
bedrock Documentation

Release 1.0

Mozilla

May 22, 2019

Contents

1	Contents	3
1.1	Installing Bedrock	3
1.2	Run the tests	5
1.3	Make it run	6
1.4	Legal Docs	7
1.5	Localization	7
1.6	Feature Flipping (aka Switches)	7
1.7	Localization	8
1.8	Developing on Bedrock	15
1.9	How to contribute	19
1.10	Continuous Integration & Deployment	22
1.11	Front-end testing	25
1.12	Managing Redirects	29
1.13	JavaScript Libraries	32
1.14	Newsletters	33
1.15	Using External Content Cards Data	35
1.16	Tabzilla	36
1.17	Mozilla.UITour	37
1.18	Send to Device widget	47
1.19	Firefox Accounts Signup Form	49
1.20	Funnel cakes and Partner Builds	50
1.21	A/B Testing	51
1.22	Analytics	53
1.23	Stub Attribution	55
1.24	Architectural Decision Records	57

bedrock is the project behind www.mozilla.org. It is as shiny, awesome, and open-sourcy as always. Perhaps even a little more.

bedrock is a web application based on [Django](#), a [Python](#) web application framework.

Patches are welcome! Feel free to fork and contribute to this project on [Github](#).

1.1 Installing Bedrock

There are two primary methods of installing bedrock: Docker and Local. Whichever you choose you'll start by getting the source:

```
$ git clone --recursive git://github.com/mozilla/bedrock.git
$ cd bedrock
```

After these basic steps you can choose your install method below. Docker is the easiest and recommended way, but local is also possible and may be preferred by people for various reasons.

1.1.1 Docker Installation

Note: This method assumes you have [Docker installed for your platform](#). If not please do that now or skip to the [Local Installation](#) section.

This is the simplest way to get started developing for bedrock. If you're on Linux or Mac (and possibly Windows 10 with the Linux subsystem) you can run a script that will pull our production and development docker images and start them:

```
$ make clean run
```

Note: You can start the server any other time with:

```
$ make run
```

You should see a number of things happening, but when it's done it will output something saying that the server is running at localhost:8000. Go to that URL in a browser and you should see the mozilla.org home page. In this mode

the site will refresh itself when you make changes to any template or media file. Simply open your editor of choice and modify things and you should see those changes reflected in your browser.

Note: It's a good idea to run `make pull` from time to time. This will pull down the latest Docker images from our repository ensuring that you have the latest dependencies installed among other things. If you see any strange errors after a `git pull` then `make pull` is a good thing to try for a quick fix.

If you don't have or want to use Make you can call the docker and compose commands directly:

```
$ docker-compose pull
$ git submodule sync
$ git submodule update --init --recursive
$ [[ ! -f .env ]] && cp .env-dist .env
```

Then starting it all is simply:

```
$ docker-compose up app assets
```

All of this is handled by the Makefile script and called by Make if you follow the above directions. You **DO NOT** need to do both.

These directions pull and use the pre-built images that our deployment process has pushed to the [Docker Hub](#). If you need to add or change any dependencies for Python or Node then you'll need to build new images for local testing. You can do this by updating the requirements files and/or package.json file then simply running:

```
$ make build
```

Asset bundles

If you make a change to `media/static-bundles.json`, you'll need to restart Docker.

Note: Sometimes stopping Docker doesn't actually kill the images. To be safe, after stopping docker, run `docker ps` to ensure the containers were actually stopped. If they have not been stopped, you can force them by running `docker-compose kill` to stop all containers, or `docker kill <container_name>` to stop a single container, e.g. `docker kill bedrock_app_1`.

1.1.2 Local Installation

These instructions assume you have Python, pip, and NodeJS installed. If you don't have *pip* installed (you probably do) you can install it with the instructions in the [pip docs](#).

You need to create a virtual environment for Python libraries. Skip the first instruction if you already have virtualenv installed:

```
$ pip install virtualenv # installs virtualenv, skip if already_
↪ have it
$ virtualenv -p python2.7 venv # create a virtual env in the folder_
↪ `venv`
$ source venv/bin/activate # activate the virtual env. On Windows,
↪ run: venv\Scripts\activate.bat
$ pip install -U pip # securely upgrade pip
$ pip install -r requirements/dev.txt # installs dependencies
```


If you are on OSX and some of the compiled dependencies fails to compile, try explicitly setting the arch flags and try again:

```
$ export ARCHFLAGS="-arch i386 -arch x86_64"
$ pip install -r requirements/dev.txt
```

If you are on Linux, you will need at least the following packages or their equivalent for your distro:

```
$ python-dev libxslt-dev
```

Sync the database and all of the external data locally. This gets product-details, security-advisories, credits, release notes, localizations, legal-docs etc:

```
$ bin/bootstrap.sh
```

Next, you need to have [Node.js](#) and [Yarn](#) installed. The node dependencies for running the site can be installed with yarn:

```
$ yarn
```

You'll also need to install the [Gulp](#) cli globally:

```
$ npm install -g gulp-cli
```

Note: Bedrock uses yarn to ensure that Node.js packages that get installed are the exact ones we meant (similar to pip hash checking mode for python). Refer to the [yarn documentation](#) for adding or upgrading Node.js dependencies.

1.2 Run the tests

Now that we have everything installed, let's make sure all of our tests pass. This will be important during development so that you can easily know when you've broken something with a change.

1.2.1 Docker

We manage our local docker environment with docker-compose and Make. All you need to do here is run:

```
$ make test
```

If you don't have Make you can simply run `docker-compose run test`.

If you'd like to run only a subset of the tests or only one of the test commands you can accomplish that with a command like the following:

```
$ docker-compose run test py.test bedrock/firefox
```

This example will run only the unit tests for the `firefox` app in bedrock. You can substitute `py.test bedrock/firefox` with most any shell command you'd like and it will run in the Docker container and show you the output. You can also just run `bash` to get an interactive shell in the container which you can then use to run any commands you'd like and inspect the file system:

```
$ docker-compose run test bash
```

1.2.2 Local

From the local install instructions above you should still have your virtualenv activated, so running the tests is as simple as:

```
$ py.test lib bedrock
```

To test a single app, specify the app by name in the command above. e.g.:

```
$ py.test bedrock/firefox
```

Note: If your local tests run fine, but when you submit a pull-request the tests fail in [CircleCI](#), it could be due to the difference in settings between what you have in `.env` and what CircleCI uses: `docker/envfiles/demo.env`. You can run tests as close to Circle as possible by moving your `.env` file to another name (e.g. `.env-backup`), then copying `docker/envfiles/demo.env` to `.env`, and running tests again.

1.3 Make it run

1.3.1 Docker

You can simply run the `make run` script mentioned above, or use `docker-compose` directly:

```
$ docker-compose up app assets
```

1.3.2 Local

To make the server run, make sure your virtualenv is activated, and then run the server:

```
$ npm start
```

If you are not inside a virtualenv, you can activate it by doing:

```
$ source venv/bin/activate
```

1.3.3 Browsersync

Both the Docker and Local methods of running the site use [Browsersync](#) to serve the development static-assets (CSS, JS, etc.) as well as refresh the browser tab for you when you change files. The refreshing of the page works by injecting a small JS snippet into the page that listens to the browsersync service and will refresh the page when it receives a signal. It also injects a script that shows a small notification in the top-right corner of the page to inform you that a refresh is happening and when the page connects to or is disconnected from the browsersync service. We've not seen issues with this, but since it is modifying the page it is possible that this could conflict with something on the page itself. Please let us know if you suspect this is happening for you. This notification can be disabled in the browsersync options in the `gulpfile.js` by setting `notify: false` in the `browser-sync` task.

1.4 Legal Docs

Legal docs (for example: the privacy policy) are generated from markdown files in the [legal-docs repo](#).

To view them or update to a more recent version update the submodule:

```
$ git submodule update --init --recursive
```

1.5 Localization

Localization (or L10n) files were fetched by the *bootstrap.sh* command you ran earlier and are included in the docker images. If you need to update them or switch to a different repo or branch after changing settings you can run the following command:

```
$ ./manage.py l10n_update
```

You can read more details about how to localize content [here](#).

1.6 Feature Flipping (aka Switches)

Environment variables are used to configure behavior and/or features of select pages on bedrock via a template helper function called `switch()`. It will take whatever name you pass to it (must be only numbers, letters, and dashes), convert it to uppercase, convert dashes to underscores, and lookup that name in the environment. For example: `switch('the-dude')` would look for the environment variable `SWITCH_THE_DUDE`. If the value of that variable is any of “on”, “true”, “1”, or “yes”, then it will be considered “on”, otherwise it will be “off”.

You can also supply a list of locale codes that will be the only ones for which the switch is active. If the page is viewed in any other locale the switch will always return `False`, even in `DEV` mode. This list can also include a “Locale Group”, which is all locales with a common prefix (e.g. “en-US, en-GB” or “zh-CN, zh-TW”). You specify these with just the prefix. So if you used `switch('the-dude', ['en', 'de'])` in a template, the switch would be active for German and any English locale the site supports.

You may also use these switches in Python in `views.py` files (though not with locale support). For example:

```
from bedrock.base.waffle import switch

def home_view(request):
    title = 'Staging Home' if switch('staging-site') else 'Prod Home'
    ...
```

1.6.1 Testing

If the environment variable `DEV` is set to a “true” value, then all switches will be considered “on” unless they are explicitly “off” in the environment. `DEV` defaults to “true” in local development and demo servers.

To test switches locally:

1. Set `DEV=False` in your `.env` file.
2. Enable the switch in your `.env` file.
3. Restart your web server.

To configure switches for a demo branch. Follow the [configuration instructions here](#).

1.6.2 Traffic Cop

Currently, these switches are used to enable/disable [Traffic Cop](#) experiments on many pages of the site. We only add the Traffic Cop JavaScript snippet to a page when there is an active test. You can see the current state of these switches and other configuration values in our [configuration repo](#).

To work with/test these experiment switches locally, you must add the switches to your local environment. For example:

```
# to switch on firstrun-copy-experiment you'd add the following to your ``.env`` file
SWITCH_FIRSTRUN_COPY_EXPERIMENT=on
```

To do the equivalent in one of the bedrock apps see the [www-config](#) documentation.

1.6.3 Notes

A shortcut for activating virtual envs in zsh or bash is `. venv/bin/activate`. The dot is the same as `source`.

There's a project called [virtualenvwrapper](#) that provides a better interface for managing/activating virtual envs, so you can use that if you want. Also if you need help managing various versions of Python on your system, the [pyenv](#) project can help.

1.7 Localization

The site is fully localizable. Localization files are not shipped with the code distribution, but are available in a separate GitHub repository. The proper repo can be cloned and kept up-to-date using the `l10n_update` management command:

```
$ ./manage.py l10n_update
```

If you don't already have a `locale` directory, this command will clone the git repo containing the translation files (either the dev or prod files depending on your `DEV` setting). If the folder is already present, it will update the repository to the latest version.

1.7.1 .lang files

Bedrock supports a workflow similar to `gettext`. You extract all the strings from the codebase, then merge them into each locale to get them translated.

The files containing the strings are called “.lang files” and end with a `.lang` extension.

To extract all the strings from the codebase, run:

```
$ ./manage.py l10n_extract
```

If you'd only like to extract strings from certain files, you may optionally list them on the command line:

```
$ ./manage.py l10n_extract bedrock/mozorg/templates/mozorg/contribute.html
```

Command line glob matching will work as well if you want all of the HTML files in a directory, for example:

```
$ ./manage.py l10n_extract bedrock/mozorg/templates/mozorg/*.html
```

That will use `gettext` to get all the needed localizations from Python and HTML files, and will convert the result into a series of `.lang` files inside `locale/templates`. This directory represents the “reference” set of strings to be translated, and you are free to modify or split up `.lang` files here as needed (just make sure they are being referenced correctly, from the code, see *Which .lang file should it use?*).

