



Monitor · Manage · Control

ContentKeeper Policy Logic And Implementation Requirements

Overview

Understanding the rules for creating customised policies and how the various types of policies affect each other is an essential skill for any ContentKeeper administrator. This paper will help you to develop that skill and therefore maximise ContentKeepers productivity.

Document Revision A

Date: 1st July 2003

Copyright © 2000, 2001,2002, 2003 ContentKeeper Technologies

ContentKeeper® Closed Loop Collaborative Filtering™ and *TrickleFeed™* are trademarks of ContentKeeper Technologies. Copyright © 2000 - 2003, ContentKeeper Technologies, Canberra, Australia. All Rights Reserved.

Linux is a registered trademark of Linus Torvalds, Red Hat Linux is a registered trademark of Red Hat Inc.

All other product names mentioned herein are trademarks or registered trademarks of their respective owners.

The products and specifications, configurations, and other technical information regarding the products contained in this manual are subject to change without notice. All statements, technical information, and recommendations contained in this manual are believed to be accurate and reliable but are presented without warranty of any kind, express or implied, and users must take full responsibility for their application of any products specified in this manual.

Document Author & Designer: Matthew R Richards

ContentKeeper Technologies
218 Northbourne Avenue
Braddon ACT 2612
Australia
PH +61-2-62614950
Fax +61-2-62579801
info@ContentKeeper.com
www.ContentKeeper.com

Policy Logic And Implementation Requirements

This document explains the logic behind ContentKeeper policy processing and describes the requirements and rules of implementing policies within ContentKeeper.

Table of Contents

CONTENTKEEPER POLICY IMPLEMENTATION REQUIREMENTS	3
CONTENTKEEPER POLICY PROCESSING LOGIC	4
(1)GLOBAL PRIORITIES	5
(2)POLICY CHOICE PRIORITIES	5
<i>Table 1. Cumulative Policy Application</i>	5
(3)POLICY COMPONENT PRIORITIES	6
(4)CATEGORY PRIORITIES	6
<i>ContentKeeper Special Categories</i>	6
CREATING, VERIFYING AND TROUBLESHOOTING POLICIES.....	7
CREATING POLICIES	7
VERIFYING POLICIES	7
TROUBLESHOOTING POLICIES	7
<i>Manual Policy-Verification Chart</i>	8
Example 1 – Using the Verification Chart	8

ContentKeeper Policy Implementation Requirements

Policies rely on ContentKeepers ability to accurately determine user and computer identities.

Network users may be identified in two ways, user credentials (which include a user name and a group name) and computer credentials (IP addresses and canonical names).

User credentials may be ascertained in the following ways:

- Extracted from the header of a http-get request [which has been sent to a proxy server that requires authentication]
- Retrieved via a NetBIOS lookup request

Computer credentials may be ascertained in the following ways:

- Retrieved via a NetBIOS lookup request
- Retrieved via a DNS lookup request
- Extracted from the header of a http-get request

Once credentials have been ascertained, they must be confirmed and matched to a policy. This may be achieved in varying ways, depending on how ContentKeeper is configured.

ContentKeeper may attempt to match credentials against one or more of the following:

- IP address and subnet mask.
- Username and/or group from local database.
- Username and/or group from a Windows NT/2000 security group.

Both the IP address / subnet mask and local database data sets are stored locally on the ContentKeeper server. The ContentKeeper Agent supplies ContentKeeper with Windows NT/2000 security group data from a Windows NT/2000 domain controller.

Refer to the sections “*Username Resolution*” and “*The ContentKeeper Agent*” in Part 3 “*Configuring ContentKeeper*” of the ContentKeeper Administration Guide for further information.

Once ContentKeeper has confirmed and matched the user computer credentials it will apply the appropriate policy.

When implementing ContentKeeper policies, the following requirements apply to policy contents and input fields:

1. Each unique username may exist in only one policy when entered into the *username* field.
2. Each unique username may exist in multiple policies when entered via the *group* field.
3. Each IP address and 32-bit subnet mask combination may exist in only one policy. An example of an IP address and 32-bit subnet mask combination is: 203.5.115.24 / 255.255.255.255
4. When a policy is created, all upper case letters will be converted to lower case in the name, description and object fields.

ContentKeeper Policy Processing Logic

A URL request may qualify for filtering if either it's source IP address or the embedded user name match one of those in a policy. All other requests will be filtered through the default policy, with the exception of those whose source IP addresses are specifically set to exclude in the *Excluded/Included IP Addresses* menu.

If a URL request qualifies for filtering through at least one policy other than the default, then none of the default policy settings are applied when filtering takes place.

A URL request may qualify for filtering through multiple policies when either its source IP address or the embedded user name exists in more than one policy.

