#The101: String Handling and Manipulation

JRahm, 2012-17-10

In this ninth and penultimate installment of #The101: iRules, we're moving ever closer to equipping you with what I hope will be a solid arsenal of iRules knowledge. With it, the idea is for you to go forth and prosper in the ways of iRules fu. So far we've covered a wide range of topics starting from the overtly simple and introductory to things more and more germane directly to iRules as a technology and language. Today's topic, strings and their manipulation, strikes deeper towards that end of the spectrum than before. First, though, let's make sure we're all on the same page here and recount what's been covered so far. Keep in mind this series builds upon itself in places and if you're not keeping up, you may find yourself confused. So, take a minute to read through the titles here and be sure you're read up on the series.

1. Introduction to Programming & Tcl

2. Introduction to F5 Technology & Terms

3. Introduction to iRules

4. Events & Priorities

5. Control Structures & Operators

6. Variables

7. Logging & Comments

8. Routing

Now…to say we're getting to the heart of the matter, dealing with string commands and parsing, re-arranging and modification, would almost be saying it too lightly…understating. String manipulation is a massive part of iRules, and is in fact a solid part of why we are using Tcl as our language of choice, along with many others that I've covered elsewhere in detail. String manipulation is useful in many ways, in many places. Whether it's re-writing a URI for an inbound HTTP request or parsing part of the TCP payload and twiddling the bits to read a bit differently, or perhaps just determining the first n characters of a string to be used for some purpose or another…all of it revolves around strings.

Fortunately for us, Tcl handles strings extremely well. In fact, everything is a string in Tcl's eyes, and as such there are many powerful tools with which you can twist strings to your desires with relative ease, and great effect. To cover all of the options would be a huge process, but we'll go over the basics here, the things seen most commonly within iRules, and you can research the more obscure wizardry at will. The publicly available documentation is good for most of the commands in question. So, in this article we will cover:

- **What is a string and why do I care?**

- **What are the most commonly used string commands?**

- **What commands besides the Tcl "string" command would be useful for strings?**

What is a string and why do I care?

A string is a particular data type, and is generally understood to be a sequence of characters, either as a literal constant, or represented in variable form. This means that basically anything can be a string. A name, an IP address, a URL…all of them are strings, especially in Tcl. Generally speaking, unless things are specifically typed as an integer or some other data type, it is a safe bet to assume they are a string. That being said, since Tcl is not a statically typed language and thereby does not allow you to specify data types explicitly, it treats everything as a string save for a few specific conditions. This is a good thing for iRules, as it means that there isn't a lot of messing about with data types, and that you can generally manipulate things in string format without much hassle. That means less programming fuss, and more getting the effect you want.

What are the most commonly used string commands?

First off, the most common and widely used command in Tcl for dealing with strings is, quite simply, "string". Mind you, in and of itself this command has little use. There are many, many permutations of this command from changing the case of a string to referencing only a portion of it, to re-ordering it and more. With this single command, and the many sub commands, you can perform the lion's share of your string work within iRules. So the question is really which "string" sub commands are most commonly used?

This one is a bit of an intense, broad sweeping question. There are so many things that you can do with a string in Tcl that listing them all here would be an indigestible amount of information, and wouldn't make sense to portray. As such, I'll do my best to list a few string commands that seem to often crop up in iRules, and discuss what each does. For a full reference on the string command and other Tcl base commands, you can find the official Tcl documentation online here (http://www.tcl.tk/man/tcl8.4/TclCmd/contents.htm)

string tolower

```
string tolower string ?first? ?last?
```

Without question the most common and widely used string command within iRules is also one of the simplest. The tolower command does pretty much what it sounds like. It converts the entirety of a string's contents to lowercase. Meaning, if you had a variable named "$uri" and the contents were, "/Admin/WebAccess", you could run the string tolower command when performing comparisons to get a different result.

