

Content Switch Rules and Their Conflict Detection

C. Edward Chow, Dept. of CS, Univ. of Colorado at Colorado Springs, Colorado Springs, CO 80933-7150, USA
Ganesh Kumar Godavari, Dept. of CS, Univ. of Colorado at Colorado Springs, Colorado Springs, CO 80933-7150
Jianhua Xie, Dept. of CS, Univ. of Colorado at Colorado Springs, Colorado Springs, CO 80933-7150, USA

Abstract

Content switch can be configured as load balancer, firewall, spam mail filter, and virus detection and removal system, by specifying a set of rules. In this paper we present our content switch rule design for a Linux-based content switch, and show how they are translated and downloaded into the switch for packet processing. One common problem in specifying rules for content switches or policy-based networks is the conflict detection problem, where rules may contradict or duplicate each other. We explore the algorithm issues related to conflict detection problem and present some performance results. An algorithm for detecting conflict in rules with regular expression matching is presented and its complexity analyzed. An interactive Java-based content switch rule editor was designed for specifying the rules and detecting the conflicts among rules. Its conflict detection algorithm and performance results are also presented.

Keywords: Content Switch, Policy-based Networking, Rule Conflict Detection

1 Introduction

With the rapid increase of Internet traffic, we have found load balancing clusters are used for improving the performance of server farm. The content switches or web switches [1,2] are new class of network devices that can be served as a front end for such server cluster. They provide load balancing service by distributing the requests based on the headers and content of the upper layers (3-7). They provide high availability by checking the server status through monitoring messages and re-routing the requests from failed servers to active servers. By rejecting packets based on their source or destination IP addresses and port numbers, they can be configured as firewall. Since they can check the upper layer information, they greatly extend the firewall coverage cases and flexibility. By putting the packets on different queues after examining their headers and content, they can be configured as policy-network devices that control the bandwidth/resource allocation.

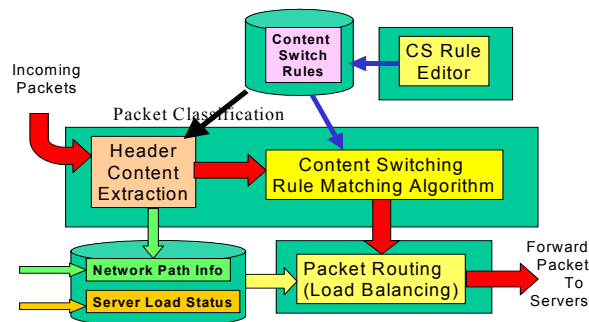


Figure 1. Content Switch Rule and Content Switch

The central mechanism of the content switch is the content switch rule. The rule specifies what class of packets it tries to match and what action should be applied to the packet. Each incoming packet is checked against a set of content switching rules. The set of content switch rules specifies the behavior of the content switch. Figure 1 shows the content switch rule and its relationship with other modules in a content switch. The administrator of the content switch use a content switch (CS) rule editor to create the content switching rules. The content switch rule editor can be a simple text editor or an enhanced editor with GUI and built-in with conflict detection function for checking the conflicts within the rule set. Examples of conflicts include duplication, two rules with the same condition but different routing actions, two rules with conditions that intersect, two rules with conditions that are subset and in improper order.

This paper is organized as follows: Section 2 presents the content switch rule design and shows how the rules can be translated and downloaded into the switch as a module. Section 3 discusses the conflict detection problems

for rules with regular expression. Section 4 shows the design of an interactive Java-based rule editor with built-in rule conflict detection. The performance results of its rule conflict detection algorithm were presented. Section 5 is the conclusion.

2 Content Switch Rule Design

Content switching rules are typically expressed in terms of content pattern or conditions that cover the class of packets to be matched, and its associated action. In the existing content switch products, there are two basic approaches that rules are specified:

1. The rules are entered using the command line interface. The syntax are typically similar to that of CISCO ACL (access control list) convention [3,4].
2. Using a special language to specify the pattern and describe the action of the service. The rule set is then translated and downloaded into the content switch [5].

