Overview and Essentials

> As of November 2023, the `Network Access` module has been renamed to `Content Filter`.

Create and manage rulesets to control the web content access of Local and Remote Devices.

Important Notes:

1. About changing default Category Allow/Block settings

The DrawBridge comes with a preset Action for each included Category. When you assign an Action (Allow/Block) to a Category, **you're simply applying a change that gets higher priority than the default setting.** This means:

1. You don't need to re-specify your Action preference for every built-in Category -- you only need to include the Categories in your Access Policy that you wish to assign a different action to than is default.
   For example: built-in Category **Sports** is set to a default action of **Block**.
   - If **Block** is the action you prefer, you *do not* need to add it to an Access Policy (eg. Company Preferences) with an action of **Block** -- the default setting is already doing this.
   - If **Allow** is the action you prefer, then you *do* need to add it to an Access Policy (eg. Company Preferences) with an action of **Allow** to override the default action.
2. In the event a custom Access Policy is removed, the filter will revert to the default Action for that Category.

2. Default Category settings are Business-focused

The default settings for the Built-in Categories are tightly scoped to business-usage needs. Depending on your usage expectations, you will want to set more categories to **Allow** in your *Company Preferences* Access Policy, or in a custom Access Policy.

Categories, Category Types, and Actions

Categories contain **Patterns**:

Pattern: a text string representing a domain or regular expression.

Categories can be one of two types:

- Classifier category
- ACL (Access Control List) category

Actions that can be assigned to a category, by type:

- Classifier category:
  - Allow
  - Block
  - Ignore
- ACL category:
  - Whitelist (allow in spite of Classifier score, above)
  - Blacklist (block in spite of Classifier score, above)
  - Blanketblock (block all requests Not matching these patterns)

Understanding Classifier categories

Classifier category patterns consist primarily of word and phrase lists (and also domains). The Redwood filter engine evaluates HTTP/S requests and responses and totals up a score for all categories with matching patterns. Then Redwood applies the action (Allow/Block) assigned to the top-scoring category.

Built-in Category patterns are managed by Compass Foundation. If you have improvements you wish to have considered for inclusion in the Built-in Categories, please send a detailed email to support@compassfoundation.io.

Understanding ACL actions & categories

Background

The Redwood filter engine analyzes all the components of a URL, including:

- Schema
- Top-level Domain (TLD), Domain, and Subdomains
- Path
- Query String

https://mobile.example.com/index.html?q=compass					
https://	mobile.	example	.com	/index.html	?q=compass
Schema	**Subdomain**	**Domain/Host**	**TLD (Top Level Domain)**	**Subdirectory/Path**	**Query String**

Also, Redwood analyzes additional parameters of the HTTP request:

- Method
- Content Type
- User Agent
- Referrer
- and more

Illustrated:

Additional HTTP parameters			
"GET" / "POST" / "HEAD" (and others)	"text/html" / "image/jpeg" (and others)	"Mozilla/5.0 (Windows NT 10.0; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/106.0.0.0 Safari/537.36"	https://duckduckgo.com
(Request) **Method**	**Content Type**	**UA (User Agent) String**	(Request) **Referrer**

In general, an ACL leverages one or more of these parameters to "tag" a specific action to a request, despite the Category score assigned to the request by the Classifier.

(In other words, this prevents an "arms-race" situation wherein competing actions are assigned by various Classifier categories; an ACL action will always take effect when the parameters match, no matter what the Classifier score and associated Category Action is.)

Note: for an ACL action to fire, the request must meet the minimum threshold score of 200 points. At that point, the action assigned by the ACL to the request is "authoritative", again, no matter the Classifier score.

Redwood ACL Actions

Action	About

`allow`	permit the request
`block`	deny the request
`ignore`	do not factor in the score assigned by this category
`censor_words`	strip out profanity
`disable_proxy_headers`	strip out the X-FORWARDED-FOR header
`hash_image`	generate a mathmatical hash of this picture
`phrase_scan`	evaluate content for matching word phrases
`require_auth`	force HTTP 407 proxy authentication response/challenge
`sslbump`	intercept the SSL/TLS encrypted session
`sslbypass`	do Not intercept the SSL/TLS encrypted session
`virus_scan`	hand off response to external analysis engine

For example, ACLs managed behind-the-scenes of the DrawBridge instruct Redwood to fire the SSL/TLS-inspection action on all requests (or not, in the case of SSLbypass/"Bypass Filter").