For instance:

```
1: set uri "/Admin/WebAccess"
2: log local0. "Uri : $uri"
3: log local0. "Lower Uri: [string tolower $uri]"
```

Would result in "Uri: /Admin/WebAccess" for the first log message, and "Lower Uri: /admin/webaccess" for the second. This is thanks to the string tolower command. Why is this so useful in iRules? Because any time you're performing a string based comparison, it is important to be sure you're comparing things in the same case. Think about comparing a host name, a URI, etc. and suddenly you may see why there's so much value in this simple command. This becomes increasingly important with things like datagroups, where you are comparing a single value against a broad range of key values. Being able to assure they are all in the proper case, and then force the incoming comparison value to that case is extremely useful.

Keep in mind that this, like most of the other string commands, does not actually modify the string itself. If you took our above example where we provided the lowercase URI and referenced $uri again, it would still maintain the original case, unaltered. For example, ensuring you directed users attempting to access the admin portion of an application while ensuring they aren't worried about proper casing gets simpler with the tolower command:

```
1: when HTTP_REQUEST {
2:   if {([HTTP::uri] starts_with "/admin") || ([HTTP::uri] starts_with "/Admin")} {
3:     pool auth_pool
4:   }
5: }
```

Becomes:

```
1: when HTTP_REQUEST {
2:   if {[string tolower [HTTP::uri]] starts_with "/admin"} {
3:     pool auth_pool
4:   }
5: }
```

string length

```
string length string
```

Much as you'd expect given the name, the string length command returns the length of the string in question. This can be used for many different things, but probably the most common use-case observed so far in iRules has been to ascertain whether or not a given command returned a proper result. For instance:

```
1: when HTTP_REQUEST {
2:   set cookie_val [HTTP::cookie "x-my-cookie"]
3:   if {[string length $cookie_val > 1} {
4:     log local0. "cookie was passed properly"
5:     pool http_pool
6:   }
7: }
```

Of course there are many ways to perform a similar check, and some are even more efficient if all you're trying to do is identify whether or not a command returned null or not, but if you want to check to see if a specific answer was set of at least n characters, or for a few other very handy purposes I've seen, the string length command can be handy.

string range

```
string range string first last
```

The string range command allows you to reference a particular portion of a given string and retrieve only that specific range of characters. This could be characters 1-10, the first character to the 3rd, or perhaps the 15th to the end of the string. There are many different ways to reference string segments and divide things up using this command, but the result is the same. It returns the value of the portion of the string you define.

This has proved useful time and time again in iRules for things like retrieving portions of a URI, ensuring that a hostname starts with a particular prefix, or dozens of other such seemingly simple requirements. Without the string range command those benign tasks would be a major headache. Note that the first character in the string starts with an ID of 0, not 1.

For instance, if you're looking at a URI that is "/myApp?user=bob" where bob is a variable username, and you're looking to return only the username you have a few options, but string range makes that quite simple:

```
1: when HTTP_REQUEST {
2:   set user [string range [HTTP::uri] 12 end]
3:   log local0. "User: $user"
4: }
```

string map

```
string map mapping string
```

Where string range allows you to select a given part of a string and return it, string map allows you to actually modify sub strings in-line. Also, instead of acting on a count or range of characters, string map works with an actual string of characters. Whereas with string range you may want to look up a particular part of a URI, such as the first 10 characters, and see if they match a string, or route based on them or…something; with string map you are able to make changes in real-time, changing one string of characters to another.

For instance with string range you may have a logic statement like "Do the first 10 characters of the URI match x". You'd supply the string to fetch the range from and the number of characters you want, by giving a beginning and end character. With string map you'd be saying something like "look for any string that looks like x, and change it to y in the given string" by providing the string to work against as well as a source and destination string, meaning "Change all cases of http to https".

```
1: when HTTP_RESPONSE {
2:   set new_uri [string map {http https} [HTTP::header "Location"]]
3:   HTTP::header replace Location $new_uri
4: }
```

string first

```
string first string1 string2 ?startIndex?
```

The string first command allows you to identify the first occurrence of a given sub string of characters within a string. This can be extremely useful for combining with the string range command. For instance, if I want to find the first occurrence of "/admin" in a URI and collect the URI from that point to the end, it would be quite difficult without the string first command. What if I don't know what the exact URI will be? What if there is a variable portion of the URI that comes before "/admin" that I don't want to collect, but have to somehow account for even though it is variable in length? I can't just set a static range and use the string range command alone, so I have to get creative and combine commands.