An example of approach 1 can be seen in Cisco Content Engine (CE) 2.20 [3]. For example, Cisco CE can support HTTP and HTTPS proxy server with the rule such as

```
rule no-cache url-regex.*cgi-bin.*
```

This rule specifies that the incoming packets with the url matching the regular expression pattern “*cgi-bin*” will not be forwarded to the cache servers. The Foundry ServIron Installation and Configuration Guide [4] provides an excellent collection of rules and web switch configuration examples.

An example of approach 2 is Intel IX-API SDK[5]. It uses network classification language (NCL) to classify the incoming traffic and describe the action on the packet. The rule syntax is presented as

```
Rule <name of the rule> {predicate} {action_method() }
```

The predicate part of the rule is a Boolean expression that describes the conditions. When a packet matches the condition, it will trigger the specified action to be performed. The action part of the rule is the name of an action function to be executed when the predicate is true, and it performs some operations upon the incoming packet. For example,

```
Rule check_src {IP.src==10.10.10.30} {action_A() }
```

The meaning of this rule is that if source IP address is 10.10.10.30, then the action function action_A() is executed

2.1 Content Switch Rule

Our content switch rule follows an approach similar to Approach 2. The rules are defined using C functions. The syntax of the defined rules is as follows:

```
RuleLabel: if (condition) { action1} [else { action2}].
```

Here are as a set of legal content switching rules. We use them to explain the content switch rule design.

```
R1: if (match(url, "purchase.pl") && xml.purchase/totalAmount > 50000) {  
    routeTo(highSpeedServers, STICKY_ON_IP_PORT); }  
R2: if (strcmp(xml.purchase/customerName, "CCL") == 0) {  
    routeTo(specialServer, NONSTICKY); }  
R3: if (match(url, "mid$") == 0) { routeTo(midiServer, NONSTICKY); }  
R4: if (match(smtp.from, "spam.com")) { discard(STICKY_ON_IP); }  
R5: if (match(smtp.to, "chow@cs.uccs.edu")) { routTo(mailServer1, STICKY_ON_IP_PORT); }  
R6: if (match(imap.login "chow")) { routeTo(mailServer1, STICKY_ON_IP_PORT); }  
R7: if (inSubnet(srcip, "128.198.60.0/24") && dstip == s2ip("128.198.192.192") &&  
    dstport == 80) { routeTo(LBServerGroup, STICKY_ON_IP_PORT); }  
R8: if (match(url, "xmlprocess.pl")) { goto R9; }  
R9: if (xml.purchase/totalAmount > 5000){routeTo(hsServers, NONSTICKY);}  
    else {routeTo(defaultServers, NONSTICKY); }
```

The rule label allows the use of goto for branching, and make referencing of rules easier. The condition of the rule are expressed in terms of Boolean expressions with relational expressions, or Boolean functions such as match(string, regexPattern) and inSubnet(), as basic terms. The relational expressions are expressed in terms of simple *variable <relational operator> value* form. Here the variable can be fields of headers, substring in the content, XML tag sequences, or a supported string functions such as strcmp() function.

In Rule R1, we have two terms in its condition. One is match(url, “purchase.pl”), which covers all packets that are to be processed by purchase.pl server side CGI script. We have implemented match(string, regexPattern) function for regular expression matching in content switching rule module. It returns true when the string contains the regular expression pattern specified in the second parameter. url is a reserved word for the absolute path parameter which appears right after the http request command. The other term is xml.purchase/totalAmount >

50000, which covers the packets that contain a purchase request in XML and contains totalAmount tag under the root tag purchase. The string value wrapped around by the <totalAmount> tag denotes the xml.purchase/totalAmount. Note that here we expect the string value to be a legal numeric value and can be converted. If the value is not numeric, an error message will return to the sender.