ACL Categories in the Console

Categories of the the type ACL enable you to leverage the "authoritative" nature of ACLs in your filter configuration.

In general, it is recommended to configure the desired content filter behavior by assigning Allow or Block to the built-in Classifier categories -- leveraging the "intelligence" built-in to these categories is a much less maintenance-intesive route to content control.

However, perhaps you want to Always assign a specific action (eg. Block) to a specific website. ACL categories are your friend in such a case: by adding a domain to an ACL category with a Block action assigned, the website will always block, even if the action assigned to the Classifier category is Allow.

The preset Always Allow and Always Block options in the Access Policy Dashboard are putting the domains in an ACL category that has the corresponding action assigned to it. These apply Company-wide.

Note: the default score assigned to all ACL category patterns is "1500". Adjusting this number will have **no impact** on the outcome of the action taken for that pattern, so long as the number is over the minimum score threshold of 200 -- the key detail here is that the pattern is part of an ACL category, so the action assigned to the ACL category is what will happen.

Advanced ACLs in the Console

Advanced ACLs simply expose many more "knobs" to apply a specific action with more granularity. Perhaps, for example, you want to only sslbypass a specific website for a specific Device Group. Advanced ACLs give you the toolset to configure that.

Filter Actions FAQ

- **Q:** What happens if I put a domain in both Always Allow and Always Block? Or what if I put a domain in two different ACL categories with competing actions assigned?
- **A:** Don't do that. :) In such a case, the outcome will be arbitrary. Decide what action you really want to have happen and adjust the policy accordingly.

What is an Access Policy?

An Access policy is the grouping of Devices, Actions, Times, (and, optionally, Applications) to create a customized DrawBridge content filter configuration.

This diagram illustrates:

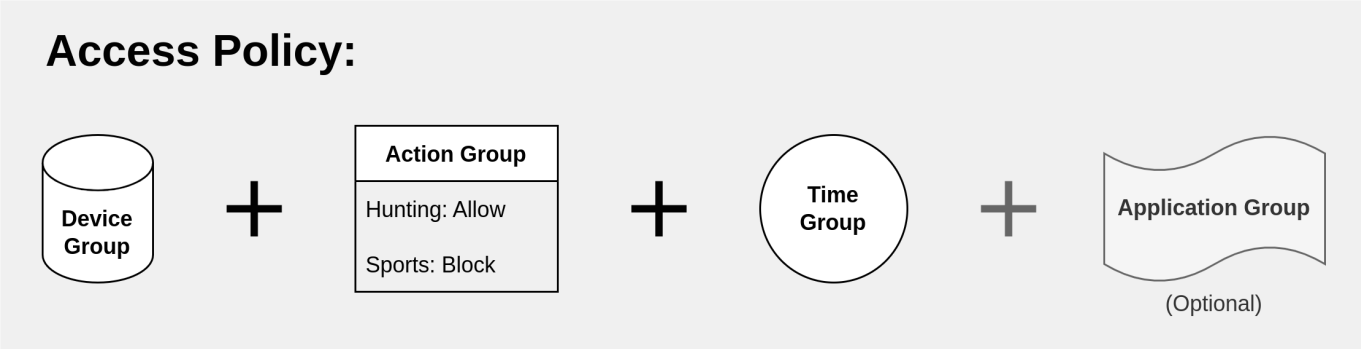

The Drawbridge supports the "stacking" or "layering" of Access Policies, enabling you to tailor the content filter experience for your users.

Access Policies by Type/Scope

Tenancy Type	Ruleset Scope
Company Access Policy	one Company
Access Policy Group	one Accountability Policy; available to apply to member Companies

Tenancy Type	Ruleset Scope
Universal Access Policy Group	globally available to all DrawBridges
System Access Policy	a specific DrawBridge; applies to all tenant Companies on that system

What is a Rating?

The DrawBridge classifies text into categories. But what is the the tone of these categories? And what do they values do they represent? A Rating system should help answer that question, as well as offer visual clues for the report reader.

But what kind of rating system? Unlike other filter projects, the DrawBridge does not rate content by who it's appropriate for - as in Everyone / Teens / Adults - but somewhat more like *where* it is appropriate. The rating names are drawn from the concept of particulate filtering - how fine or coarse is the filter mesh that would permit the content to traverse it.

A key assumption here is that the Internet is most frequently being used in a workplace environment, facilitating the everyday tasks of research, transactions, and commerce. Usage reports are colorized according to the Category Ratings of the content that was accessed.