Once you have extracted `.lang` files locally, they can then be added via pull request to the [mozilla.org l10n](https://www.mozilla.org/l10n) repository for translation. You can read the [full documentation](#) for more information on the extraction workflow.

Translating with .lang files

To translate a string from a `.lang` file, simply use the `gettext` interface.

In a `jinja2` template:

```
<div>{{ _('Hello, how are you?') }}</div>

<div>{{ _('<a href="%s">Click here</a>') |format('http://mozilla.org/') }}</div>

<div>{{ _('<a href="%s">Click here</a>') |format(url='http://mozilla.org/') }}</div>
```

Note the usage of variable substitution in the latter examples. It is important not to hardcode URLs or other parameters in the string. `jinja`’s `format` filter lets us apply variables outside of the string.

You can provide a one-line comment to the translators like this:

```
{# L10n: "like" as in "similar to", not "is fond of" #}
{{ _('Like this:') }}
```

The comment will be included in the `.lang` files above the string to be translated.

In a Python file, use `lib.l10n_utils.dotlang._` or `lib.l10n_utils.dotlang._lazy`, like this:

```
from lib.l10n_utils.dotlang import _lazy as _

sometext = _('Foo about bar.')
```

You can provide a one-line comment to the translators like this:

```
# L10n: "like" as in "similar to", not "is fond of"
sometext = _('Like this:')
```

The comment will be included in the `.lang` files above the string to be translated.

There’s another way to translate content within `jinja2` templates. If you need a big chunk of content translated, you can put it all inside a `trans` block.

```
{% trans %}
  <div>Hello, how are you</div>
{% endtrans %}

{% trans url='http://mozilla.org' %}
  <div><a href="{{ url }}">Click here</a></div>
{% endtrans %}
```

Note that it also allows variable substitution by passing variables into the block and using template variables to apply them.

Which .lang file should it use?

Translated strings are split across several .lang files to make it easier to manage separate projects and pages. So how does the system know which one to use when translating a particular string?

- All translations from Python files are put into `main.lang`. This should be a very limited set of strings and most likely should be available to all pages.
- Templates always load `main.lang` and `download_button.lang`.
- Additionally, each template has its own .lang file, so a template at `mozorg/firefox.html` would use the .lang file at `<locale>/mozorg/firefox.lang`.
- Templates can override which .lang files are loaded. The above global ones are always loaded, but instead of loading `<locale>/mozorg/firefox.lang`, the template can specify a list of additional lang files to load with a template block:

```
{% add_lang_files "foo" "bar" %}
```

That will make the page load `foo.lang` and `bar.lang` in addition to `main.lang` and `download_button.lang`.

When strings are extracted from a template, they are added to the template-specific .lang file. If the template explicitly specifies .lang files like above, it will add the strings to the first .lang file specified, so extracted strings from the above template would go into `foo.lang`.

You can similarly specify extra .lang files in your Python source as well. Simply add a module-level constant in the file named `LANG_FILES`. The value should be either a string, or a list of strings, similar to the `add_lang_files` tag above.

```
# forms.py

from lib.l10n_utils.dotlang import _

LANG_FILES = ['foo', 'bar']

sometext = _('Foo about bar.')
```

This file's strings would be extracted to `foo.lang`, and the lang files `foo.lang`, `bar.lang`, `main.lang` and `download_button.lang` would be searched for matches in that order.

1.7.2 l10n blocks

Bedrock also has a block-based translation system that works like the `{% block %}` template tag, and marks large sections of translatable content. This should not be used very often; lang files are the preferred way to translate content. However, there may be times when you want to control a large section of a page and customize it without caring very much about future updates to the English page.

A Localizers' guide to l10n blocks

Let's look at how we would translate an example file from **English** to **German**.

The English source template, created by a developer, lives under `apps/appname/templates/appname/example.html` and looks like this:

```
{% extends "base-pebbles.html" %}

{% block content %}
  
```

(continues on next page)

(continued from previous page)

```
{% l10n foo, 20110801 %}
<h1>Hello world!</h1>
{% endl10n %}

<hr>

{% l10n bar, 20110801 %}
<p>This is an example!</p>
{% endl10n %}
{% endblock %}
```

The `l10n` blocks mark content that should be localized. Realistically, the content in these blocks would be much larger. For a short string like above, please use lang files. We'll use this trivial code for our example though.

The `l10n` blocks are named and tagged with a date (in ISO format). The date indicates the time that this content was updated and needs to be translated. If you are changing trivial things, you shouldn't update it. The point of `l10n` blocks is that localizers completely customize the content, so they don't care about small updates. However, you may add something important that needs to be added in the localized blocks; hence, you should update the date in that case.

When the command `./manage.py l10n_extract` is run, it generates the corresponding files in the `locale` folder (see below for more info on this command).

The German version of this template is created at `locale/de/templates/appname/example.html`. The contents of it are:

```
{% extends "appname/example.html" %}

{% l10n foo %}
<h1>Hello world!</h1>
{% endl10n %}

{% l10n bar %}
<p>This is an example!</p>
{% endl10n %}
```

This file is an actual template for the site. It extends the main template and contains a list of `l10n` blocks which override the content on the page.

The localizer just needs to translate the content in the `l10n` blocks.

When the reference template is updated with new content and the date is updated on an `l10n` block, the generated `l10n` file will simply add the new content. It will look like this:

```
{% extends "appname/example.html" %}

{% l10n foo %}
<h1>This is an English string that needs translating.</h1>
{% was %}
<h1>Dies ist ein English string wurde nicht.</h1>
{% endl10n %}

{% l10n bar %}
<p>This is an example!</p>
{% endl10n %}
```

Note the `was` block in `foo`. The old translated content is in there, and the new content is above it. The `was` content is always shown on the site, so the old translation still shows up. The localizer needs to update the translated content

and remove the `was` block.

Generating the locale files

```
$ ./manage.py l10n_check
```

This command will check which blocks need to be translated and update the locale templates with needed translations. It will copy the English blocks into the locale files if a translation is needed.

You can specify a list of locales to update:

```
$ ./manage.py l10n_check fr
$ ./manage.py l10n_check fr de es
```

1.7.3 Currency

When dealing with currency, make a separate gettext wrapper, placing the amount inside a variable. You should also include a comment describing the intent. For example:

```
{# L10n: Inserts a sum in US dollars, e.g. '$100'. Adapt the string in your_
↳translation for your locale conventions if needed, ex: %(sum)s US$ #}
{{ _('${sum}s')|format(sum='15') }}
```

1.7.4 CSS

If a localized page needs some locale-specific style tweaks, you can add the style rules to the page's stylesheet like this:

```
html[lang="it"] #features li {
    font-size: 20px;
}

html[dir="rtl"] #features {
    float: right;
}
```

If a locale needs site-wide style tweaks, font settings in particular, you can add the rules to `/media/css/l10n/{{LANG}}/intl.css`. Pages on Bedrock automatically includes the CSS in the base templates with the `l10n_css` helper function. The CSS may also be loaded directly from other Mozilla sites with such a URL: `//mozorg.cdn.mozilla.net/media/css/l10n/{{LANG}}/intl.css`.

Open Sans, the default font on mozilla.org, doesn't offer non-Latin glyphs. `intl.css` can have `@font-face` rules to define locale-specific fonts using custom font families as below:

- *X-LocaleSpecific-Light*: Used in combination with *Open Sans Light*. The font can come in 2 weights: normal and optionally bold
- *X-LocaleSpecific*: Used in combination with *Open Sans Regular*. The font can come in 2 weights: normal and optionally bold
- *X-LocaleSpecific-Extrabold*: Used in combination with *Open Sans Extrabold*. The font weight is 800 only

Here's an example of `intl.css`:


```

@font-face {
  font-family: X-LocaleSpecific-Light;
  font-weight: normal;
  font-display: swap;
  src: local(mplus-2p-light), local(Meiryo);
}

@font-face {
  font-family: X-LocaleSpecific-Light;
  font-weight: bold;
  font-display: swap;
  src: local(mplus-2p-medium), local(Meiryo-Bold);
}

@font-face {
  font-family: X-LocaleSpecific;
  font-weight: normal;
  font-display: swap;
  src: local(mplus-2p-regular), local(Meiryo);
}

@font-face {
  font-family: X-LocaleSpecific;
  font-weight: bold;
  font-display: swap;
  src: local(mplus-2p-bold), local(Meiryo-Bold);
}

@font-face {
  font-family: X-LocaleSpecific-Extrabold;
  font-weight: 800;
  font-display: swap;
  src: local(mplus-2p-black), local(Meiryo-Bold);
}

```

Localizers can specify locale-specific fonts in one of the following ways:

- Choose best-looking fonts widely used on major platforms, and specify those with the `src: local(name)` syntax
- Find a best-looking free Web font, add the font files to `/media/fonts/`, and specify those with the `src: url(path)` syntax
- Create a custom Web font to complement missing glyphs in *Open Sans*, add the font files to `/media/fonts/l10n/`, and specify those with the `src: url(path)` syntax. *M+ 2c* offers various international glyphs and looks similar to *Open Sans*, while *Noto Sans* is good for the bold and italic variants. You can create subsets of these alternative fonts in the WOFF and WOFF2 formats using a tool found on the Web. See [Bug 1360812](#) for the Fulah (ff) locale's example

Developers should use the `.open-sans` mixin instead of `font-family: 'Open Sans'` to specify the default font family in CSS. This mixin has both *Open Sans* and *X-LocaleSpecific* so locale-specific fonts, if defined, will be applied to localized pages. The variant mixins, `.open-sans-light` and `.open-sans-extrabold`, are also available.

1.7.5 Staging Copy Changes

The need will often arise to push a copy change to production before the new copy has been translated for all locales. To prevent locales not yet translated from displaying English text, you can use the `l10n_has_tag` template function. For example, if the string “Firefox features” needs to be changed to “Firefox benefits”:

```
{% if l10n_has_tag('firefox_products_headline_spring_2016') %}
  <h1>{{ _('Firefox features') }}</h1>
{% else %}
  <h1>{{ _('Firefox benefits') }}</h1>
{% endif %}
```

This function will check the `.lang` file(s) of the current page for the tag `firefox_products_headline_spring_2016`. If it exists, the translation for “Firefox features” will be displayed. If not, the pre-existing translation for “Firefox benefits” will be displayed.

When using `l10n_has_tag`, be sure to coordinate with the localization team to decide on a good tag name. Always use underscores instead of hyphens if you need to visually separate words.

1.7.6 Locale-specific Templates

While the `l10n_has_tag` template function is great in small doses, it doesn’t scale particularly well. A template filled with conditional copy can be difficult to comprehend, particularly when the conditional copy has associated CSS and/or JavaScript.

In instances where a large amount of a template’s copy needs to be changed, or when a template has messaging targeting one particular locale, creating a locale-specific template may be a good choice.

Locale-specific templates function simply by naming convention. For example, to create a version of `/firefox/new.html` specifically for the `de` locale, you would create a new template named `/firefox/new.de.html`. This template can either extend `/firefox/new.html` and override only certain blocks, or be entirely unique.

When a request is made for a particular page, bedrock’s rendering function automatically checks for a locale-specific template, and, if one exists, will render it instead of the originally specified (locale-agnostic) template.