Rule R1 contains routeTo (highSpeedServers, STICKY_ON_IP_PORT). HighSpeedServers is the name representing a cluster of fast servers. The user can specify what servers constitute the highSpeedServers cluster and what load balancing algorithm is used to select one of them for serving the request. STICY_ON_IP_PORT option specifies that the subsequent packets with the same source IP address and source port number will be forwarded to the same selected server. The connection is called a *sticky connection* and is to be sticky on IP address and port number.

```
R2: if (strcmp(xml.purchase/customerName, "CCL") == 0) {
    routeTo(specialServer, NONSTICKY); }
```

Here we use strcmp() to compare the string value of XML tag sequence purchase/customerName with that of "CCL". Any packet with purchase/customerName tag value equals to "CCL" will be forwarded to specialServer without regarding its url or other header values.

```
R3: if (match(url, "mid$") == 0) { routeTo(midiServer, NONSTICKY); }
```

Here any HTTP request packet with file extension mid will be routed to midiServer. The \$ indicates the mid must appear at the end of the url string. This rule illustrates that we can spread the server load based on the media type of their request documents.

```
R4: if (match(smtp.from, "spam.com")) { discard(STICKY_ON_IP); }
```

Here we demonstrate the rule for configuring the content switch to discard packet from a specific domain and actually any subsequent request from the same node will be discarded. smtp is the reserved word for the Simple Mail Transfer Protocol. From is a reserved word for the from header field of the email. Any packet with the email address containing "spam.com" will be covered by this rule. If we would still like the spam node to access others of our services, we can change the option to STICKY_ON_IP_PORT.

```
R5: if (match(smtp.to, "chow")) { routeTo(mailServer1, STICKY_ON_IP_PORT); }
```

```
R6: if (match(imap.login, "chow")) { routeTo(mailServer1, STICKY_ON_IP_PORT); }
```

Here we demonstrate how to route the mails of a user to a specific mail server. IMAP protocol is used to retrieve emails. Login is a reserved word for the login field in the IMAP request.

With a slight modification, the following six rules will spread the load of mail services to three servers.

```
R5a: if (match(smtp.to, "[A-I,a-i]")) { routeTo(mailServer1, STICKY_ON_IP_PORT); }
```

```
R6a: if (match(imap.login, "[A-I,a-i]")) { routeTo(mailServer1, STICKY_ON_IP_PORT); }
```

```
R5j: if (match(smtp.to, "[J-R,j-r]")) { routeTo(mailServer2, STICKY_ON_IP_PORT); }
```

```
R6j: if (match(imap.login, "[J-R,j-r]")) { routeTo(mailServer2, STICKY_ON_IP_PORT); }
```

```
R5s: if (match(smtp.to, "[S-Z,s-z]")) { routeTo(mailServer3, STICKY_ON_IP_PORT); }
```

```
R6s: if (match(imap.login, "[S-Z,s-z]")) { routeTo(mailServer3, STICKY_ON_IP_PORT); }
```

```
R7: if (inSubnet(srcip, "128.198.60.0/24") && dstip == s2ip("128.198.192.192") &&
    dstport == 80) { routeTo(LBServerGroup, STICKY_ON_IP_PORT); }
```

In Rule R7, Srcip, dstip, srport, dstport are reserved words representing the fields in the headers of protocols. Since we want to cover packets from a subnet, an Boolean function inSubnet() for checking if the net address of the IP address of the packet falls in a specific one. A s2ip() function is used to convert IP address in dot notation to the equivalent 32 bit value.

```
R8: if (match(url, "xmlprocess.pl")) { goto R9; }
```

```
R9: if (xml.purchase/totalAmount > 5000) { routeTo(hsServers, NONSTICKY); }
    else { routeTo(defaultServers, NONSTICKY); }
```

Rule R8 demonstrates the use of goto. Rule R9 demonstrates the use of else branches.

2.2 The Rule action and syntax of the sticky (non-sticky) connections.

In our Linux-based Content Switch prototype (LCS), there are three different options related to the sticky connections. These rules are based on the different content of the packets.

