

Setting Up a Web Server

While setting up a Macintosh-based Web server from scratch is a little beyond the scope of this Help text, I will cover how the names and locations of your folders and files are related to their usage within your HTML documents.

Before You Begin Editing

Before you begin creating HTML documents, I'd recommend going to the Preferences card under Paths & Labels Preferences, and specifying the location of your HTTP server folder (or at least typing in the intended location) by clicking on the "Set" button. This is the path to your HTTP server folder. Then enter your server name in the HTTP Server Name field, as described below.

Domain Name System (DNS)

Every computer connected to the Internet has a unique IP (Internet Protocol) number called a host address. This is usually assigned by your Internet service provider or site systems administrator, and consists of four numbers between 0 and 255 connected by periods, for example "192.188.119.14". The host address is the method by which computers refer to each other on the Internet. While this works fine for computers, it's difficult for people to remember these numbers, so an Internet-wide database called the Domain Name System (or DNS) exists that keeps track of the system names assigned to each computer, basically a big phone book for the Internet.

While you can have people refer to your Web server by its IP number, such as

```
http://192.188.119.14/default.html
```

this is obviously not the way the pros do it. You'll probably need to talk to your site systems administrator to assign a system name to your server. There is somewhat of a convention of prefixing World Wide Web servers' names with "www", as in "www.ignatz.foo.edu", although this is not strictly necessary: the "http:" designates the service. Once named, you can use the system name in place of the host address, as in

```
http://ignatz.foo.edu/default.html
```

The DNS system name (which I'll henceforth call the "server name") and the Macintosh path to the HTTP server folder (the location of your HTTP application) are both set under the Paths & Labels Preferences.

The Relationship between Macintosh and URL Paths

This is easiest to describe this is by way of our current example. If our HTTP server application resides in a folder "www" on a hard drive named "earth", the server root is

```
earth:www
```

If the server name is "ignatz.foo.edu", then the document "gbateson.html" residing in a folder "people" on the server root would have a Macintosh path of

```
earth:www:people:gbateson.html
```

and a URL of

```
http://ignatz.foo.edu/people/gbateson.html
```

You can see that by substituting the forward slash character for the colon delimiter, we can easily understand the relative location of any file by looking at it's URL. And just what is a URL? Next...the URL.

Uniform Resource Locator

I once had an argument with someone at NASA Headquarters over the real words behind the acronym "URL". While some very authoritative books (yes, even the NutShell® book Managing Internet Information Services) seem to think the "U" stands for universal (as it did with URIs), it really stands for uniform. While the complete URL specification may be considered insomniac reading by some*, the part we need to understand is pretty simple, really.

A URL can be broken up into six parts, some of which are optional:

```
http://ignatz.foo.edu/people/gbateson.html#bali
```

```
http: // ignatz.foo.edu /people/ gbateson.html #bali
scheme      server name      path      document name      fragment-ID
```

Take a deep breath:

Scheme

The scheme is the type of Internet service you are requesting. Is the file to be transferred by ftp, gopher, or http? If the document is coming from your http server, you'll use "http:" for the scheme.

Server Name

As discussed in the DNS section above, this will either be an assigned server name (such as "ignatz.foo.edu") or a host address (such as "192.188.119.14").

Path

This is the location of the document within the hierarchy of directories below the http server folder. In this example, the document "gbateson.html" resides in a folder "people" located in the server folder. There is a discussion of this topic in the section Paths and Filenames, as well as below.

Document Name

This is the actual document name, such as "gbateson.html".

Fragment-ID

The optional use of a fragment id can refer to a specific named anchor (in this case "``") marking a location within the document. If present, the browser will open the document and scroll to the specified named anchor. If absent, the URL points to the top of the document.

Absolute and Relative

There are two types of URLs: absolute (or full), and relative. An absolute URL is fully specified, that is, it includes the complete information necessary for retrieving the document over the Internet. A relative URL includes only the relative path from the current location, including only those parts of the URL that have changed.

For example, if the current document "`gbateson.html`" has a URL of:

```
http://ignatz.foo.edu/people/gbateson/gbateson.html
```

and we want to add a link to a section labelled with a named anchor "`bali`" within the same document, the relative URL would be simply

```
#bali
```

Whereas, if we wanted to refer to that same location from a different document within the same folder as "`gbateson.html`", the relative URL would be

```
gbateson.html#bali
```

By the same token, we could refer to any document on the same server by specifying the relative path from the server root (the server folder).

```
/people/tcleary.html  
/projects/sanskrit.html  
/images/previous.GIF
```

In addition to being shorter, the relative URL has the advantage of making your documents more portable. If you move an entire project to a different location, you may only need modify the top level document to suit its new location.

* RFC 1630 T. Berners-Lee, "Universal Resource Identifiers in WWW: A Unifying Syntax for the Expression of Names and Addresses of Objects on the Network as used in the World-Wide Web", 06/09/1994. (Pages=28) (This remark not meant to disparage Tim in any way. I have read the document and it didn't put me to sleep. But both my sleeping habits and level of patience are abnormal, so I cannot fairly judge what others might find either dull or offensive. I very much enjoyed Miller's Rosy Crucifixion , while my wife wouldn't get near it.)

Back to Why Don't My Links Work?, on to Hosts File, or return to Design Issues.