InternetConnection` keeps this information around so it always knows how to properly destroy a stream; you can call `CreateStream` on an `InternetConnection`, change the `ServiceChoice` attribute, and then call `DestroyStream`, and the `InternetConnection` will still know how to disconnect the stream from the original service provider.

I'm connected. Now what?

Now that you know about creating and destroying streams with `Connection` objects, and how to fill in the fields of these objects, it's time to look at the object returned by `CreateStream`: the stream. Streams will be the topic of the next chapter.

Streams

Subclasses of **Stream** are used for sending data between Magic Cap and a remote host. There are two key methods that are used with streams: **Read** and **Write**. There are other handy methods for finding out how many bytes are waiting to be read, or writing the contents of a null-terminated c string, and we will discuss many of these.

Essential methods of Stream

Streams are quite easy to use. **CountReadPending** returns the number of bytes available to read, **Read** reads the bytes, and **Write** writes bytes.

CountReadPending

```
operation CountReadPending(): Unsigned, noFail;
```

This method is used to find out how many bytes are in the stream's local buffer and are therefore available for immediate reading. Reading is a synchronous operation, so code that does not want to block while reading should always call **CountReadPending** first to check how many bytes are available.

Read

```
operation Read(buffer: Pointer; count: Unsigned): Unsigned, noFail;
```

This method is used to read a number of bytes from a stream into a buffer. If the number of bytes requested by **Read** are not available yet, for example the server has not sent them, **Read** will block until either all the bytes are received or an error has occurred. The return value is the number of bytes that were read, and comparing this value to the number of bytes requested lets the application know if an error occurred; if fewer bytes were returned than requested, then an error occurred but **Read** still returned as much data as it could get.

One useful tactic when waiting for data is to block on a one character read and then read any other pending data. For example, the following code is used in the `CujoTerm` template application:

```
count = Read(stream, &buffer, 1);

if (count != 1)
{
    /* error case */
    Fail(iServerAborted);
}
else /* count == 1; no errors */
{
    Unsigned count = CountReadPending(stream);
    if (count != 0)
    {
        if (count > 254) count = 254;
        Read(stream, &buffer[1], count);
    }

    /* Remember that we already read one byte up above! */
    count++;

    HandleBytes(client, (Pointer)buffer, count);
}
```

In this case, the one character read will block until either data is available or an error occurred, for example the remote host closing the stream. When `Read` returns, a check is made to make sure that `Read` returned a character, and if not an exception is thrown. If the character was indeed read okay, the rest of the code will get the rest of the pending bytes, up to 255 total, and process them.

Write

```
operation Write(buffer: Pointer; count: Unsigned), noFail;
```

`Write`, as one can imagine, is used to write data to a stream. With TCP streams, write will immediately return after placing the bytes into TCP's outgoing buffer. If the stream was closed for some reason, or some other error occurs, `Write` will instead throw an `iServerAborted` exception. As a result, all calls to `Write` should be prepared to catch `iServerAborted` exceptions and handle the error case.

Other useful methods

The stream class defines several additional methods for your convenience. All of these methods are based on the essential three methods discussed above, so error handling tactics are identical; methods that Read data should check return values, and methods that Write data should catch `iServerAborted` exceptions.

ReadUntil

```
operation ReadUntil(buffer: Pointer; endChar: UnsignedByte; maxLen: Unsigned;
    var numRead: Unsigned; VAR numInBuf: Unsigned): Boolean, noFail;
```

ReadUntil is just like `Read`, except that is used to read until a specified character is found, or until a specified number of characters is read. This method can be very useful for reading data one line at a time; your code can `ReadUntil` the EOL character very easily. The boolean return value determines if the character was found. If the return value is true, the character was found. This character will not be in the buffer, so the `*numInBuf` parameter will be one less than `*numRead`. If the return value is false, `ReadUntil` read until the `maxLength` parameter was reached but did not find the character, or `ReadUntil` encountered an error. To tell which problem caused `ReadUntil` to return false, compare `*numRead` to `maxLength`. If these values are the same, `ReadUntil` did not find the character, but if `*numRead` is less than `maxLength`, and error occurred while reading from the stream.