Note: Creating a locale-specific template for `en-US` was not possible when this feature was introduced, but it is now. So you can create your `en-US`-only template and the rest of the locales will continue to use the default.

Important: Note that the presence of an L10n template (e.g. `locale/de/templates/firefox/new.html`) will take precedence over a locale-specific template in bedrock.

1.7.7 Specifying Active Locales in Views

Normally we rely on activation tags in our translation files (`.lang` files) to determine in which languages a page will be available. This will almost always be what we want for a page. But sometimes we need to explicitly state the locales available for a page. The *impresum* page for example is only available in German and the template itself has German hard-coded into it since we don’t need it to be translated into any other languages. In cases like these we can send a list of locale codes with the template context and it will be the final list. This can be accomplished in a few ways depending on how the view is coded.

For a plain view function, you can simply pass a list of locale codes to `l10n_utils.render` in the context using the name `active_locales`. This will be the full list of available translations. Use `add_active_locales` if you want to add languages to the existing list:

```
def french_and_german_only(request):
    return l10n_utils.render(request, 'home.html', {'active_locales': ['de', 'fr']})
```

If you don't need a custom view and are just using the `page()` helper function in your `urls.py` file, then you can similarly pass in a list:

```
page('about', 'about.html', active_locales=['en-US', 'es-ES']),
```

Or if your view is even more fancy and you're using a Class-Based-View that inherits from `LangFilesMixin` (which it must if you want it to be translated) then you can specify the list as part of the view Class definition:

```
class MyView(LangFilesMixin, View):
    active_locales = ['zh-CN', 'hi-IN']
```

Or in the `urls.py` when using a CBV:

```
url(r'^about/$', MyView.as_view(active_locales=['de', 'fr'])),
```

The main thing to keep in mind is that if you specify `active_locales` that will be the full list of localizations available for that page. If you'd like to add to the existing list of locales generated from the lang files then you can use the `add_active_locales` name in all of the same ways as `active_locales` above. It's a list of locale codes that will be added to the list already available. This is useful in situations where we would have needed the l10n team to create an empty `.lang` file with an active tag in it because we have a locale-specific-template with text in the language hard-coded into the template and therefore do not otherwise need a `.lang` file.

1.7.8 Development

In local development environments and on demo servers all `l10n_has_tag` calls evaluate to true. If the content has not been translated it will display the English strings.

To test l10n locally you can set `DEV=False` in your `.env` file.

If you are running your local server you will need to restart it after altering your `.env` file.

1.8 Developing on Bedrock

1.8.1 Writing URL Patterns

URL patterns should be as strict as possible. It should begin with a `^` and end with `/` to make sure it only matches what you specify. It also forces a trailing slash. You should also give the URL a name so that other pages can reference it instead of hardcoding the URL. Example:

```
url(r'^channel/$', channel, name='mozorg.channel')
```

Bedrock comes with a handy shortcut to automate all of this:

```
from bedrock.mozorg.util import page
page('channel', 'mozorg/channel.html')
```

You don't even need to create a view. It will serve up the specified template at the given URL (the first parameter). You can also pass template data as keyword arguments:

```
page('channel', 'mozorg/channel.html', latest_version=product_details.firefox_versions['LATEST_FIREFOX_VERSION'])
```

The variable *latest_version* will be available in the template.

1.8.2 Embedding images

Images should be included on pages using helper functions.

static()

For a simple image, the *static()* function is used to generate the image URL. For example:

```

```

will output an image:

```

```

high_res_img()

For images that include a high-resolution alternative for displays with a high pixel density, use the *high_res_img()* function:

```
high_res_img('firefox/new/firefox-logo.png', {'alt': 'Firefox', 'width': '200',  
→ 'height': '100'})
```

The *high_res_img()* function will automatically look for the image in the URL parameter suffixed with *-high-res*, e.g. *firefox/new/firefox-logo-high-res.png* and switch to it if the display has high pixel density.

high_res_img() supports localized images by setting the *l10n* parameter to *True*:

```
high_res_img('firefox/new/firefox-logo.png', {'l10n': True, 'alt': 'Firefox', 'width  
→ ': '200', 'height': '100'})
```

When using localization, *high_res_img()* will look for images in the appropriate locale folder. In the above example, for the *de* locale, both standard and high-res versions of the image should be located at *media/img/l10n/de/firefox/new/*.

l10n_img()

Images that have translatable text can be handled with *l10n_img()*:

```

```

The images referenced by *l10n_img()* must exist in *media/img/l10n/*, so for above example, the images could include *media/img/l10n/en-US/firefox/os/have-it-all/messages.jpg* and *media/img/l10n/es-ES/firefox/os/have-it-all/messages.jpg*.

platform_img()

Finally, for outputting an image that differs depending on the platform being used, the `platform_img()` function will automatically display the image for the user's browser:

```
platform_img('firefox/new/browser.png', {'alt': 'Firefox screenshot'})
```

`platform_img()` will automatically look for the images `browser-mac.png`, `browser-win.png`, `browser-linux.png`, etc. Platform image also supports hi-res images by adding `'high-res': True` to the list of optional attributes.

`platform_img()` supports localized images by setting the `'l10n'` parameter to `True`:

```
platform_img('firefox/new/firefox-logo.png', {'l10n': True, 'alt': 'Firefox screenshot
→'})
```

When using localization, `platform_img()` will look for images in the appropriate locale folder. In the above example, for the `es-ES` locale, all platform versions of the image should be located at `media/img/l10n/es-ES/firefox/new/`.

1.8.3 Writing Views

You should rarely need to write a view for mozilla.org. Most pages are static and you should use the `page` expression documented above.

If you need to write a view and the page has a newsletter signup form in the footer (most do), make sure to handle this in your view. Bedrock comes with a function for doing this automatically:

```
from bedrock.mozorg.util import handle_newsletter
from django.views.decorators.csrf import csrf_exempt

@csrf_exempt
def view(request):
    ctx = handle_newsletter(request)
    return l10n_utils.render(request, 'app/template.html', ctx)
```

You'll notice a few other things in there. You should use the `l10n_utils.render` function to render templates because it handles special l10n work for us. Since we're handling the newsletter form post, you also need the `csrf_exempt` decorator.

Make sure to namespace your templates by putting them in a directory named after your app, so instead of `templates/template.html` they would be in `templates/blog/template.html` if `blog` was the name of your app.

Variation Views

We have a generic view that allows you to easily create and use a/b testing templates. If you'd like to have either separate templates or just a template context variable for switching, this will help you out. For example:

```
# urls.py

from django.conf.urls import url

from bedrock.utils.views import VariationTemplateView

urlpatterns = [
    url(r'^testing/$',
        VariationTemplateView.as_view(template_name='testing.html',
```

(continues on next page)

(continued from previous page)

```

        name='testing'),
        template_context_variations=['a', 'b']),
    ]

```

This will give you a context variable called *variation* that will either be an empty string if no param is set, or *a* if *?v=a* is in the URL, or *b* if *?v=b* is in the URL. No other options will be valid for the *v* query parameter and *variation* will be empty if any other value is passed in for *v* via the URL. So in your template code you'd simply do the following:

```
{% if variation == 'b' %}<p>This is the B variation of our test. Enjoy!</p>{% endif %}
```

If you'd rather have a fully separate template for your test, you can use the *template_name_variations* argument to the view instead of *template_context_variations*:

```

# urls.py

from django.conf.urls import url

from bedrock.utils.views import VariationTemplateView

urlpatterns = [
    url(r'^testing/$',
        VariationTemplateView.as_view(template_name='testing.html',
                                     template_name_variations=['1', '2']),
        name='testing'),
]

```

This will not provide any extra template context variables, but will instead look for alternate template names. If the URL is *testing/?v=1*, it will use a template named *testing-1.html*, if *v=2* it will use *testing-2.html*, and for everything else it will use the default. It simply puts a dash and the variation value between the template file name and file extension.

It is theoretically possible to use the template name and template context versions of this view together, but that would be an odd situation and potentially inappropriate for this utility.

You can also limit your variations to certain locales. By default the variations will work for any localization of the page, but if you supply a list of locales to the *variation_locales* argument to the view then it will only set the variation context variable or alter the template name (depending on the options explained above) when requested at one of said locales. For example, the template name example above could be modified to only work for English or German like so:

```

# urls.py

from django.conf.urls import url

from bedrock.utils.views import VariationTemplateView

urlpatterns = [
    url(r'^testing/$',
        VariationTemplateView.as_view(template_name='testing.html',
                                     template_name_variations=['1', '2'],
                                     variation_locales=['en-US', 'de']),
        name='testing'),
]

```

Any request to the page in for example French would not use the alternate template even if a valid variation were given in the URL.

Note: If you'd like to add this functionality to an existing Class-Based View, there is a mixin that implements this pattern that should work with most views: *bedrock.utils.views.VariationMixin*.

1.8.4 Coding Style Guides

- [Mozilla Python Style Guide](#)
- [Mozilla HTML Style Guide](#)
- [Mozilla JS Style Guide](#)
- [Mozilla CSS Style Guide](#)

Use the `.open-sans`, `.open-sans-light` and `.open-sans-extrabold` mixins to specify font families to allow using international fonts. See the `:ref: CSS<110n>` section in the 110n doc for details.

Use the `.font-size()` mixin to generate root-relative font sizes. You can declare a font size in pixels and the mixin will convert it to an equivalent `rem` (root em) unit while also including the pixel value as a fallback for older browsers that don't support `rem`. This is preferable to declaring font sizes in either fixed units (`px`, `pt`, etc) or element-relative units (`em`, `%`). See [this post](#) by [Jonathan Snook](#) for more info.

When including CSS blocks, use `{% block page_css %}` for page specific inclusion of CSS. `{% block site_css %}` should only be touched in rare cases where base styles need to be overwritten.

1.8.5 Configuring your code editor

Bedrock includes an `.editorconfig` file in the root directory that you can use with your code editor to help maintain consistent coding styles. Please see [editorconfig.org](#) for a list of supported editors and available plugins.

1.9 How to contribute

Before diving into code it might be worth reading through the *Developing on Bedrock* documentation, which contains useful information and links to our coding guidelines for Python, Django, JavaScript and CSS.

1.9.1 Git workflow

When you want to start contributing, you should create a branch from master. This allows you to work on different project at the same time:

```
git checkout master
git checkout -b topic-branch
```

To keep your branch up-to-date, assuming the mozilla repository is the remote called mozilla:

```
git fetch mozilla
git checkout master
git merge mozilla/master
git checkout topic-branch
git rebase master
```

If you need more Git expertise, a good resource is the [Git book](#).

Once you're done with your changes, you'll need to describe those changes in the commit message.

1.9.2 Git commit messages

Commit messages are important when you need to understand why something was done.

- First, learn [how to write good git commit messages](#).
- All commit messages must include a bug number. You can put the bug number on any line, not only the first one.
- If you use the syntax `bug xxx`, Github will reference the commit into Bugzilla. With `fix bug xxx`, it will even close the bug once it goes into master.

If you're asked to change your commit message, you can use these commands:

```
git commit --amend
# -f is doing a force push because you modified the history
git push -f my-remote topic-branch
```

1.9.3 Submitting your work

In general, you should submit your work with a pull request to master. If you are working with other people or you want to put your work on a demo server, then you should be working on a common topic branch.

Once your code has been positively reviewed, it will be deployed shortly after. So if you want feedback on your code but it's not ready to be deployed, you should note it in the pull request.

1.9.4 Squashing your commits

Should your pull request contain more than one commit, sometimes we may ask you to squash them into a single commit before merging. You can do this with *git rebase*.