Misc Rating

The Misc Rating is used when no category of interest could be found. Perhaps the request incident was not text-based, or perhaps a category needs to be extended or created to for this type of situation.

Base Rating

The Base Rating is the most general grade, including categories like Search Engines or Technology Services. Any more specific category and rating would be preferred. For example, it's great to know that a body of text is about Search Engines, but it's better to know what is being searched for.

Silt Rating

The Silt Rating is expected usage in the workplace environment. While not every workplace will commonly access every category in the Silt Rating, any given user in the business environment will periodically need most categories found here.

It is recommended that all categories in the Silt Rating be "allowed" in the workplace, although policies can be created to limit access to given devices.

Sand Rating

The Sand Rating will still be frequently used in the workplace environment, although the industry type will very much determine how much categories in Sand Rating are accessed.

Categories in the Sand Rating can be "allowed" or "blocked" per the business owner's preferences or the preferences established by the Accountability Policy.

Pebble Rating

The Pebble Rating contains categories that generally fall outside the workplace, while remaining universally pertinent to other areas of life, such as Medical, News, Clothing, etc.

Categories in the Pebble Rating can be "allowed" or "blocked" per the business owner's preferences or the preferences established by the Accountability Policy.

Stone Rating

The Stone Rating contains categories that are increasingly beyond the scope of any type of workplace, reaching more into popular culture and society at large.

Categories in the Stone Rating will typically be blocked by most business owners and school administrators.

Rock Rating

The Rock Rating contains categories that tend to represent the rougher edges of popular culture and general society.

Categories in the Rock Rating will typically be blocked by all business owners and school administrators.

Boulder Rating

The Boulder Rating categories that represent the "redlight" district of the Internet. These categories cannot be enabled in the Redwood Console even by administrators.

Categories in the Boulder Rating are always blocked, and cannot be allowed in the DrawBridge.

Actions for Classifier categories

Action	When this category is the top-scoring one on a web request:
`allow`	web request content loads as expected
`block`	web request is served a block page instead of the original destination webpage
`ignore`	web request action referred to next-to-top scoring category

When to use the `ignore` Action

In most situations, the category action should be `allow` or `block`, but in some situations the next-to-top scoring category is more meaningful. For example, an automotive shop may perform work that overlaps with the Racing category. If Racing is set to `block`, the shop's activities will be hampered. If Racing is set to `allow`, then access may be wider than desired.

Solution - set Racing to `ignore`. If next-to-top-scoring category is Automotive, the page will be allowed, and if it's Sports, the page will be blocked as Sports.