WriteLiteral

```
operation WriteLiteral(dataToWrite: Literal);
```

`WriteLiteral` is used to write a null-terminated c-style string to a stream. It will write everything in the literal except for the null terminator.

WriteTextAsASCII, WriteTextAsUnicode

```
operation WriteTextAsASCII(text: HasText);  
operation WriteTextAsUnicode(text: HasText);
```

These two methods write **Text** objects to a stream. The first method, `WriteTextAsASCII`, writes the contents of a `Text` object as 8-bit characters and will replace all non-ASCII characters with '?'. The second method, `WriteTextAsUnicode`, writes the text object as 16-bit unicode characters.

Synchronous stream issues

All stream methods are synchronous, meaning they will not return until they have either finished their task or have encountered an error. This detail may not be very important for writing data since writing can usually send data down a stream fairly quickly, but it is very important for reading. The `Read` method will block until all bytes requested are actually read, so if the remote host is slow sending data, your application must make provisions for not making the user interface halt while waiting.

There are two means of ensuring that your application does not block user interaction. The first is to periodically check `CountReadPending` until all data that you want can be read, and then call `Read` knowing that the data is available in the stream's local buffer. The second, and by far most preferable way, is to do all blocking communications work on a different thread. With this tactic, user interaction takes place on its own tread, so your communications thread can block and the user will not notice; they can go on using other parts of your application or `Magic Cap`. Multithreading is the topic of our next chapter.

Multithreading with Actors

Magic Cap is a multi-threading platform, and threads are very handy for communications. As you may have noticed with the Finger template, blocking all user interaction while dialing the phone or waiting for a remote server is, at best, very annoying for the user. Instead of executing time consuming code on the thread that handles user interaction, applications should create their own threads for these tasks.

In Magic Cap, a thread is called an **Actor**. Applications wishing to create their own actors must subclass the Actor class and override its **Main** method to make it do what they want. This section will briefly cover how to use actors, but the *Magic Cap Package Development Guide* provides a more complete discussion of this topic. We recommend that you read that chapter as you have time.

Actor concepts

Like everything in Magic Cap, an actor is an object. The base Actor class does not really do anything when you create it, but rather it is meant to be subclassed. The first method that you should override is Main. Main is the heart of the actor; when the actor is created, Main is executed. When Main returns, the actor is destroyed.

Magic Cap uses cooperative multitasking with its actors, so each actor should be sure to give up time to other actors. This is done by calling **RunNext** on the scheduler, referenced by its class number **Scheduler_**. RunNext tells the scheduler to run the next waiting actor.

Some methods that you might call will automatically call RunNext if they are going to take a while to return. For example, if you call Read on a TCPStream object, and the Read cannot be immediately satisfied because all the bytes requested are not available, Read will call RunNext to let other actors do their stuff.

Creating an actor

Actors are created using the **NewTransient** method. Here's an example:

```
newActor = NewTransient(CommsActor_, nil);
```

This code will create a new `CommsActor` class. The second `nil` parameter is for parameters that one might want to pass for the new actor, and passing `nil` tells Magic Cap to use the default parameters. If you want to pass in parameters, for example the size of the execution stack that the actor should have, you can do so just like for any other object. Here's an example of setting the stack size a bit larger than the default:

```
NewActorParameters newActorParams;  
ZeroMem(&newActorParams, sizeof(NewActorParameters));  
  
newActorParams.stackSize = 0x2000; /* default is 0x1000 (4K) */  
  
newActor = NewTransient(CommsActor_, newActorParams);
```

Once the actor is created, it will be ready to run in the scheduler. Keep in mind that the code in your actor's `Main` method will not start executing until the scheduler switches to the actor. This will happen the next time that you, or the system, calls `RunNext`.

Using an actor

As mentioned above, all the real work of an actor is performed in the `Main` method. Here's a sample main method:

```
Method void  
MyActor_Main(Reference self, Pointer UNUSED(params))  
{  
    while (FeelingPeachy(self))  
    {  
        DoSomeStuff(self);  
  
        if (SomeFatalErrorOccurred(self))  
        {  
            return;  
        }  
    }  
}
```

There is one essential caveat to using actors that you should be aware of: code running on an actor that isn't the User Actor cannot call methods that deal with drawing on the screen. This means that code in the above `MyActor_Main` method cannot move viewables on the screen, call `RedrawNow`, or otherwise change stuff on screen. If you need to modify viewables or draw on the screen, you must use a special method called `RunSoon` to execute a completion function on the User Actor. `RunSoon` will be discussed in the `Moving between actors` section of this chapter.