As an example, let's say your pull request contains two commits. To squash them into a single commit, you can follow these instructions:

```
git rebase -i HEAD~2
```

You will then get an editor with your two commits listed. Change the second commit from *pick* to *fixup*, then save and close. You should then be able to verify that you only have one commit now with *git log*.

To push to GitHub again, because you "altered the history" of the repo by merging the two commits into one, you'll have to *git push -f* instead of just *git push*.

1.9.5 Server architecture

Demos

- *URLs*: - <http://www-demo1.allizom.org/> - <http://www-demo2.allizom.org/> - <http://www-demo3.allizom.org/> - <http://www-demo4.allizom.org/> - <http://www-demo5.allizom.org/>
- *Bedrock locales*: dev repo
- *Bedrock Git branch*: demo/1, demo/2, etc.

On-demand demos

- *URLs*: Demo instances can also be spun up on-demand by pushing a branch to the mozilla bedrock repo that matches a specific naming convention (the branch name must start with `demo/`). Jenkins will then automate spinning up a demo instance based on that branch. For example, pushing a branch named `demo/feature` would create a demo instance with the following URL: `https://bedrock-demo-feature.us-west.moz.works/`
- *Bedrock locales*: dev repo
- *Bedrock Git branch*: any branch named starting with `demo/`

Note: Deployed demo instances are not yet automatically cleaned up when branches are deleted, so to avoid lots of instances piling up it is currently recommended to try and limit a single demo instance per developer, reusing a branch such as `demo/<your_username>`.

Dev

- *URL*: <http://www-dev.allizom.org/>
- *Bedrock locales*: dev repo
- *Bedrock Git branch*: master, deployed on git push

Stage

- *URL*: <http://www.allizom.org/>
- *Bedrock locales*: prod repo
- *Bedrock Git branch*: prod, deployed on git push with date-tag

Production

- *URL*: <http://www.mozilla.org/>
- *Bedrock locales*: prod repo
- *Bedrock Git branch*: prod, deployed on git push with date-tag

You can check the currently deployed git commit by checking <https://www.mozilla.org/revision.txt>.

1.9.6 Pushing to production

We're doing pushes as soon as new work is ready to go out.

After doing a push, those who are responsible for implementing changes need to update the bugs that have been pushed with a quick message stating that the code was deployed.

If you'd like to see the commits that will be deployed before the push run the following command:

```
./bin/open-compare.py
```

This will discover the currently deployed git hash, and open a compare URL at github to the latest master. Look at `open-compare.py -h` for more options.

We automate pushing to production via tagged commits (see *Push to prod branch (tagged)*)

1.10 Continuous Integration & Deployment

Bedrock runs a series of automated tests as part of continuous integration workflow and [Deployment Pipeline](#). You can learn more about each of the individual test suites by reading their respective pieces of documentation:

- Python unit tests (see [Run the tests](#)).
- JavaScript unit tests (see [Front-end testing](#)).
- Redirect tests (see [Testing redirects](#)).
- Functional tests (see [Front-end testing](#)).

1.10.1 Tests in the lifecycle of a change

Below is an overview of the tests during the lifecycle of a change to bedrock:

Local development

The change is developed locally, and all integration tests can be executed against a locally running instance of the application.

Pull request

Once a pull request is submitted, [CircleCI](#) will run both the Python and JavaScript unit tests, as well as the suite of redirect headless HTTP(s) response checks.

Push to master branch

Whenever a change is pushed to the master branch, a new image is built and deployed to the dev environment, and the full suite of headless and UI tests are then run against Firefox on Windows 10 using [Sauce Labs](#). This is handled by the pipeline, and is subject to change according to the settings in the [master.yml](#) file in the repository.

Push to stage branch

Whenever a change is pushed to the stage branch, a production docker image is built, published to [Docker Hub](#), and deployed to a [public staging environment](#). Once the new image is deployed, the full suite of UI tests is run against it with Firefox, Chrome, and Internet Explorer on Windows 10, and the sanity suite is run with IE9.

Push to prod branch (tagged)

When a tagged commit is pushed to the prod branch, a production docker image is built and published to [Docker Hub](#) if needed (usually this will have already happened as a result of a push to the stage branch), and deployed to each [production](#) deployment. After each deployment is complete, the full suite of UI tests is run against Firefox, Chrome and Internet Explorer on Windows 10, and the sanity suite is run against IE9. As with untagged pushes, this is all handled by the pipeline, and is subject to change according to the settings in the [prod.yml](#) file in the repository.

Push to prod cheat sheet

1. Check out the `master` branch
2. Make sure the `master` branch is up to date with `mozilla/bedrock master`

3. Check that dev deployment is green:

1. View [deployment pipeline](#) and look at `master` branch
4. Tag and push the deployment by running `bin/tag-release.sh --push`

Note: By default the `tag-release.sh` script will push to the `origin` git remote. If you'd like for it to push to a different remote name you can either pass in a `-r` or `--remote` argument, or set the `MOZ_GIT_REMOTE` environment variable. So the following are equivalent:

```
$ bin/tag-release.sh --push -r mozilla
$ MOZ_GIT_REMOTE=mozilla bin/tag-release.sh --push
```

And if you'd like to just tag and not push the tag anywhere, you may omit the `--push` parameter.

1.10.2 Pipeline integration

Our [Jenkinsfile](#) will run the integration tests based on information in our [branch-specific YAML files](#). These files specify various test names per branch that will cause it to use different parameters, allowing it to be called in many different ways to cover the testing needs. The job executes [this script](#), which then runs [this Docker image](#), and ultimately runs [another script](#). The two scripts can also be executed locally to replicate the way Jenkins operates.

During the **Test Images** stage, the Test Runner job is called without a `BASE_URL`. This means that a local instance of the application will be started, and the URL of this instance will be used for testing. The `DRIVER` parameter is set to `Remote`, which causes a local instance of Selenium Grid to be started in Docker and used for the browser-based functional UI tests.

The test scripts above will be run once for each properties name specified in the [branch-specific YAML files](#) for the branch being built and tested. Pushes to `master` will run different tests than pushes to `prod` for example.

Configuration

Many of the options are configured via environment variables passed from the initial script to the Docker image and onto the final script. Many of these options can be set in the [branch-specific YAML files](#) in the repository. In the [branch-specific YAML files](#) folder you can copy any file there to match the name of your branch and modify it to set how it should be built by jenkins. Take the following example:

```
# jenkins/branches/change-all-the-things.yml
apps:
  - bedrock-probably-broken
```

This configuration would cause commits pushed to a branch named `change-all-the-things` to have docker images built for them, have the unit tests run, and deploy to a deis app named `bedrock-probably-broken` in our `us-west` deis cluster. If you'd like it to create the deis app and pre-fill a local database for your app, you can set `demo: true` in the file. Note that if the app already exists it must have the `jenkins` user added via the `deis perms:create jenkins -a <your app name>` command.

The available branch configuration options are as follows:

- `push_public_registry`: boolean. Set to `true` to cause the built images to be pushed to the public docker hub.
- `require_tag`: boolean. Set to `true` to require that the commit being built have a git tag in the format `YYYY-MM-DD.X`.

- `regions`: list. A list of strings indicating the deployment regions for the set of apps. The valid values are in the `regions` area of the `jenkins/global.yml` file. If omitted a deployment to only `oregon-b` is assumed.
- `apps`: list. A list of strings indicating the deis app name(s) to which to deploy. If omitted no deployments will occur.
- `demo`: boolean. Set to `true` to have the deployed app in demo mode, which means it will have a pre-filled local database and the deis app will be created and configured for you if it doesn't already exist.
- `integration_tests`: list. A list of strings indicating the types of integration tests to run. If omitted no tests will run.

Configure Demo Servers

You can also set app configuration environment variables via deployment as well for demos. The default environment variables are set in `jenkins/branches/demo/default.env`. To modify your app's settings you can create an env file named after your branch (e.g. `jenkins/branches/demo/pmac-110n.env` for the branch `demo/pmac-110n.env`). The combination of values from `demo/default.env`, your branch specific env file, and a region specific env file (e.g. `jenkins/regions/virginia.env`) will be used to configure the app. So you only need to add the variables that differ from the default files to your file, and you can override any values from the default files as well.

Instance Configuration & Switches

Beyond setting environment variables in demo instances as described above, we have a [separate repo](#) for configuring our primary instances (dev, stage, and prod). The [docs for updating configurations](#) in that repo are on their own page, but there is a way to tell what version of the configuration is in use on any particular instance of bedrock. You can go to the `/healthz-cron/` URL on an instance (see [prod](#) for example) to see the current commit of all of the external Git repos in use by the site and how long ago they were updated. The info on that page also includes the latest version of the database in use, the git revision of the bedrock code, and how long ago the database was updated. If you recently made a change to one of these repos and are curious if the changes have made it to production, this is the URL you should check.

Updating Selenium

There are two components for Selenium, which are independently versioned. The first is the Python client, and this can be updated via the [test dependencies](#). The other component is the server, which in the pipeline is either provided by a Docker container or [Sauce Labs](#). The `SELENIUM_VERSION` environment variable controls both of these, and they should ideally use the same version, however it's possible that availability of versions may differ. You can check the [Selenium Docker versions](#) available. If needed, the global default can be set and then can be overridden in the individual job configuration.

Adding test runs

Test runs can be added by creating a new properties section in the [integration tests script](#) with the parameters of the new test run. This is simply a bash script and you can duplicate a clause of the case statement. For example, if you wanted to run tests in Firefox on both Windows 10 and OS X, you could create the following clauses:

```
case $1 in
  osx-firefox)
    BROWSER_NAME=firefox
    PLATFORM="OS X 10.11"
  ;;
```

(continues on next page)

(continued from previous page)

```
win10-firefox)
  BROWSER_NAME=firefox
  PLATFORM="Windows 10"
;;
```

You can use [Sauce Labs platform configurator](#) to help with the parameter values.

If you have commit rights to our Github repo (mozilla/bedrock) you can simply push your branch to the branch named `run-integration-tests`, and the `bedrock-integration-tests` app will be deployed and all of the integration tests defined in the `jenkins.yml` file for that branch will be run. Please announce in our IRC channel ([#www on irc.mozilla.org](#)) that you'll be doing this so that we don't get conflicts.

1.10.3 Known issues in Jenkins

Jenkins stalls after global configuration changes

When using the IRC plugin for notifications, global configuration changes can cause Jenkins to become unresponsive. To make such changes it can be necessary to first restart Jenkins, as this issue only appears some time after Jenkins has been started. A [bug for the IRC plugin](#) has been raised.

1.11 Front-end testing

Bedrock runs a suite of front-end [Jasmine](#) behavioral/unit tests, which use [Karma](#) as a test runner. We also have a suite of functional tests using [Selenium](#) and [pytest](#). This allows us to emulate users interacting with a real browser. All these test suites live in the `tests` directory.

The `tests` directory comprises of:

- `/functional` contains `pytest` tests.
- `/pages` contains Python page objects.
- `/unit` contains the `Jasmine` tests and `Karma` config file.

1.11.1 Installation

First follow the [installation instructions for bedrock](#), which will install the specific versions of `Jasmine/Karma` which are needed to run the unit tests, and guide you through installing `pip` and setting up a virtual environment for the functional tests. The additional requirements can then be installed by using the following commands:

```
$ source venv/bin/activate
$ pip install -r requirements/dev.txt
```

1.11.2 Running Jasmine tests using Karma

To perform a single run of the `Jasmine` test suite using `Firefox`, type the following command:

```
$ gulp js:test
```

See the [Jasmine](#) documentation for tips on how to write JS behavioral or unit tests. We also use [Sinon](#) for creating test spies, stubs and mocks.