1. Option for sticky connection based on the source IP address.

Example: `If(source_ip==128.198.192.194) { routeTo(server2, STICKY_ON_IP); }`

The condition of this rule is related to the source IP address. The action inside routeTo() will assign the real server server2 to the connection, and add this connection to the sticky connection database. When the new connection comes, the rule matching process will look for the data entry with the same IP address in sticky

database first, if the data entry is found, the connection will be routed to the same server directly without carrying out the rule matching.

- Option for sticky connection based on source IP address and TCP port number.

Example: *If*((source_ip==128.198.192.194)&&(source_port==9872)) {
 routeTo(server4, STICKY_ON_IP_PORT);}

The condition of this rule includes the source IP and the port number. This rule is for multiple requests in one TCP keep alive connection. The action process will add this entry to the keep-alive connection hash table using the IP address and port number as the hash key. If the new request arrives from the same connection, the request will be routed to the same server without rule matching.

- The option for non-sticky connection.

Example: *If*(URL=="*jpg") { *RouteTo*(imageServer, NON_STICKY);}

This rule specifies the connection to be non-sticky connection. So either the request from the same connection or the new connection all need to go through the rule matching process to select the real server.

2.3 Translation and Update of the Content Switch Rule Set

The CS rules can be specified by a text editor. The rule set file will be translated by a ruleTranslate program into a data structure containing XML tag sequences and a function with translated if statements of the rule set. The data structure and the rule function form the basis of a rule module. To update to a new rule set, *rmmod* command is first executed to remove the old rule set module from the kernel. Before the new rule set is installed, the content switch schedule control module will call a default function NO_CS(). Next, the rule module is inserted using *insmod* command. The content switch is then switched to use the new rule set. To speed up the rule matching, we also replace the domain names and sever names in rules with their IP addresses in the translation process.

3 Conflict Detection Problem for Rules with Regular Expressions

Recent work [7] analyzed extensively the conflict types in distributed systems management and the conflict detection study on IPSec/VPN security policies was presented in [8]. Even though we focus on detecting rule conflict in content switches. Our results can be applied in those areas. When the content switch rules are defined, unintentionally conflicts may be introduced into the rule set. For example, if there are two rules in the rule set:

```
if ( match(url, "^*.gif$) {routeTo(Server1); }
if ( match(url, "^a*.gif$) {routeTo(server2);}
```

The second rule will never be executed, because the second regular express is a subset of the first one. A useful content switch rule editor should detect this type of errors and prompt the administrator with informative messages. To detect this kind of conflicts, one needs to know the relationship between two regular expressions. There are five potential relations between two regular expressions, say, *reg1* and *reg2*:

| | |
|-----------------------------------|------------------------------------|
| <i>reg1</i> = <i>reg2</i> | <i>reg1</i> equals to <i>reg2</i> |
| <i>reg1</i> \supset <i>reg2</i> | <i>reg1</i> contains <i>reg2</i> |
| <i>reg1</i> \subset <i>reg2</i> | <i>reg2</i> contains <i>reg1</i> |
| $(reg1 \cap reg2) \neq \phi$ | <i>reg1</i> intersects <i>reg2</i> |
| $(reg1 \cap reg2) = \phi$ | <i>reg1</i> disjoins <i>reg2</i> |

For detecting content switch rule conflict, we need only detect the first three cases. Because “x contains y” is symmetric to “y contains x” and “x equals y” is the same as “x contains y and y contains x”, so we need only design an algorithm to detect if “x contains y”, then all the problems are solved.

3.1 Algorithm for detecting “x contains y”

Definition 1: The relation “assignable to” between two characters *x* and *y* is defined as:

| | |
|---------------------------------------|---|
| <i>c</i> assignable to <i>c</i> | for any character or meta-character <i>c</i> |
| <i>c</i> assignable to ? | for any character or range character([<i>x-y</i>]) <i>c</i> |
| <i>c</i> assignable to * | for any character or meta-character or empty string <i>c</i> |
| <i>c</i> assignable to [<i>x-y</i>] | for any character or the range character within the range from <i>x</i> to <i>y</i> |

If *x* is assignable to *y*, we write it as: $x \rightarrow y$.