Filter processing flowchart: Category Filtering

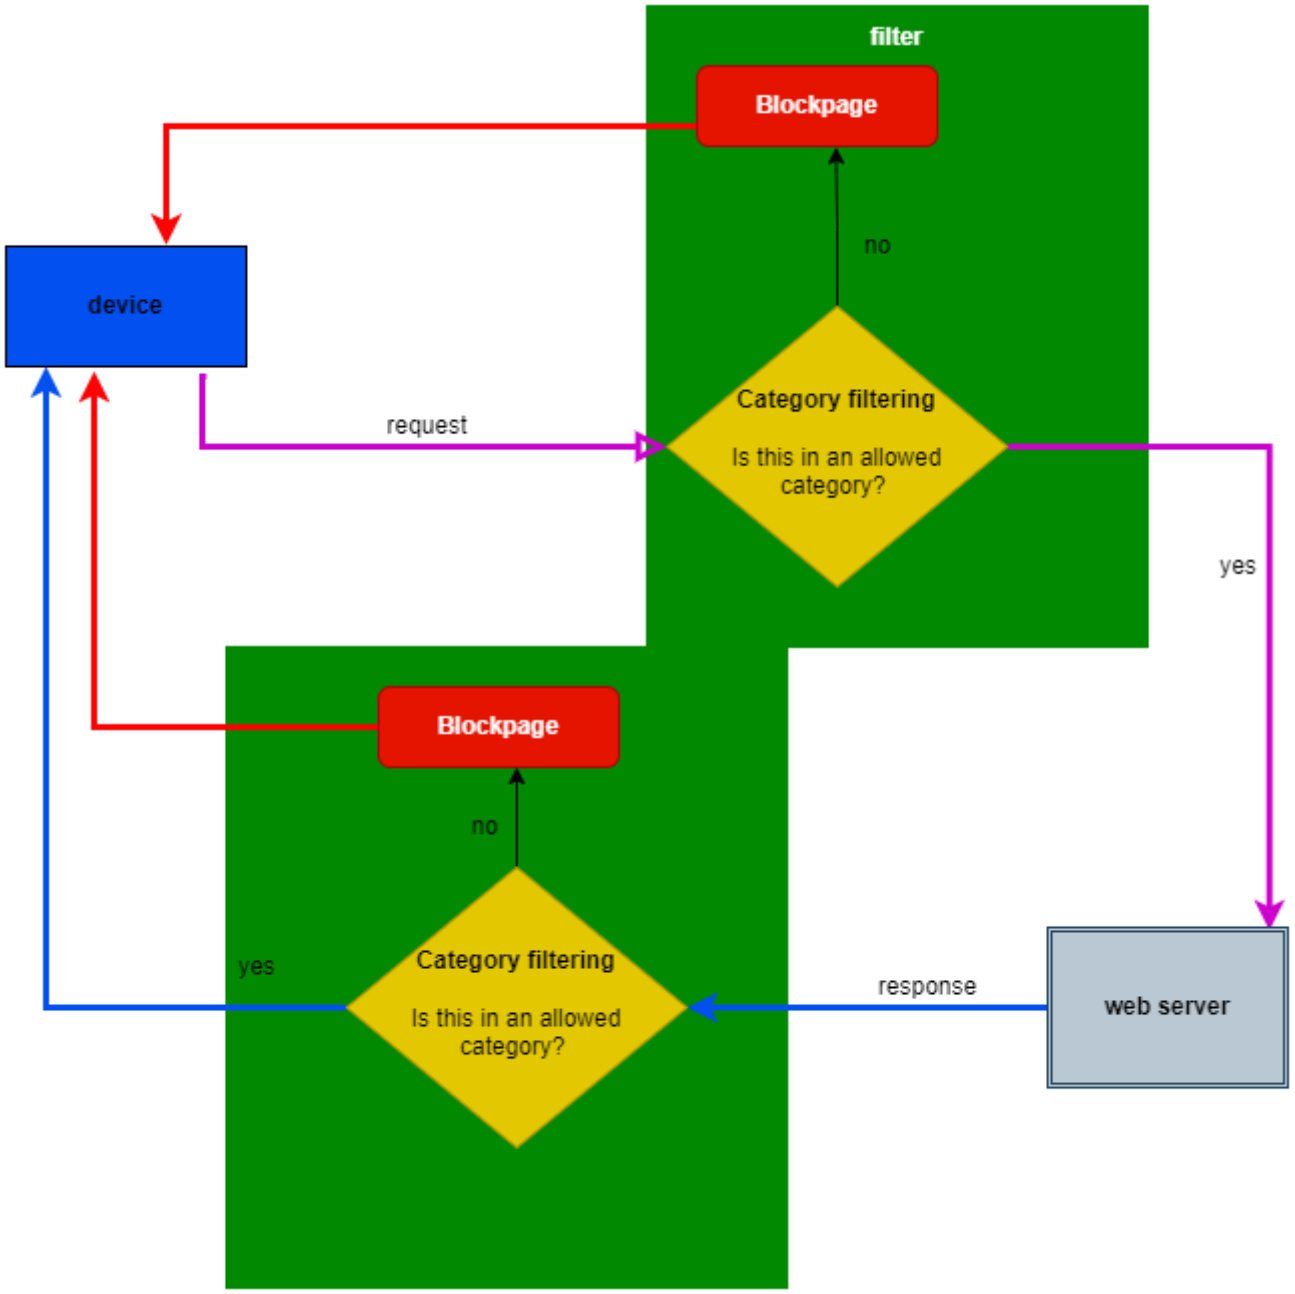

Actions for ACL categories

Action	About
whitelist	A Category consisting of domains (and/or regular expression patterns) that the DrawBridge will Always Allow, in spite of the content scores. ***Use with caution!***
blacklist	A Category consisting of domains (and/or regular expression patterns) that the DrawBridge will Always Block, in spite of the content scores.