1.11.3 Running functional tests

Note: Before running the functional tests, please make sure to follow the bedrock [installation docs](#), including the database sync that is needed to pull in external data such as event/blog feeds etc. These are required for some of the tests to pass. To run the tests using Firefox, you must also first download [geckodriver](#) and make it available in your [system path](#). You can alternatively specify the path to geckodriver using the command line (see the [pytest-selenium documentation](#) for more information).

To run the full functional test suite against your local bedrock instance:

```
$ py.test --base-url http://localhost:8000 --driver Firefox --html tests/functional/
↳results.html tests/functional/
```

This will run all test suites found in the `tests/functional` directory and assumes you have bedrock running at localhost on port 8000. Results will be reported in `tests/functional/results.html`.

Note: If you omit the `--base-url` command line option then a local instance of bedrock will be started, however the tests are not currently able to run against bedrock in this way.

By default, tests will run one at a time. This is the safest way to ensure predictable results, due to [bug 1230105](#). If you want to run tests in parallel (this should be safe when running against a deployed instance), you can add `-n auto` to the command line. Replace `auto` with an integer if you want to set the maximum number of concurrent processes.

Note: There are some functional tests that do not require a browser. These can take a long time to run, especially if they're not running in parallel. To skip these tests, add `-m 'not headless'` to your command line.

To run a single test file you must tell `py.test` to execute a specific file e.g. `tests/functional/test_newsletter.py`:

```
$ py.test --base-url http://localhost:8000 --driver Firefox --html tests/functional/
↳results.html tests/functional/test_newsletter.py
```

To run a single test you can filter using the `-k` argument supplied with a keyword e.g. `-k test_successful_sign_up`:

```
$ py.test --base-url http://localhost:8000 --driver Firefox --html tests/functional/
↳results.html tests/functional/test_newsletter.py -k test_successful_sign_up
```

You can also easily run the tests against any bedrock environment by specifying the `--base-url` argument. For example, to run all functional tests against dev:

```
$ py.test --base-url https://www-dev.allizom.org --driver Firefox --html tests/
↳functional/results.html tests/functional/
```

Note: For the above commands to work, Firefox needs to be installed in a predictable location for your operating system. For details on how to specify the location of Firefox, or running the tests against alternative browsers, refer to the [pytest-selenium documentation](#).

For more information on command line options, see the [pytest documentation](#).

Running tests in Sauce Labs

You can also run tests in Sauce Labs directly from the command line. This can be useful if you want to run tests against Internet Explorer when you're on Mac OSX, for instance.

1. Sign up for an account at <https://saucelabs.com/opensauce/>.
2. Log in and obtain your Remote Access Key from user settings.
3. Run a test specifying SauceLabs as your driver, and pass your credentials.

For example, to run the home page tests using Internet Explorer via Sauce Labs:

```
$ SAUCELABS_USERNAME=thedude SAUCELABS_API_KEY=123456789 py.test --base-url https://
↳www-dev.allizom.org --driver SauceLabs --capability browserName 'internet explorer'
↳-n auto --html tests/functional/results.html tests/functional/test_home.py
```

1.11.4 Writing Selenium tests

Tests usually consist of interactions and assertions. Selenium provides an API for opening pages, locating elements, interacting with elements, and obtaining state of pages and elements. To improve readability and maintainability of the tests, we use the [Page Object](#) model, which means each page we test has an object that represents the actions and states that are needed for testing.

Well written page objects should allow your test to contain simple interactions and assertions as shown in the following example:

```
def test_sign_up_for_newsletter(base_url, selenium):
    page = NewsletterPage(base_url, selenium).open()
    page.type_email('noreply@mozilla.com')
    page.accept_privacy_policy()
    page.click_sign_me_up()
    assert page.sign_up_successful
```

It's important to keep assertions in your tests and not your page objects, and to limit the amount of logic in your page objects. This will ensure your tests all start with a known state, and any deviations from this expected state will be highlighted as potential regressions. Ideally, when tests break due to a change in bedrock, only the page objects will need updating. This can often be due to an element needing to be located in a different way.

Please take some time to read over the [Selenium documentation](#) for details on the Python client API.

Destructive tests

By default all tests are assumed to be destructive, which means they will be skipped if they're run against a [sensitive environment](#). This prevents accidentally running tests that create, modify, or delete data on the application under test. If your test is nondestructive you will need to apply the `nondestructive` marker to it. A simple example is shown below, however you can also read the [pytest markers](#) documentation for more options.

```
import pytest

@pytest.mark.nondestructive
def test_newsletter_default_values(base_url, selenium):
    page = NewsletterPage(base_url, selenium).open()
    assert '' == page.email
    assert 'United States' == page.country
    assert 'English' == page.language
```

(continues on next page)

```
assert page.html_format_selected
assert not page.text_format_selected
assert not page.privacy_policy_accepted
```

Sanity tests

Sanity tests are considered to be our most critical tests that must pass in a wide range of web browsers, including old versions of Internet Explorer. Sanity tests are run automatically post deployment on a wider range of browsers & platforms than we run the full suite against. The number of sanity tests we run should remain small, but cover our most critical pages where legacy browser support is important. Sanity tests are typically run after a tagged commit to master (see *Push to prod branch (tagged)*).

```
import pytest

@pytest.mark.sanity
@pytest.mark.nondestructive
def test_click_download_button(base_url, selenium):
    page = FirefoxNewPage(base_url, selenium).open()
    page.download_firefox()
    assert page.is_thank_you_message_displayed
```

You can run sanity tests only by adding `-m sanity` when running the test suite on the command line.

Waits and Expected Conditions

Often an interaction with a page will cause a visible response. While Selenium does its best to wait for any page loads to be complete, it's never going to be as good as you at knowing when to allow the test to continue. For this reason, you will need to write explicit `waits` in your page objects. These repeatedly execute code (a condition) until the condition returns true. The following example is probably the most commonly used, and will wait until an element is considered displayed:

```
from selenium.webdriver.support import expected_conditions as expected
from selenium.webdriver.support.ui import WebDriverWait as Wait

Wait(selenium, timeout=10).until(
    expected.visibility_of_element_located(By.ID, 'my_element'))
```

For convenience, the Selenium project offers some basic `expected conditions`, which can be used for the most common cases.

1.11.5 Debugging Selenium

Debug information is collected on failure and added to the HTML report referenced by the `--html` argument. You can enable debug information for all tests by setting the `SELENIUM_CAPTURE_DEBUG` environment variable to `always`.

1.11.6 Guidelines for writing functional tests

- Try and keep tests organized and cleanly separated. Each page should have its own page object and test file, and each test should be responsible for a specific purpose, or component of a page.

- Avoid using sleeps - always use waits as mentioned above.
- Don't make tests overly specific. If a test keeps failing because of generic changes to a page such as an image filename or `href` being updated, then the test is probably too specific.
- Avoid string checking as tests may break if strings are updated, or could change depending on the page locale.
- When writing tests, try and run them against a staging or demo environment in addition to local testing. It's also worth running tests a few times to identify any intermittent failures that may need additional waits.

See also the [Web QA style guide](#) for Python based testing.

1.11.7 Testing Basket email forms

When writing functional tests for front-end email newsletter forms that submit to [Basket](#), we have some special case email addresses that can be used just for testing:

1. Any newsletter subscription request using the email address “[success@example.com](#)” will always return success from the basket client.
2. Any newsletter subscription request using the email address “[failure@example.com](#)” will always raise an exception from the basket client.

Using the above email addresses enables newsletter form testing without actually hitting the Basket instance, which reduces automated newsletter spam and improves test reliability due to any potential network flakiness.

1.11.8 Link Checks

A full link checker is run over the production environments, which uses a tool named [LinkChecker](#) to crawl the entire website and reports any broken or malformed links both internally and externally. These jobs are run once a day in the [Jenkins instance](#) and are named with the `bedrock_linkchecker_` prefix.

In addition, there are targeted functional tests for the [download](#) and [localized download](#) pages. These tests do not use the LinkChecker tool, and are run as part of the pipeline, which ensures that any broken download links are noticed much earlier, and also do not depend on a crawler to find them.

1.12 Managing Redirects

We have a redirects app in bedrock that makes it easier to add and manage redirects. Due to the size, scope, and history of mozilla.org we have quite a lot of redirects. If you need to add or manage redirects read on.

1.12.1 Add a redirect

You should add redirects in the app that makes the most sense. For example, if the source url is `/firefox/...` then the `bedrock.firefox` app is the best place. Redirects are added to a `redirects.py` file within the app. If the app you want to add redirects to doesn't have such a file, you can create one and it will automatically be discovered and used by bedrock as long as said app is in the `INSTALLED_APPS` setting (see `bedrock/mozorg/redirects.py` as an example).

Once you decide where it should go you can add your redirect. To do this you simply add a call to the `bedrock.redirects.util.redirect` helper function in a list named `redirectpatterns` in `redirects.py`. For example:

```
from bedrock.redirects.util import redirect

redirectpatterns = [
    redirect(r'^rubble/barny/$', '/flintstone/fred/'),
]
```

This will make sure that requests to /rubble/barny/ (or with the locale like /pt-BR/rubble/barny/) will get a 301 response sending users to /flintstone/fred/.

The `redirect()` function has several options. Its signature is as follows:

```
def redirect(pattern, to, permanent=True, locale_prefix=True, anchor=None, name=None,
             query=None, vary=None, cache_timeout=12, decorators=None):
    """
    Return a url matcher suited for urlpatterns.

    pattern: the regex against which to match the requested URL.
    to: either a url name that `reverse` will find, a url that will simply be
    ↪ returned,
        or a function that will be given the request and url captures, and return the
        destination.
    permanent: boolean whether to send a 301 or 302 response.
    locale_prefix: automatically prepend `pattern` with a regex for an optional locale
        in the url. This locale (or None) will show up in captured kwargs as 'locale'.
    anchor: if set it will be appended to the destination url after a '#'.
    name: if used in a `urls.py` the redirect URL will be available as the name
        for use in calls to `reverse()`. Does NOT work if used in a `redirects.py`
    ↪ file.
    query: a dict of query params to add to the destination url.
    vary: if you used an HTTP header to decide where to send users you should include
    ↪ that
        header's name in the `vary` arg.
    cache_timeout: number of hours to cache this redirect. just sets the proper
    ↪ `cache-control`
        and `expires` headers.
    decorators: a callable (or list of callables) that will wrap the view used to
    ↪ redirect
        the user. equivalent to adding a decorator to any other view.

    Usage:
    urlpatterns = [
        redirect(r'projects/$', 'mozorg.product'),
        redirect(r'^projects/seamonkey$', 'mozorg.product', locale_prefix=False),
        redirect(r'apps/$', 'https://marketplace.firefox.com'),
        redirect(r'firefox/$', 'firefox.new', name='firefox'),
        redirect(r'the/dude$', 'abides', query={'aggression': 'not_stand'}),
    ]
    """
```

1.12.2 Differences

This all differs from urlpatterns in urls.py files in some important ways. The first is that these happen first. If something matches in a redirects.py file it will always win the race if another url in a urls.py file would also have matched. Another is that these are matched before any locale prefix stuff happens. So what you're matching against in the redirects files is the original URL that the user requested. By default (unless you set

`locale_prefix=False`) your patterns will match either the plain url (e.g. `/firefox/os/`) or one with a locale prefix (e.g. `/fr/firefox/os/`). If you wish to include this locale in the destination URL you can simply use python's string `format()` function syntax. It is passed to the `format` method as the keyword argument `locale` (e.g. `redirect('^stuff/$', '{locale}whatnot/')`). If there was no locale in the url the `{locale}` substitution will be an empty string. Similarly if you wish to include a part of the original URL in the destination, just capture it with the regex using a named capture (e.g. `r'^stuff/(?P<rest>.*)$'` will let you do `'/whatnot/{rest}'`).