Note: There is one exception to the rule of not messing with stuff on screen from outside the User Actor: announcements. You can always call the **Announce** method regardless of the current actor since it will automatically use `RunSoon` when needed

There is one more interesting caveat to using actors in Magic Cap: an actor is a transient object, so it may or may not get destroyed when the device is power cycled. If Magic Cap decides to clean up transient clusters while powering up or down, the actor will be destroyed, but otherwise it will stay around and start executing when the device is powered up again. As a result, special care must be taken to ensure that

the state of an actor does not get confused if the power is turned off and then back on. This is particularly important with communications, of course, because any data links will get shut down when the power is turned off.

The key to managing transient actors is overriding the **ResetClass** method of an object. `ResetClass` gets called when the device is powered on, so you can use this method to hunt down any actors that need to be managed and do whatever is appropriate. For communications this usually means destroying the actor.

Source Code Note: The Magic Internet Kit's `CujoTerm` template illustrates how to override `ResetClass` and destroy remaining actors at power-up time. See the `CommsManager` class in `CujoTerm:CommsManager.[cdef/cpp]`.

Destroying an actor

Code can destroy an actor using Magic Cap's **DestroyActor** method. If the code wanting to kill the actor is running on the actor in question, though, it should instead force the actor's `Main` method to return.

Moving between actors

Magic Cap provides mechanisms for code executing on one actor to talk to other actors, so we'll briefly cover those mechanisms here. There are a few different ways to manage multiple threads, including semaphores, cross-actor exception throwing, and the `RunSoon` method mentioned above. `RunSoon` is not covered in this document at the time, but it will be in a future revision to be released shortly.

Semaphores

A **Semaphore** object is very handy for controlling access to a resource. Magic Cap's semaphores are quite a bit more intelligent than the typical semaphore in operating system theory, in fact they behave almost like monitors. Semaphores provide automatic queueing and dequeuing as needed to control access to a single resource from multiple threads, so there is no need to poll a semaphore.

There are two key methods for using a Semaphore: **Access** and **Release**. `Access` is used to hold down the semaphore. If the semaphore being accessed is not already held down by code on another actor, `Access` will return immediately. If the semaphore is already accessed, though, `Access` will block until someone else releases the semaphore using `Release`. `Release` will tell the semaphore that you're done messing with it, and it will then wake up the next actor in the queue, if any, that wanted to access the semaphore.

Cross-actor exceptions

If you're not already familiar with exception handling in Magic Cap, you should read the "Handling Exceptions" chapter of the *Magic Cap Package Developers Guide* for a good introduction. This section briefly describes how to use exceptions with actors.

Every actor has its own exception stack, so an exception thrown on one actor using `Fail` cannot be caught from another actor. This is very useful for localizing exception handling code; for example an `iServerAborted` exception thrown on a communicating application's comms actor won't be caught accidentally by the Post Office actor in the system.

If code executing on one actor wants to throw an exception for a different actor to catch, it should use the **FailSoon** method. `FailSoon` will cause a specified exception to get thrown on the other actor the next time that the actor comes up in the scheduler. If the target actor is not ready to run, e.g. it is blocked, the actor will be awakened and the exception thrown.

Source Code Note: See `CommsManager_DestroyCommsActor` in the `CujoTerm` template for an example of using `FailSoon`.

Actors in the Magic Internet Kit

If all of this multithreading stuff looks intimidating, don't fret; the Magic Internet Kit comes to the rescue. The `CujoTerm` template makes heavy use of actors. All connecting and reading is performed on its `CommsActor` object. Additionally, there is a `CommsManager` class that is used to create and destroy `CommsActor` objects.

Source Code Note: See the `CommsActor.[cdef/cpp]` and `CommsManager.[cdef/cpp]` files in `CujoTerm` for its use of actors.