3.2 Definition 2: Single size character: all characters and meta-characters except *

With the definitions above, we have the following algorithm.

3.3 Algorithm: Detecting if one regular expression contains another

Input: regular expression *reg1* and *reg2*

Output: true if $reg1 \supset reg2$, false otherwise

Method:

1. $p1=0, p2=length\ of\ reg1, q1=0, q2=length\ of\ reg2$
2. if $q1\ equals\ q2$ and $p1\ equals\ p2$, finish and return true.
3. If $reg1[p1]$ is a single size character and $reg2[q1] \rightarrow reg1[p1]$, then let $p1 = p1+1$ and $q1 = q1+1$, goto step 2
4. for t from $q1+1$ to $q2$, execute the algorithm with $p1+1, p2, t, q2$

The pseudo-code of the algorithm is as follows:

```
contains(string reg1, string reg2, int p1, int p2, int q1, int q2){
    if (q1 == q2) {
        if (p1 == p2) return true;
        else return false
    }
    if (X[p1] is a single size character) {
        if (! Y[q1] → X[p1]) return false;
        else return contains(p1+1, p2, q1+1, q2);
    }
    else {
        for (t=q1; t<=q2; t++)
            if (contains(p1+1, p2, t, q2)) return true;
        return false;
    }
}
```

3.4 Complexity analysis

In the above algorithm, the recursion occurs only when the character to be compared is asterisk. So in the worst case, i.e., half of the characters in the reg1 are asterisk, the complexity is $o(n^n)$, where n is the length of the regular expression. But in the use of the content switch rule editing, the number of asterisks in a regular expression is typically quite small, normally less than 3, and the length of regular expression is also very limited, normally less than 15 characters. For example, if there are 2 asterisks in the regular expression, then the complexity is $o(n^3)$, which is good enough for cases where n is less than 15.

4 Java-based Content Switch Rule Editor

We have implemented an interactive Java-based Content Switch Rule Editor, RuleEdit, for editing the content switch rule set and detecting the conflicts between the rule being specified and the existing rule set. Figure 2 shows the screendump of RuleEdit. It can open the text file of an existing rule set, and load the rule set into its internal structure. It currently only implements rule conditions with one basic term. The user can enter the variable, the relational operator, and the value of a relational expression, or a Boolean expression. On the right hand side is a set of buttons for inserting or appending the rule to the rule set based on the line number, and for deleting certain rule. The text window below shows the rules in the existing rule set.

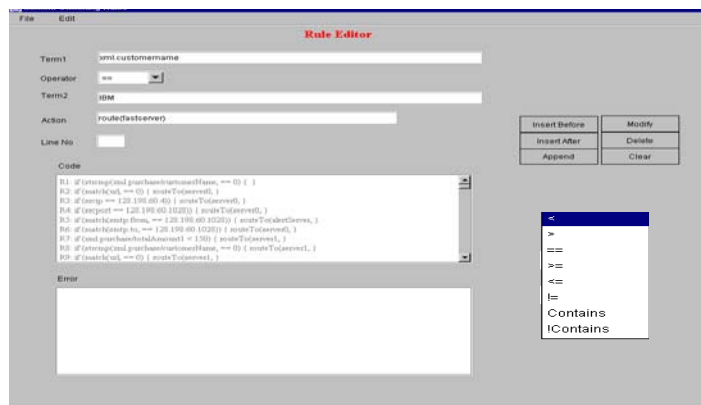


Figure 2. Java-based Interactive Content Switch Rule Editor




Figure 3. Time Performance of Rule Conflict Detection algorithm.

The operator List box contains various options. When the contents of term2 are not numeric, the rule editor

automatically converts the condition into a strcmp statement. The options “contains” and “!contains” are used for handling regular expressions, e.g. `if (term1 contains term2) {sequence of actions to be performed}`
Here term2 can be a regular expression like “*.gif”.