1.12.3 Utilities

There are a couple of utility functions for use in the `to` argument of `redirect` that will return a function to allow you to match something in an HTTP header.

ua_redirector

`bedrock.redirects.util.ua_redirector` is a function to be used in the `to` argument that will use a regex to match against the User-Agent HTTP header to allow you to decide where to send the user. For example:

```
from bedrock.redirects.util import redirect, ua_redirector

redirectpatterns = [
    redirect(r'^rubble/barny/$',
            ua_redirector('firefox(os)?', '/firefox/', '/not-firefox/'),
            cache_timeout=0),
]
```

You simply pass it a regex to match, the destination url (substitutions from the original URL do work) if the regex matches, and another destination url if the regex does not match. The match is not case sensitive unless you add the optional `case_sensitive=True` argument.

Note: Be sure to include the `cache_timeout=0` so that you won't be bitten by any caching proxies sending all users one way or the other. Do not set the `Vary: User-Agent` header; this will not work in production.

header_redirector

This is basically the same as `ua_redirector` but works against any header. The arguments are the same as above except that there is an additional first argument for the name of the header:

```
from bedrock.redirects.util import redirect, header_redirector

redirectpatterns = [
    redirect(r'^rubble/barny/$',
            header_redirector('cookie', 'been-here', '/firefox/', '/firefox/new/'),
            vary='cookie'),
]
```

1.12.4 Testing redirects

A suite of tests exists for redirects, which is intended as a reference of the redirects we expect to work on www.mozilla.org. This will become a base for implementing these redirects in the bedrock app and allow us to test them before release.

Installation

First follow the *installation instructions for bedrock*, which will guide you through installing pip and setting up a virtual environment for the tests. The additional requirements can then be installed by using the following commands:

```
$ source venv/bin/activate
$ pip install -r requirements/dev.txt
```

Running the tests

If you wish to run the full set of tests, which requires a deployed instance of the site (e.g. www.mozilla.org) you can set the `--base-url` command line option:

```
$ py.test --base-url https://www.mozilla.org tests/redirects/
```

By default, tests will run one at a time. If you intend to run the suite against a remote instance of the site (e.g. production) it will run a lot quicker by running the tests in parallel. To do this, you can add `-n auto` to the command line. Replace `auto` with an integer if you want to set the maximum number of concurrent processes.

1.13 JavaScript Libraries

1.13.1 Mozilla Pager

mozilla-pager.js

1.13.2 Mozilla Accordion

mozilla-accordion.js

1.13.3 Mozilla Image Lazy Loader

mozilla-lazy-load.js

1.13.4 Mozilla Notification Banner

mozilla-notification-banner.js

1.14 Newsletters

Bedrock includes support for signing up for and managing subscriptions and preferences for Mozilla newsletters.

By default, every page's footer has a form to signup for the default newsletter, "Firefox & You".

1.14.1 Features

- ability to subscribe to a newsletter from a page's footer area. Many pages on the site might include this.
- whole pages devoted to subscribing to one newsletter, often with custom text, branding, and layout
- newsletter preference center - allow user to change their email address, preferences (e.g. language, HTML vs. text), which newsletters they're subscribed to, etc. Access is limited by requiring a user-specific token in the URL (it's a UUID). The full URL is included as a link in each newsletter sent to the user, which is the only way (currently) they can get the token.
- landing pages that user ends up on after subscribing. These can vary depending on where they're coming from.

1.14.2 Newsletters

Newsletters have a variety of characteristics. Some of these are implemented in Bedrock, others are transparent to Bedrock but implemented in the basket back-end that provides our interface to the newsletter vendor.

- Public name - the name that is displayed to users, e.g. "Firefox Weekly Tips".
- Internal name- a short string that is used internal to Bedrock and basket to identify a newsletter. Typically these are lowercase strings of words joined by hyphens, e.g. "firefox-tips". This is what we send to basket to identify a newsletter, e.g. to subscribe a user to it.
- Show publicly - pages like the newsletter preferences center show a list of unsubscribed newsletters and allow subscribing to them. Some newsletters aren't included in that list by default (though they are shown if the user is already subscribed, to let them unsubscribe).
- Languages - newsletters are available in a particular set of languages. Typically when subscribing to a newsletter, a user can choose their preferred language. We should try not to let them subscribe to a newsletter in a language that it doesn't support.

The backend only stores one language for the user though, so whenever the user submits one of our forms, whatever language they last submitted is what is saved for their preference for everything.

- Welcome message - each newsletter can have a canned welcome message that is sent to a user when they subscribe to it. Newsletters should have both an HTML and a text version of this.
- Drip campaigns - some newsletters implement so-called drip campaigns, in which a series of canned messages are dribbled out to the user over a period of time. E.g. 1 week after subscribing, they might get message 1; a week later, message 2, and so on until all the canned messages have been sent.

Because drip campaigns depend on the signup date of the user, we're careful not to accidentally change the signup date, which could happen if we sent redundant subscription commands to our backend.

1.14.3 Bedrock and Basket

Bedrock is the user-facing web application. It presents an interface for users to subscribe and manage their subscriptions and preferences. It does not store any information. It gets all newsletter and user-related information, and makes updates, via web requests to the Basket server.

The Basket server implements an HTTP API for the newsletters. The front-end (Bedrock) can make calls to it to retrieve or change users' preferences and subscriptions, and information about the available newsletters. Basket implements some of that itself, and other functions by calling the newsletter vendor's API. Details of that are outside the scope of this document, but it's worth mentioning that both the user token (UUID) and the newsletter internal name mentioned above are used only between Bedrock and Basket.

1.14.4 URLs

Here are a few important URLs implemented. These were established before Bedrock came along and so are unlikely to be changed.

(Not all of these might be implemented in Bedrock yet.)

/newsletter/ - subscribe to 'mozilla-and-you' newsletter (public name: "Firefox & You")

/newsletter/existing/USERTOKEN/ - user management of their preferences and subscriptions

1.14.5 Configuration

Currently, information about the available newsletters is configured in Basket. See Basket for more information.

1.14.6 Footer signup

Customize the footer signup form by overriding the `email_form` template block. For example, to have no signup form:

```
{% block email_form %}{% endblock %}
```

The default is:

```
{% block email_form %}{{ email_newsletter_form() }}{% endblock %}
```

which gives a signup for Firefox & You. You can pass parameters to the macro `email_newsletter_form` to change that. For example, the `newsletter_id` parameter controls which newsletter is signed up for, and `title` can override the text:

```
{% block email_form %}
    {{ email_newsletter_form('app-dev',
                             _('Sign up for more news about the Firefox Marketplace.
→')) }}
{% endblock %}
```

Pages can control whether country or language fields are included by passing `include_language=[True|False]` and/or `include_country=[True|False]`.

You can also use the same form outside a page footer by passing `footer=False` to the macro.

You can also specify one of three color variants for the "Sign Up Now" button. The options are:

- default - Which sets the border and font color to a light blue [#00afe5]
- dark - Which sets the border and font color to the dark Firefox blue [00539F]
- white - Which sets the border and font color to white [fff]

This is done in a template as follows:

```
# default
{% block email_form %}
    {{ email_newsletter_form() }}
{% endblock %}

# dark
{% block email_form %}
    {{ email_newsletter_form(button_class='button-dark') }}
{% endblock %}

# white
{% block email_form %}
    {{ email_newsletter_form(button_class='button-light') }}
{% endblock %}
```

1.14.7 Creating a signup page

Start with a template that extends 'newsletter/one_newsletter_signup.html'. It's probably simplest to copy an existing one, like 'newsletter/mobile.html'.

Set the `newsletter_title` and `newsletter_id` variables and override at least the `page_title` and `newsletter_content` blocks:

```
{% set newsletter_title = _('Firefox and You') %}
{% set newsletter_id = 'mozilla-and-you' %}

{% block page_title %}{{ newsletter_title }}{% endblock %}

{% block newsletter_content %}
<div id="main-feature">
  <h2>Subscribe to <span>about:mobile</span>!</h2>
  <p>Our about:mobile newsletter brings you the latest and greatest news
    from the Mozilla contributor community.
  </p>
</div>
{% endblock %}
```

Then add a url to `newsletter/urls.py`:

```
# "about:mobile"
page('newsletter/about_mobile', 'newsletter/mobile.html'),
```

1.15 Using External Content Cards Data

The [www-admin](#) repo contains data files and images that are synced to bedrock and available for use on any page. The docs for updating said data is available via that repo, but this page will explain how to use the cards data once it's in the bedrock database.

1.15.1 Add to a View

The easiest way to make the data available to a page is to add the `page_content_cards` variable to the template context:

```
from bedrock.contentcards.models import get_page_content_cards

def view_with_cards(request):
    locale = l10n_utils.get_locale(request)
    ctx = {'page_content_cards': get_page_content_cards('home', locale)}
    return l10n_utils.render(request, 'sweet-words.html', ctx)
```

The `get_page_content_cards` returns a dict of card data dicts for the given page (home in this case) and locale. The dict keys are the names of the cards (e.g. `card_1`). If the `page_content_cards` context variable is available in the template, then the `content_card()` macro will discover it automatically.

Note: The `get_page_content_cards` function is not all that clever as far as l10n is concerned. If you have translated the cards in the `www-admin` repo that is great, but you should have cards for every locale for which the page is active or the function will return an empty dict. This is especially tricky if you have multiple English locales enabled (en-US, en-CA, en-GB, etc.) and want the same cards to be used for all of them. You'd need to do something like `if locale.startswith('en-')`: then use `en-US` in the function call.

Alternately you could just wrap the section of the template using cards to be optional in an `{% if page_content_cards %}` statement, and that way it will not show the section at all if the dict is empty if there are no cards for that page and locale combination.

1.15.2 Add to the Template

Once you have the data in the template context, using a card is simple:

```
{% from "macros-protocol.html" import content_card with context %}

{{ content_card('card_1') }}
```

This will insert the data from the `card_1.en-US.md` file from the `www-admin` repo into the template via the `card()` macro normally used for protocol content cards.

If you don't have the `page_content_cards` variable in the template context and you don't want to create or modify a view, you can fetch the cards via a helper function in the template itself, but you have to pass the result to the macro:

```
{% from "macros-protocol.html" import content_card with context %}
{% set content_cards = get_page_content_cards('home', LANG) %}

{{ content_card('card_1', content_cards) }}
```

1.16 Tabzilla

The latest version of Tabzilla is no longer part of the bedrock repository, and can now be included directly as part of any Mozilla site. Please see the Tabzilla repository for more information:

<https://github.com/mozilla/tabzilla/>

1.16.1 Translation and Update bar

The Translation and Update Bar functionality that used to be part of Tabzilla can now also be found in its own repository:

<https://github.com/mozilla/mozilla-infobar/>

1.17 Mozilla.UITour

1.17.1 Introduction

`Mozilla.UITour` is a JS library that exposes an event-based Web API for communicating with the Firefox browser chrome. It can be used for tasks such as opening menu panels and highlighting the position of buttons in the toolbar. It is supported in Firefox 29 onward.

For security reasons `Mozilla.UITour` will only work on white-listed domains and over a secure connection. The white-listed domains are <https://www.mozilla.org> and <https://support.mozilla.org> and the special `about:home` page.