When a URL request qualifies for filtering through multiple policies, the least restrictive policy settings apply. ContentKeeper decides how and whether or not to filter a URL request based on the priorities below.

These priorities are processed in the order that they are presented here. If a URL request does not match the first priority the ContentKeeper will move on and process the request against the second, third and fourth priorities until a match is found.

(1)Global Priorities

1. Requests from excluded IP addresses are not filtered.
2. If the default only check box is enabled then only the default policy is applied. In this case there are no custom policies defined therefore all requests that are not made by an excluded IP match the Default Policy and processing continues with group (3) Policy Component Priorities.

(2)Policy Choice Priorities

1. Any policy containing an IP address and 32-bit subnet mask that match the IP address of the request is applied exclusively. This means that with the exception of the Global Policy, settings are only applied from the matching policy.
2. Any policy containing a username that matches the username embedded in the request is applied exclusively. This means that with the exception of the Global Policy, settings are only applied from the matching policy.
3. Any policies with groups that contain a username that matches the username embedded in the request are applied collectively. This means that settings from all matching policies are applied in a least restrictive manner, i.e. the opposite order to which they appear in group (4) Category Priorities. Refer to the Table 1. below.

Table 1. Cumulative Policy Application

This table demonstrates how policies are applied cumulatively, such as when a user belongs to a group and that group has been associated with two policies.

	Policy 1				Policy 2			Cumulative Result
1	Adult Content	Block		1	Adult Content	Block		Block
2	News	Coach		2	News	Allow		Allow
3	Job Search	Authenticate		3	Job Search	Coach		Coach
4	Gambling	Block		4	Gambling	Block		Block
5	Travel/Tourism	Block		5	Travel/Tourism	Allow		Allow

(3)Policy Component Priorities

1. Global Policy settings.

If the global Policy is enabled, any of it's settings not set to *Ignore* take priority over corresponding settings in all other policies. Within the global policy, policy components are prioritised in the following order:

- a. Custom and Trusted URLs
- b. Category States
- c. Custom File Types

All other policy components are prioritised in the following order:

2. Custom and Trusted URLs
3. Category States
4. Custom File Types

(4)Category Priorities

Category priorities are an exception to the *Least Restrictive* rule. Category priorities are applied in Most Restrictive order based on category state.

1. Block
2. Authenticate
3. Coach
4. Allow

There are two additional states, Time of Day and Block Discard. The Block Discard state applies to individual URLs and has the same priority as the Block state. The Time of Day state equates to, and is processed with the same priority as, one of the four main states, based on the current time and day.

ContentKeeper Special Categories

ContentKeeper processes the *Educational*, *News*, *Search Sites* and *Business Oriented* categories differently to its other categories. The blocking rules associated with each of these categories are outlined below.

Educational, News & Search Sites – If a URL exists in either the *News* or *Search Sites* categories and any of these categories is set to allow, then it will not be blocked under the other category. This is to cope with the high level of similarity between these categories.

Business Oriented – This category has been added to help ensure that business sites are correctly classified. URLs in this category are subject to two default rules:

1. URLs in this category are excluded from the real-time analysis engine.
2. If a URL exists in **one** of the *Business Oriented*, *News* or *Search Sites* categories and any of these categories is set to allow, then it will not be blocked under the remaining two categories regardless of whether or not the URL is classified under all of these categories.

Creating, Verifying and Troubleshooting Policies

Creating Policies

Creating Policies is covered in detail in the ContentKeeper Administration Guide. There is also a manual verification chart included at the end of this document to help you construct and understand new and existing policies.

Refer to the section “*Creating Policies*” in Part 3 “*Configuring ContentKeeper*” of the ContentKeeper Administration Guide for further information.

Verifying Policies

Verifying Policies is covered in detail in the ContentKeeper Administration Guide. There is also a manual verification chart included at the end of this document to help you construct and understand new and existing policies.

Refer to the section “*Verifying Policies*” in Part 3 “*Configuring ContentKeeper*” of the ContentKeeper Administration Guide for further information.

Troubleshooting Policies

While the underlying logic is somewhat complex, on the surface ContentKeeper policies are relatively simple! Although, you may occasionally wish to understand why a URL is or is not being blocked, or have need to carefully plan the construction of a new policy.

The best tool for helping you understand the effects of your policies is the ContentKeeper Policy Verification Facility. It will allow you to test a policy on its own or in conjunction with all other policies, and in any possible scenario. Refer to the section above titled *Verifying Policies* for more information.

When ContentKeeper generates a block page, information about what, why and who caused the block page is displayed in a table in the middle of the page. This information is repeated and expanded upon in the ContentKeeper Blocking Activity reports and logs. Both sources of information can be useful when troubleshooting policies.

Determining if a URL has been categorised and which category it belongs to is essential to understanding the effect of a policy. The *Block/Unblock URLs in Control List* facility allows you to quickly and accurately determine whether a URL has been categorised and which category it belongs to.