By making use of the string first command, if I have a URI that looks something like "/users/apps/bob/bobsapp? user=admin" where the username is variable length and I can't be certain of the length of the URI because of it, but I wanted to retrieve the user argument being passed in, I could do something like:

```
1: set user [string range [HTTP::uri] [expr {[string first "user=" [HTTP::uri]] + 5}] end]
```

What the above is doing is finding the first occurance of "user=" in the URI and returning the index of the first character. Then adding 5 to that, since that is the length of the string "user=", and we want to reference what comes after "user=", not include it, then take the range of the string from that point to the end, and return that as the value of the username being passed in. It looks a bit complex, but if you break it down command by command, you're really just stringing together several smaller commands to get the functionality you want. And now you can start to see why string commands are so important and powerful in iRules.

string trim

```
string trim string ?chars?
```

Again, a command true to its name, the string trim command allows you to manipulate the beginning and end of a given string to trim off unwanted characters. Perhaps you have a string that you need to ensure doesn't begin or end with white space, or you're looking at URI comparisons and need to be sure that you don't have a trailing slash in some cases and not others. Regardless, the string trim command makes that easy. All you'd do is specify which characters you want to ensure are removed, either whitespace or slashes in the examples just mentioned, and you'd be set to ensure standardized comparisons. For instance, if you want to ensure that someone is making a request to the admin section of your domain you can be sure that they are going to have a slash at the beginning of the URI, but they may or may not include a trailing slash. You could use the starts_with comparison operator to ignore this, but that would also ignore anything else they may include after the specific string. If you want an exact match for either "/admin" or "/admin/" you could use an or, or could trim the URI, like so:

```
1: when HTTP_REQUEST {
2:   if{[string trim [HTTP::uri] "/"] eq "admin"} {
3:     pool admin_pool
4:   }
5: }
```

Note that by using the trim command it removed both the preceding and trailing slashes. There are trimright and trimleft versions as well, for only trimming one end of a string, if that's necessary.

What commands besides the Tcl "string" command would be useful for strings?

While there is a massive amount that can be done with the string command and the many subcommands thereof, there are some limitations. Some functional, others performance, that should be accounted for. In some situations you may need more power than a simple string command, or even a set of string commands. If that is the case then you're looking at either the scan command, or using regular expressions. I won't cover them in detail here, as there's just too much and it gets advanced very quickly, but I'll give a brief overview.

scan

The scan command also has many permutations and a near endless set of possibilities. It is insanely high performance, but almost equally as unintuitive. It is not for the feint of heart to begin tackling the advanced end of the pool with the scan command's options, but the easier things are more palatable, and can serve very useful purposes. We encourage people looking to do slightly more advanced string matching or mapping to look at scan as an alternative to multiple string commands, and for those things that string commands just can't accomplish. It is very fast, very efficient code-wise, and offers a wide range of possibilities from string conversion to mapping and more. You can read more about the scan command here (http://www.tcl.tk/man/tcl8.4/TclCmd/scan.htm) for official documentation and here (https://devcentral.f5.com/Tutorials/TechTips/tabid/63/articleType/ArticleView/articleId/197/Scan--Making-string-manipulation-efficient.aspx) for some more DC info.

regex

Regular expressions are also available to you within iRules. They are amazingly powerful, as is always the case with regex style matching, even though lookbacks are disabled for efficiency in this implementation. They have amazing potential and are highly portable, as regex syntax is fairly ubiquitous in the *Nix world whether in programming or otherwise. Unfortunately they are also amazingly inefficient, as far as iRules commands go. Regular expressions are the single most expensive operation you can perform in an iRule, suspending commands aside. Firing up the regex interpreter alone runs about the same as the cost of 10 standard string commands, and that's before you've done any actual comparisons. For this reason we tend to encourage people to only use regular expressions when necessary, and to seek out other solutions where it's possible.