The `Mozilla.UITour` library is maintained on [Mozilla Central](#).

Important: The API is supported only on the desktop versions of Firefox. It doesn't work on Firefox for Android and iOS.

1.17.2 Local development

To develop or test using `Mozilla.UITour` locally you need to create some custom preferences in `about:config`.

- `browser.uitour.testingOrigins` (string) (value: local address e.g. `http://127.0.0.1:8000`)
- `browser.uitour.requireSecure` (boolean) (value: `false`)

Note that `browser.uitour.testingOrigins` can be a comma separated list of domains, e.g.

```
'http://127.0.0.1:8000, https://www-demo2.allizom.org'
```

Important: Prior to Firefox 36, the testing preference was called `browser.uitour.whitelist.add.testing` (Bug 1081772). This old preference does not accept a comma separated list of domains, and you must also exclude the domain protocol e.g. `https://`. A browser restart is also required after adding a whitelisted domain.

1.17.3 Tracking Protection UITour

In order to test the Tracking Protection tour on a local server or domain other than `www.mozilla.org`, you must first set an additional preference in `about:config` in addition to white listing `UITour` for the domain.

- `privacy.trackingprotection.introURL` value e.g. `http://127.0.0.1:8000/%LOCALE%/firefox/%VERSION%/tracking-protection/start/`

Once this preference has been set, the tour can be accessed by opening a new Private Window and then by clicking the “See how this works” CTA button.

1.17.4 JavaScript API

registerPageID(pageId)

Register an ID for use in `Telemetry`. `pageId` must be a string unique to the page:

```
var pageId = 'firstrun-page-firefox-29';  
Mozilla.UITour.registerPageID(pageId);
```

showHighlight(target, effect)

Highlight a button in the browser chrome. `target` is the string ID for the button and `effect` is the animation type:

```
Mozilla.UITour.showHighlight('appMenu', 'wobble');
```

Target types:

- 'accountStatus'
- 'addons'
- 'appMenu'
- 'backForward'
- 'bookmarks'
- 'customize'
- 'help'
- 'home'
- 'quit'
- 'search'
- 'searchIcon' (Firefox 34 and above)
- 'urlbar'
- 'loop'
- 'forget'
- 'privateWindow'
- 'trackingProtection' (Firefox 42 and above)
- 'controlCenter-trackingUnblock' (Firefox 42 and above)
- 'controlCenter-trackingBlock' (Firefox 42 and above)

Effect types:

- 'random'
- 'wobble'
- 'zoom'
- 'color'
- 'none' (default)

hideHighlight()

Hides the currently visible highlight:

```
Mozilla.UITour.hideHighlight();
```

showInfo(target, title, text, icon, buttons, options)

Displays a customizable information panel pointing to a given target:

```

var buttons = [
  {
    label: 'Cancel',
    style: 'link',
    callback: cancelBtnCallback
  },
  {
    label: 'Confirm',
    style: 'primary',
    callback: confirmBtnCallback
  }
];

var icon = '//mozorg.cdn.mozilla.net/media/img/firefox/australis/logo.png';

var options = {
  closeButtonCallback: closeBtnCallback
};

Mozilla.UITour.showInfo('appMenu', 'my title', 'my text', icon, buttons, options);

```

Available targets:

Any target that can be highlighted can have an information panel attached.

Additional parameters:

- title panel title (string).
- text panel description (string).
- icon panel icon absolute url (string). Icon should be 48px x 48px.
- buttons array of buttons (object)
- options (object)

buttons array items can have the following properties:

- label button text (string)
- icon button icon url (string)
- style button style can be either *primary* or *link* (string)
- callback to be executed when the button is clicked (function)
- options (object)

options can have the following properties:

- closeButtonCallback to be executed when the (x) close button is clicked (function)

hideInfo()

Hides the currently visible info panel:

```
Mozilla.UITour.hideInfo();
```

showMenu(target, callback)

Opens a targeted menu in the browser chrome.

```
Mozilla.UITour.showMenu('appMenu', function() {  
    console.log('menu was opened');  
});
```

Available targets:

- 'appMenu'
- 'bookmarks'
- 'loop' (Firefox 35 and above)
- 'controlCenter' (Firefox 42 and above)

Optional parameters:

- callback function to be called when the menu was successfully opened.

hideMenu(target)

```
Mozilla.UITour.hideMenu('appMenu');
```

Closes a menu panel.

previewTheme(theme)

Previews a Firefox theme. theme should be a JSON literal:

```
var theme = {  
    "category":      "Firefox",  
    "iconURL":      "https://addons.mozilla.org/_files/18066/preview_small.jpg?  
↪1241572934",  
    "headerURL":    "https://addons.mozilla.org/_files/18066/1232849758499.jpg?  
↪1241572934",  
    "name":         "Dark Fox",  
    "author":       "randomaster",  
    "footer":       "https://addons.mozilla.org/_files/18066/1232849758500.jpg?  
↪1241572934",  
    "previewURL":   "https://addons.mozilla.org/_files/18066/preview.jpg?1241572934",  
    "updateURL":    "https://versioncheck.addons.mozilla.org/en-US/themes/update-  
↪check/18066",  
    "accentcolor":  "#000000",  
    "header":       "https://addons.mozilla.org/_files/18066/1232849758499.jpg?  
↪1241572934",  
    "version":      "1.0",  
    "footerURL":    "https://addons.mozilla.org/_files/18066/1232849758500.jpg?  
↪1241572934",
```

(continues on next page)

(continued from previous page)

```
"detailURL": "https://addons.mozilla.org/en-US/firefox/addon/dark-fox-18066/",
"textcolor": "#ffffff",
"id": "18066",
"description": "My dark version of the Firefox logo."
};

Mozilla.UITour.previewTheme(theme);
```

resetTheme()

Removes the previewed theme and resets back to default:

```
Mozilla.UITour.resetTheme();
```

cycleThemes(themes, delay, callback)

Cycles through an array of themes at a set interval with a callback on each step:

```
var themes = [
  ...
];

var myCallback = function () {
  ...
};

Mozilla.UITour.cycleThemes(themes, 5000, myCallback);
```

- themes (array)
- delay in milliseconds (number)
- callback to execute at each step (function)

getConfiguration(type, callback)

Queries the current browser configuration so the web page can make informed decisions on available highlight targets.

Available type values:

- 'sync'
- 'availableTargets'
- 'appinfo'
- 'selectedSearchEngine'
- 'search'
- 'canReset'

Other parameters:

- callback function to execute and return with the queried data

Specific use cases:

sync

If 'sync' is queried the object returned can be used to determine if the user has Sync enabled, and also metrics on the number and types of devices used.

```
Mozilla.UITour.getConfiguration('sync', function (config) {
  console.log(config) // { setup: true, desktopDevices: 2, mobileDevices: 1,
↳totalDevices: 3 }
});
```

Important: Sync device count metrics only available in Firefox 50 onwards.

availableTargets

If 'availableTargets' is queried the object returned by the callback contain array called targets. This can be used to determine what highlight targets are currently available in the browser chrome:

```
Mozilla.UITour.getConfiguration('availableTargets', function (config) {
  console.dir(config.targets);
});
```

appinfo

If 'appinfo' is queried the object returned gives information on the users current Firefox version.

```
Mozilla.UITour.getConfiguration('appinfo', function (config) {
  console.dir(config); //{defaultUpdateChannel: "aurora", version: "48.0a2",
↳distribution: "default", defaultBrowser: true}
});
```

The defaultUpdateChannel key has many possible values, the most important being:

- 'release'
- 'beta'
- 'aurora'
- 'nightly'
- 'default' (self-build or automated testing builds)

The distribution key holds the value for the Firefox distributionId property. This value will be default in most cases but can differ for repack or funnelcake builds.

The profileCreatedWeeksAgo key returns the number of weeks since the profile was created, starting from 0 for profiles dating less than seven days old.

The profileResetWeeksAgo key returns the number of weeks since the profile was last reset, starting from 0 for profiles reset less than seven days ago. If the profile has never been reset it returns null.

Important: appinfo is only available in Firefox 35 onward. The defaultBrowser property will only be returned on Firefox 40 or later. The distribution property will only be returned on Firefox 48 or later. Properties profileCreatedWeeksAgo and profileResetWeeksAgo will only be returned on Firefox 56 or later.

selectedSearchEngine

If 'selectedSearchEngine' is queried the object returned gives the currently selected default search provider.

```
Mozilla.UITour.getConfiguration('selectedSearchEngine', function (data) {
  console.log(data.searchEngineIdentifier); // 'google'
});
```

Important: selectedSearchEngine is only available in Firefox 34 onward.

search

This is an alias to 'selectedSearchEngine' that also returns an array of available search engines.

```
Mozilla.UITour.getConfiguration('search', function (data) {
  console.log(data); // { searchEngineIdentifier: "google", engines: Array[8] }
});

.. Important::

   ``search`` is only available in Firefox 43 onward.
```

canReset

If 'canReset' is queried the callback returns a boolean value to indicate if a user can refresh their Firefox profile via resetFirefox()

```
Mozilla.UITour.getConfiguration('canReset', function (canReset) {
  console.log(canReset); // true
});
```

Important: canReset is only available in Firefox 48 onward.

setConfiguration(name, value);

Sets a specific browser preference using a given key value pair.

Available key names:

- 'defaultBrowser'

Specific use cases:

defaultBrowser

Passing defaultBrowser will set Firefox as the default web browser.

```
Mozilla.UITour.setConfiguration('defaultBrowser');
```

Important: setConfiguration('defaultBrowser') is only available in Firefox 40 onward.

showFirefoxAccounts(extraURLCampaignParams);

Allows a web page to navigate directly to <https://accounts.firefox.com/signup?entrypoint=uitour>. In Firefox 47 and beyond, optionally accepts an object of utm_* key/values, which will be encoded and appended to the <https://accounts.firefox.com/signup> querystring.

Important: All keys in `extraURLCampaignParams` must begin with `utm_`. If an invalid key is present, the call to `showFirefoxAccounts` will fail.

```
// no extra utm_ campaign params. will open
// https://accounts.firefox.com/signup?entrypoint=uitour
Mozilla.UITour.showFirefoxAccounts();

// with extra utm_ campaign params. will open
// https://accounts.firefox.com/signup?&entrypoint=uitour&utm_foo=bar&utm_bar=baz
Mozilla.UITour.showFirefoxAccounts({
  'utm_foo': 'bar',
  'utm_bar': 'baz'
});
```

Important: `showFirefoxAccounts()` is only available in Firefox 31 onward. `extraURLCampaignParams` parameter only functional in Firefox 47 onward.

Note: A convenience method named `utmParamsFxA` exists in `js/base/search-params.js` that pulls all `utm_` params from the current page's URL and places them in an object (along with pre-defined defaults) ready to pass to `showFirefoxAccounts`.

resetFirefox();

Opens the Firefox reset panel, allowing users to choose to remove add-ons and customizations, as well as restore browser defaults.

```
Mozilla.UITour.resetFirefox();
```

Important: `resetFirefox()` should be called only in Firefox 48 onwards, and only after first calling `getConfig('canReset')` to determine if the user profile is eligible.

addNavBarWidget(target, callback);

Adds an icon to the users toolbar

- `target` can be an highlight target e.g. `forget` (string)
- `callback` to execute once icon added successfully (function)

```
Mozilla.UITour.addNavBarWidget('forget', function (config) {
  console.log('forget button added to toolbar');
});
```

Important: Only available in Firefox 33.1 onward.

setDefaultSearchEngine(id);

Sets the browser default search engine provider.