Refer to the section “*Block/Unblock URLs in Control List*” in Part 3 “*Configuring ContentKeeper*” of the ContentKeeper Administration Guide for further information.

Manual Policy-Verification Chart

Another useful tool for helping you to construct or decipher policies has been included with this paper. It is the manual verification chart below. Use this chart in conjunction with the information presented on the four Priority Groups under the *Policy Processing Logic* section above.

1. To Use the manual verification chart to create or understand a policy you will need the following information:
 - A URL or Content Type.
 - Knowledge of ContentKeepers current *Username Resolution* method.
 - An IP address.
 - A username.
2. Follow the chart from left to right to understand how a policy works.
3. Refer to the information about the corresponding Priority Group at the beginning of each step.
4. Apply the URL and username information at the appropriate junctures.

Important: This chart has been designed specifically for use with the information presented on the four Priority Groups. It is essential that you refer to the information about the corresponding Priority Group at the beginning of each step in the chart as using the chart alone may not lead to accurate results.

Example 1 – Using the Verification Chart

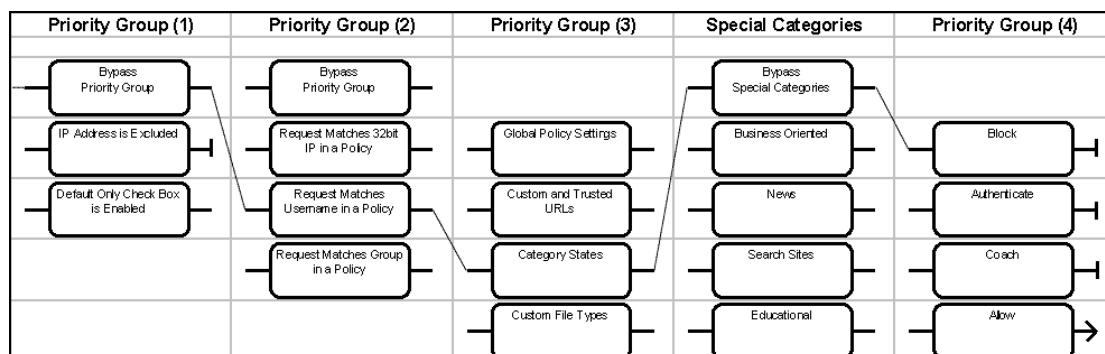
In the following example an administrator is attempting to determine exactly what steps lead up to the URL www.playboy.com being blocked and if those steps conform to the intended function of the governing policy. The administrator has obtained the following information:

URL: www.playboy.com

Username Resolution Method: Proxy Header, Basic Authentication Scheme

IP Address: 10.10.0.45/255.255.255.0

Username: John



Step 1: Examine Priority Group 1.

First the Administrator determines that the Requesting IP address is not in the ContentKeeper Excluded/Included IP Addresses List, the administrator also knows that there are multiple customised policies.

Result: *Bypass Priority Group* is selected.

Step 2: Examine Priority Group 2.

The administrator knows that each individual in the company has their own policy.

Result: *Request Matches Username in a Policy* is selected.

Step 3: Examine Priority Group 3.

The ContentKeeper blocking screen had reported that the URL was List Blocked. This means that the URL was not a Custom URL, nor was it blocked based on File Type settings. The administrator also knows that the Global Policy is not enabled.

Result: *Category States* is selected.

Step 4: Examine Special Categories.

By using the ContentKeeper *Block/Unblock URLs in Control List* facility the administrator is able to determine that the URL does not belong to one of the special categories.

Result: *Bypass Special Categories* is selected.

Step 5: Examine Priority Group 4.

By examining the policy responsible for causing the block page the administrator is able to determine that one of the categories to which the URL belongs is set to block.

Result: *Block* is selected.

Conclusion:

The administrator now has a clear roadmap of the how the policy has functioned and interacted with other policies. The administrator has been able to determine if the configuration of the policy has been successful and make any necessary adjustments.

Using the manual verification chart in conjunction with the information presented on the four Priority Groups under the *Policy Processing Logic* section will enable you to verify a policy step by step, component by component; identifying and correcting any mis-configured policy components.

Priority Group (1)	Priority Group (2)	Priority Group (3)	Special Categories	Priority Group (4)
Bypass Priority Group	Bypass Priority Group		Bypass Special Categories	
IP Address is Excluded	Request Matches 32bit IP in a Policy	Global Policy Settings	Business Oriented	Block
Default Only Check Box is Enabled	Request Matches Username in a Policy	Custom and Trusted URLs	News	Authenticate
	Request Matches Group in a Policy	Category States	Search Sites	Coach
		Custom File Types	Educational	Allow