When the rules are entered in the rule editor. The rule editor outputs a flag whenever a potential conflict is detected. Potential conflicts can be

- 1) Duplication of the condition i.e. for same condition specifying same/different actions.
- 2) Numerical comparison errors, these are easily committed and difficult to detect manually.

For example,

```
Rule111: (xml.purchase/totalAmount > 5000) { routeTo( miniServer, NONSTICKY);}  
Rule122: if (xml.purchase/totalAmount > 20000) { routeTo(crayServer, NOSTICKY);}
```

When rule matching occurs and `xml.purchase/totalAmount = 30000`, Rule111 will be executed first rather than Rule122, which may not be the intention of the user.

Similar numerical comparison conflict can occur with “<” operator. The Rule Editor will help detect these potential conflicts. Duplication Errors can easily be detected by checking the condition of the current rule with the existing conditions in the rule set. Numeric Comparison Errors are detected using the following algorithm.

4.1 Conflict Rule Detection Algorithm

When a new rule is being entered,

- 1) Check the logical operator, if it is of not “>”, “<” goto step 6.
- 2) Check the existing rules whether the logical operator and term1 match, if no match is found goto step 6.
- 3) If the logical operator is “>”, check whether `match.term2 < term2` and `match.ruleposition < ruleposition`. If true, flag the user of potential conflict and goto step 6.
- 4) If the logical operator is “<”, check whether `match.term2 > term2` and `match.ruleposition < ruleposition`. If true, flag the user of potential conflict and goto step 6.
- 5) Accept the rule to the rule set.
- 6) Exit.

4.2 Performance Results of the Rule Conflict Detection Algorithm

To understand the performance, a Perl script was written to create rule sets of varying sizes and load them in RuleEdit and compute the time it takes for the RuleEdit to finish the conflict detection. Figure 3 shows the performance results. It takes about 1.6 minutes to finish the comparison between one rule and an existing 8000 rules. Below 2000 rules, the user did not see noticeable delay.

4.3 Features and Current Limitations of RuleEdit

The RuleEdit can inform the end user of a potential conflict in if statements, which will be tiring to debug manually. It automatically converts string data type into strcmp format, the native form of c string handling. It also handles regular expression comparisons. Currently it handles only one term in the condition. We are working on the multiple term cases.

5 Conclusion

We have presented a flexible content switch rule design for specifying the operations of content switches and demonstrated rules for various configurations. The rule conflict detection problem and their algorithms are presented together with the performance and design of an interactive Java-based Rule Editor.

6 References

- [1] George Apostolopoulos, David Aubespin, Vinod Peris, Prashant Pradhan, Debanjan Saha, “ Design, Implementation and Performance of a Content-Based Switch”, Proc. Infocom2000, Tel Aviv, March 26 - 30, 2000, <http://www.ieee-infocom.org/2000/papers/440.ps>
- [2] Gregory Yerxa and James Hutchinson, “Web Content Switching”, <http://www.networkcomputing.com>.
- [3] “Release Notes for Cisco Content Engine Software”. <http://www.cisco.com>”.
- [4] “Foundry ServIron Installation and Configuration Guide,” May 2000. <http://www.foundrynetworks.com/techdocs/SI/index.html>
- [5] “Intel IXA API SDK 4.0 for Intel PA 100,” <http://www.intel.com/design/network/products/software/ixapi.htm> and http://www.intel.com/design/ixa/whitepapers/ixa.htm#IXA_SDK.
- [6] Emil C. Lupu and Morris Sloman, “Conflicts in Policy-Based Distributed Systems Management,” IEEE Trans. On Software Engineering, Vol 25, No. 6, Nov/Dec 1999, pp. 852-869.
- [7] Z. Fu, S. Felix Wu, H. Huang, K. Loh, F. Gong, I. Baldine, and C. Xu, “IPSec/VPN Security Policy: Correctness, Conflict Detection and Resolution,” to appear in Policy 2001 conference.