Action	About
`blanketblock`	A Category consisting of domains (and/or regular expression patterns) to which the DrawBridge will apply *regular category-based filtering* **and** block access to all other sites **not specified** in the blanketblock category (or a linked category)

See below for more information:

Whitelist

A Category consisting of domains (and/or regular expression patterns) that the DrawBridge will Always Allow, in spite of the content scores. ***Use with caution!***

Filter processing flowchart:

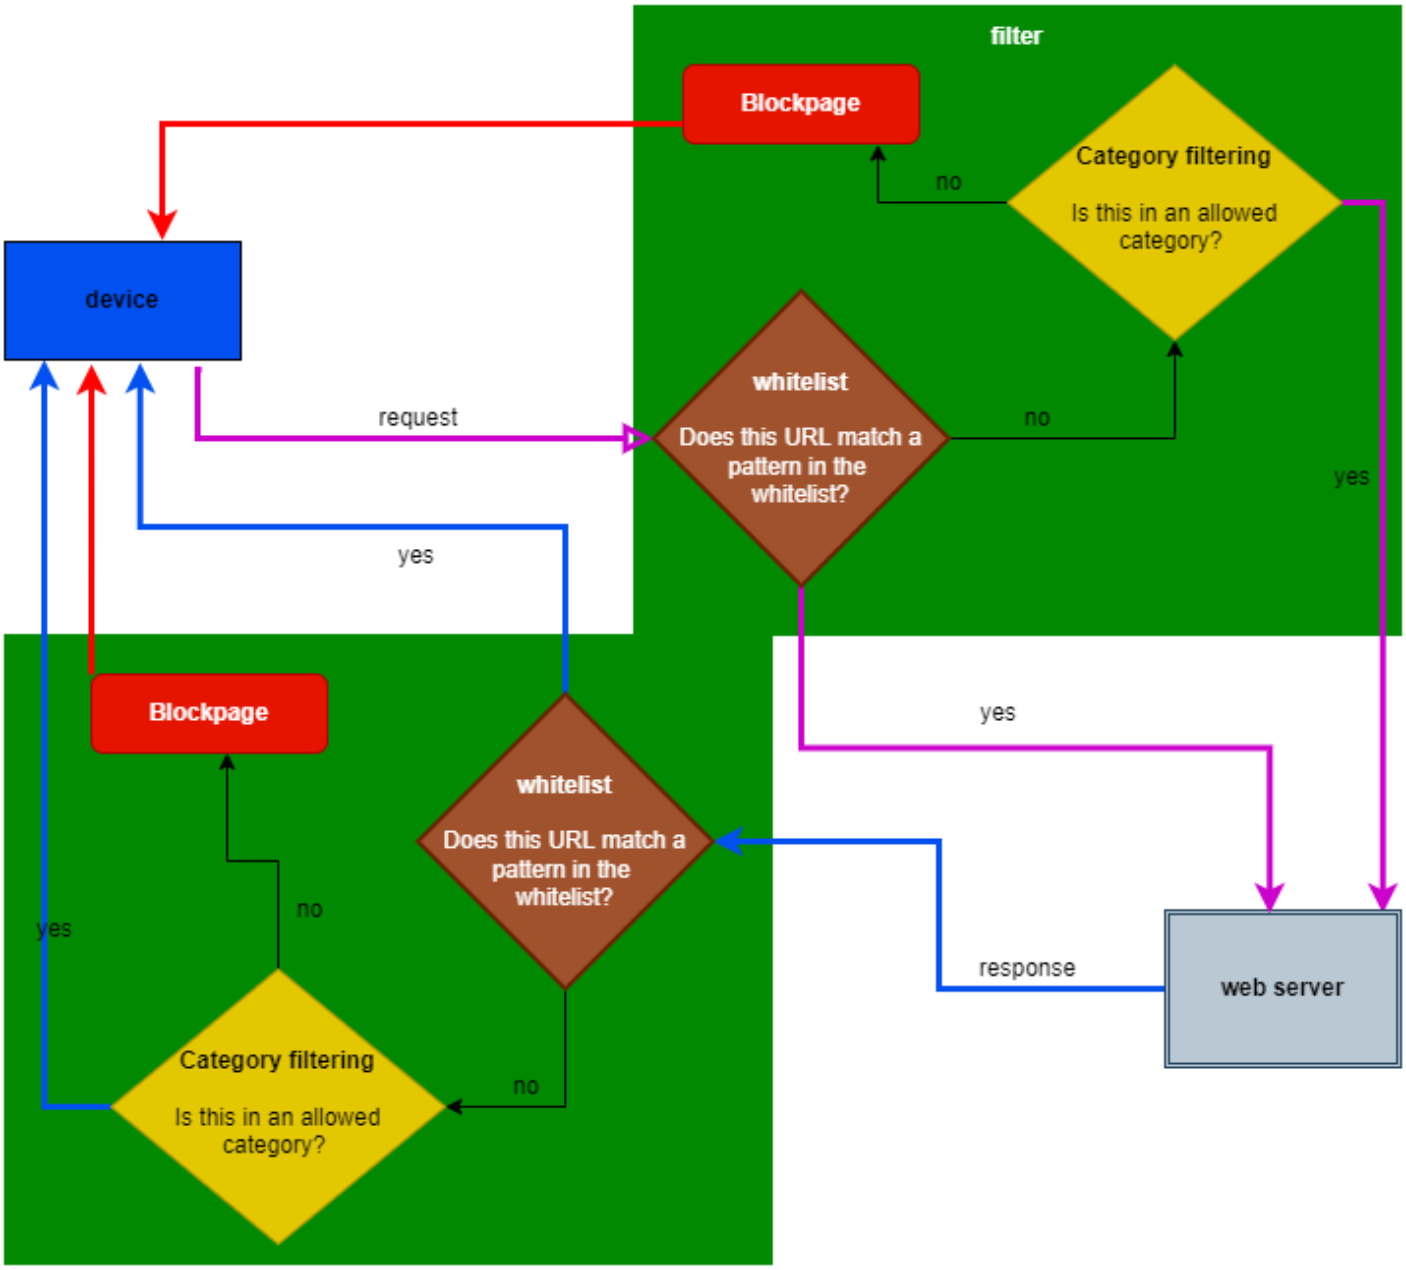

Blacklist

A Category consisting of domains (and/or regular expression patterns) that the DrawBridge will Always Block, in spite of the content scores.

Filter processing flowchart:

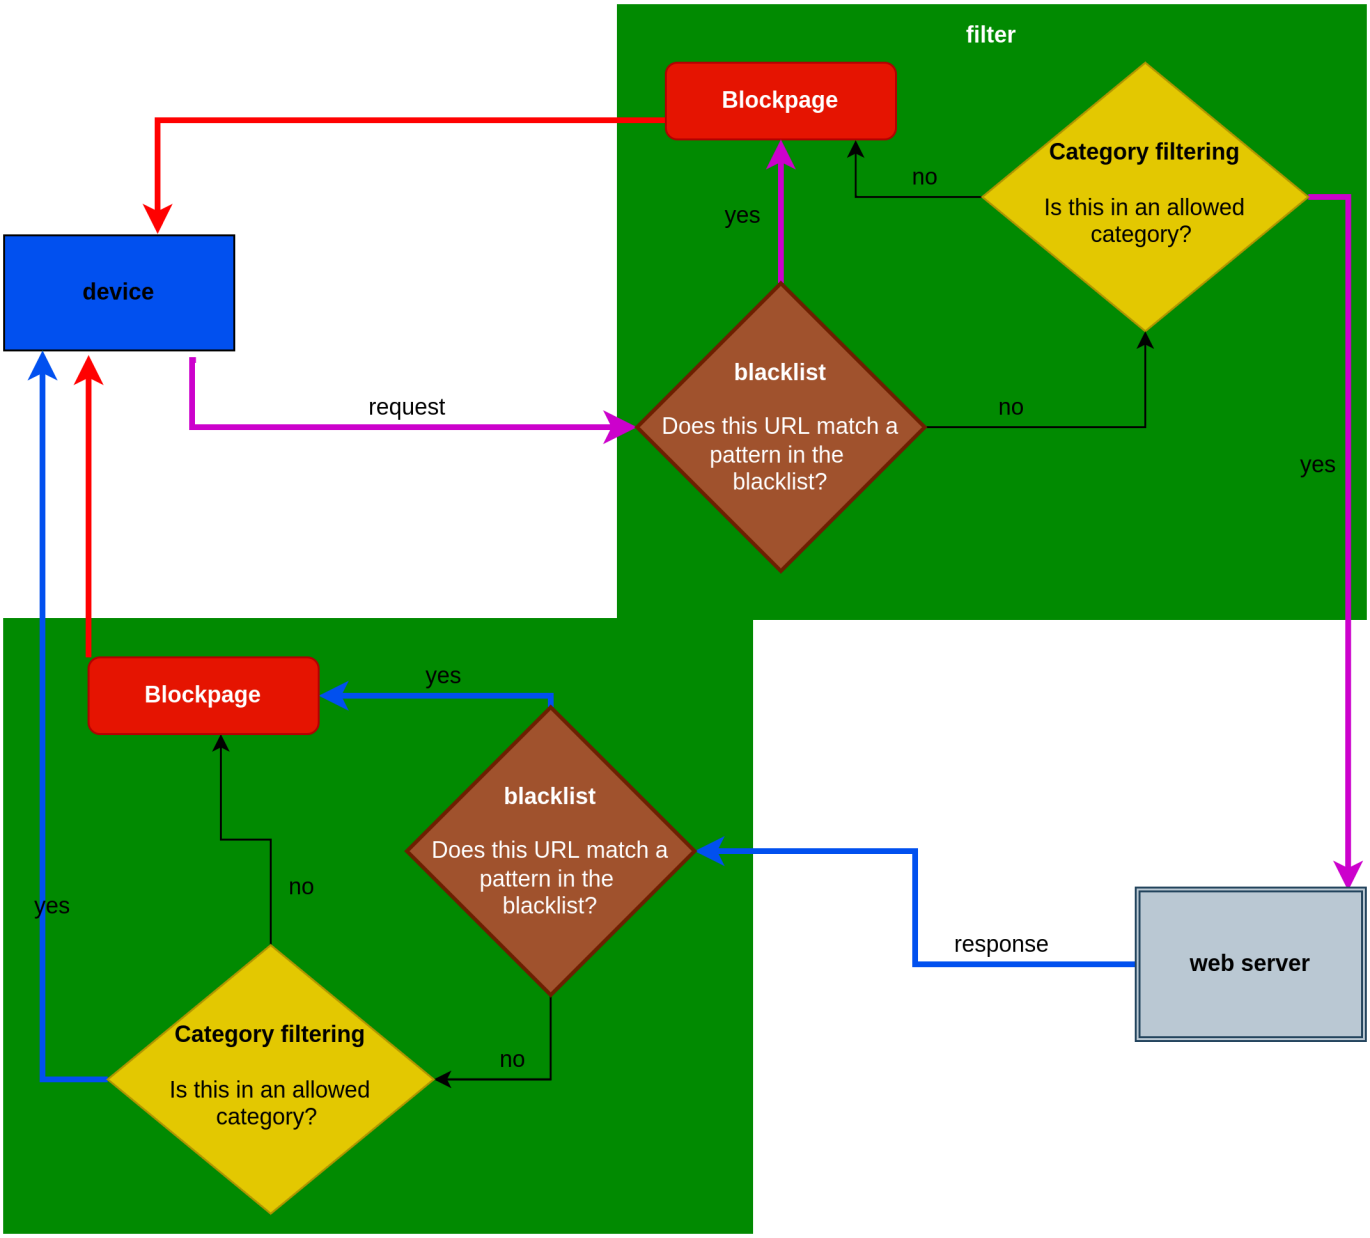

Blanketblock

A Category consisting of domains (and/or regular expression patterns) to which the DrawBridge will apply *regular category-based filtering* **and** block access to all other sites **not specified** in the blanketblock category (or a linked category)

Filter processing flowchart:

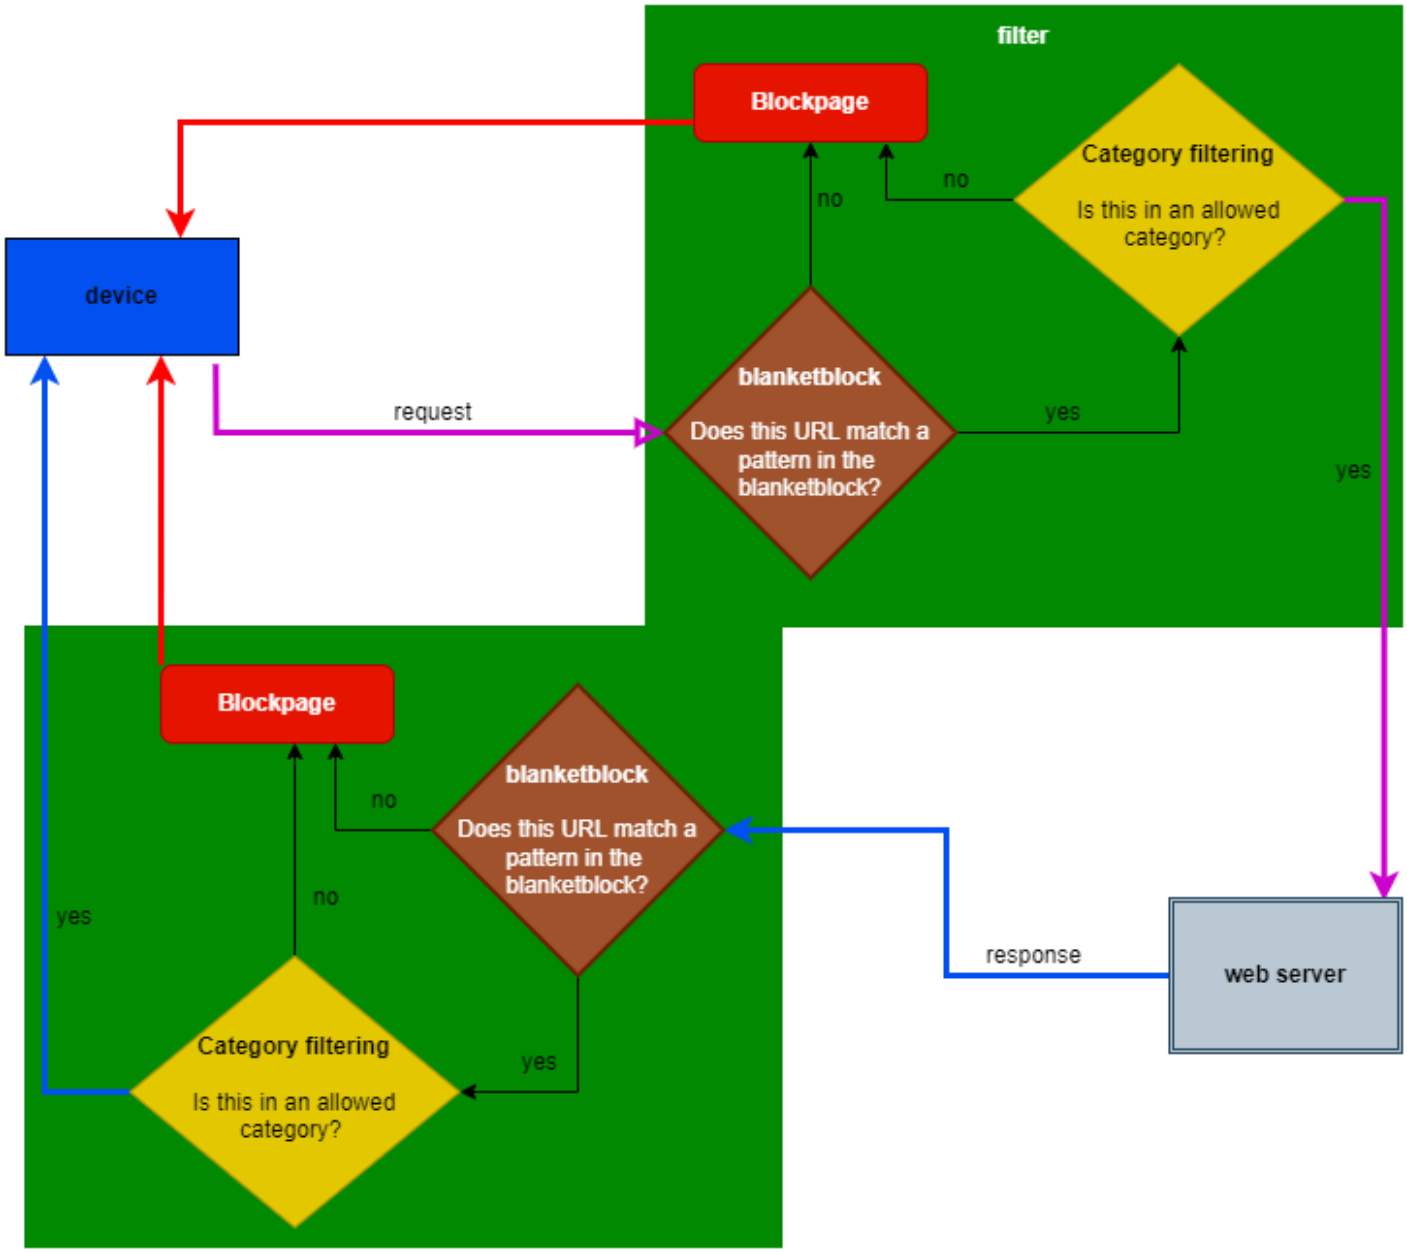

Revision #31
Created 15 September 2022 21:07:24 by Marvin M.
Updated 13 February 2024 14:53:43 by james