- `id` string identifier e.g. 'yahoo' or 'google'.

```
Mozilla.UITour.setDefaultSearchEngine('yahoo');
```

- Identifiers for en-US builds: <https://mxr.mozilla.org/mozilla-release/source/browser/locales/en-US/searchplugins/list.txt>
- Identifiers for other locales: <https://mxr.mozilla.org/10n-mozilla-release/find?string=browser%2Fsearchplugins%2Flist.txt>

Important: Only available in Firefox 34 onward.

setSearchTerm(string);

Populates the search UI with a given search term.

- `string` search term e.g. 'Firefox'

```
Mozilla.UITour.setSearchTerm('Firefox');
```

Important: Only available in Firefox 34 onward.

openSearchPanel(callback);

Opens the search UI drop down panel.

- `callback` function to execute once the search panel has opened

```
Mozilla.UITour.openSearchPanel(function() {  
    console.log('search panel opened');  
});
```

Important: Only available in Firefox 34 onward.

setTreatmentTag(name, value);

Sets a key value pair as a treatment tag for recording in FHR.

- `name` tag name for the treatment
- `value` tag value for the treatment

```
Mozilla.UITour.setTreatmentTag('srch-chg-action', 'Switch');
```

Important: Only available in Firefox 34 onward.

getTreatmentTag(name, callback);

Retrieved the value for a set **FHR** treatment tag.

- name tag name to be retrieved
- callback function to execute once the data has been retrieved

```
Mozilla.UITour.getTreatmentTag('srch-chg-action', function(value) {
  console.log(value);
});
```

Important: Only available in Firefox 34 onward.

ping(callback);

Pings Firefox to register that the page is using UiTour API.

- callback function to execute when Firefox has acknowledged the ping.

```
Mozilla.UITour.ping(function() {
  console.log('UiTour is working!');
});
```

Important: Only available in Firefox 35 onward.

openPreferences(id);

Opens the Firefox Preferences tab at a specified section. Accepts one of the following options to be passed as an *id*:

- 'general'
- 'search'
- 'content'
- 'applications'
- 'privacy'
- 'security'
- 'sync'
- 'advanced'

```
Mozilla.UITour.openPreferences('privacy');
```

Important: Only available in Firefox 42 onward.

closeTab();

Closes the current tab.

```
Mozilla.UITour.closeTab();
```

Important: This function will do nothing when called from the last browser window when it contains only one tab. You may need to provide a work around for this edge case in your code. This function is also only available in Firefox 46 onward.

showNewTab();

Opens `about:newtab` in the same tab.

```
Mozilla.UITour.showNewTab();
```

Important: This function is only available in Firefox 51 onward.

1.18 Send to Device widget

The *Send to Device* widget is a single macro form which facilitates the sending of a download link for either Firefox for iOS, Firefox for Android, or both. The form allows sending via SMS or Email, although the SMS copy & messaging is shown only to those in the configured countries. Geo-location is handled in JavaScript using a `bedrock view` that gets the request country from `CloudFlare CDN headers`. For users without JavaScript, the widget falls back to a standard Email form.

Important: This widget should only be shown to a limited set of locales who are set up to receive the emails. For those locales not in the list, direct links to the respective app stores should be shown instead. If a user is on iOS or Android, CTA buttons should also link directly to respective app stores instead of showing the widget. This logic should be handled on a page-by-page basis to cover individual needs.

Note: A full list of supported locales can be found in `settings/base.py` under `SEND_TO_DEVICE_LOCALES`, which can be used in the template logic for each page to show the form. The countries enabled for SMS for each message are defined in the `running configuration` per the environment names in `SEND_TO_DEVICE_MESSAGE_SETS` in the settings file.

1.18.1 Usage

1. Include this macro:

```
{% from "macros.html" import send_to_device with context %}
```

2. Add the appropriate lang file to the page template:

```
{% add_lang_files "firefox/sendto" %}
```

3. Make sure necessary files are in your CSS/JS bundles:

- 'css/base/send-to-device.less'
- 'js/base/send-to-device.js'

4. Include the macro in your page template:

```
{{ send_to_device() }}
```

The macro defaults to sending links for both Android and iOS apps. You can also pass a 'platform' to specify different configuration options:

```
{{ send_to_device(platform='select') }}
```

- `select` shows a drop down so the user can choose their platform.
- `ios` sends the user a link to Firefox for iOS only.
- `android` sends the user a link to Firefox for Android only.

If the page requires a custom title for the widget, you can also pass an optional heading:

```
{{ send_to_device(title_text='Foo Bar') }}
```

If you do not want to show a title, you can pass `include_title=False`:

```
{{ send_to_device(include_title=False) }}
```

To add a logo and rounded corners to the widget for display in a modal:

```
{{ send_to_device(include_logo=True) }}
```

If you need a customized App Store URL (e.g. including page-specific parameters), you can pass `ios_link`:

```
{{ send_to_device(ios_link=firefox_ios_url('mozorg-ios_page-appstore-button_↪sd') ) }}
```

5. Initialize the widget:

In your page JS, initialize the widget using:

```
var form = new Mozilla.SendToDevice();  
form.init();
```

1.18.2 Geolocation Callback

You can piggy-back on the widget's geolocation call by providing a callback function to be executed when the lookup has completed:

```
var form = new Mozilla.SendToDevice();  
form.geoCallback = function(countryCode) {  
  console.log(countryCode);  
}  
form.init();
```

The callback function will be passed a single argument - the country code returned from the geolocation lookup.

If the geolocation lookup fails, the country code passed to the callback function will be an empty string.

1.18.3 Micro embedded form

A micro embedded version of the send to device form is also available when targeting a single platform (e.g. `platform='android'` or `platform='ios'`).

The styles can be applied by using the following LESS file (instead of the regular stylesheet):

- `'css/base/send-to-device-micro.less'`

1.19 Firefox Accounts Signup Form

1.19.1 Introduction

Certain bedrock pages may feature a Firefox Accounts signup form. As this form has conditional functionality based on distribution (e.g. the China re-pack), the form should only be displayed to users of Firefox 48 and up.

Note: Firefox 48 is the minimum version needed to support the `distribution` property from UITour (which is how we detect the China re-pack).

Testing the signup flow on a non-production environment requires additional steps.

1.19.2 Configuring bedrock

Set the following in your local `.env` file:

```
FXA_ENDPOINT=https://latest.dev.lcip.org/
```

1.19.3 Configuring a demo Server

Demo servers must have the same `.env` setting as above. See the *Configure Demo Servers* docs.

1.19.4 Local and Demo Server Testing

Follow the [instructions](#) provided by the FxA team. These instructions will launch a new Firefox instance with the necessary config already set. In the new instance of Firefox:

1. Navigate to the page containing the Firefox Accounts form
2. If testing locally, be sure to use `127.0.0.1` instead of `localhost`

1.20 Funnel cakes and Partner Builds

1.20.1 Funnel cakes

In addition to being an [American delicacy](#) funnel cakes are what we call special builds of Firefox. They can come with extensions preinstalled and/or a custom first-run experience.

“The whole funnelcake system is so marred by history at this point I don’t know if anyone fully understands what it’s supposed to do in all situations” - pmac

Funnelcakes are configured by the Release Engineering team. You can see the configs in the [funnelcake git repo](#)

Currently bedrock only supports funnelcakes for “stub installer platforms”. Which means they are windows only. However, funnelcakes can be made for all platforms so [bedrock support may expand](#).

We signal to bedrock that we want a funnelcake when linking to the download page by appending the query variable *f* with a value equal to the funnelcake number being requested.

```
https://www.mozilla.org/en-US/firefox/download/thanks/?f=137
```

Bedrock checks to see if the funnelcake is configured (this is handled in the [www-config repo](#))

```
FUNNELCAKE_135_LOCALES=en-US
FUNNELCAKE_135_PLATFORMS=win,win64
```

Bedrock then converts that into a request to download a file like so:

Windows:

```
https://download.mozilla.org/?product=firefox-stub-f137&os=win&lang=en-US
```

Mac (You can see the mac one does not pass the funnelcake number along.):

```
https://download.mozilla.org/?product=firefox-latest-ssl&os=osx&lang=en-US
```

Someone in Release Engineering needs to set up the redirects on their side to take the request from here.

Places things can go wrong

As with many technical things, the biggest potential problems are with people:

- Does it have executive approval?
- Did legal sign off?
- Has it had a security review?

On the technical side:

- Is the switch enabled?
- Is the variable being passed?

1.20.2 Partner builds

Bedrock does not have an automated way of handling these, so you’ll have to craft your own download button:

```
<a href="https://download.mozilla.org/?product=firefox-election-edition&os=win&
↳lang=en-US">
Download</a>
```

Bugs that might have useful info:

- https://bugzilla.mozilla.org/show_bug.cgi?id=1450463
- https://bugzilla.mozilla.org/show_bug.cgi?id=1495050

PRs that might have useful code:

- <https://github.com/mozilla/bedrock/pull/5555>

1.21 A/B Testing

A/B testing on bedrock is done with the help of [Traffic Cop](#). A small javascript library which will direct site traffic to different variants in a/b experiments and make sure a visitor always sees the same variation.

It's possible to test more than 2 variants.

Traffic Cop sends users to experiments and then we use Google Analytics (GA) to analyze which variation is more successful. (If the user has DNT enabled they do not participate in experiments.)

All a/b tests should have a [mana page](#) detailing the experiment and recording the results.

1.21.1 Coding the variants

Traffic cop supports two methods of a/b testing. Executing different on page javascript or redirecting to the same URL with a query string appended. We mostly use the redirect method in bedrock. This makes testing easier.

Create a [variation view](#) for the a/b test.

The view can handle the url redirect in one of two ways:

1. the same page, with some different content based on the *variation* variable
2. a totally different page

Content variation

Useful for small focused tests.

This is explained on the [variation view](#) page.

New page

Useful for large page changes where content and assets are dramatically different.

Create the variant page like you would a new page. Make sure it is `noindex` and does not have a `canonical` url.

```
{% block canonical_urls %}<meta name="robots" content="noindex, follow">{% endblock %}
```

Configure as explained on the [variation view](#) page.

1.21.2 Traffic Cop

Create a .js file where you initialize Traffic Cop and include that in the experiments block in the template that will be doing the redirection. Wrap the extra js include in a `switch`.

```
{% block experiments %}
  {% if switch('experiment-berlin-video', ['de']) %}
    {{ js_bundle('firefox_new_berlin_experiment') }}
  {% endif %}
{% endblock %}
```

Switches

See the traffic cop section of the [switch docs](#) for instructions.

1.21.3 Recording the data

Including the `data-ex-variant` and `data-ex-name` in the analytics reporting will add the test to an auto generated report in GA. The variable values may be provided by the analytics team.

```
if(href.indexOf('v=a') !== -1) {
  window.dataLayer.push({
    'data-ex-variant': 'de-page',
    'data-ex-name': 'Berlin-Campaign-Landing-Page'
  });
} else if (href.indexOf('v=b') !== -1) {
  window.dataLayer.push({
    'data-ex-variant': 'campaign-page',
    'data-ex-name': 'Berlin-Campaign-Landing-Page'
  });
}
```

Make sure any buttons and interaction which are being compared as part of the test and will report into GA.

1.21.4 Tests

Write some tests for your a/b test. This could be simple or complex depending on the experiment.

Some things to consider checking:

- Requests for the default (non variant) page call the correct template.
- Requests for a variant page call the correct template.
- Locales excluded from the test call the correct (default) template.

1.21.5 A/B Test PRs that might have useful code to reuse

- <https://github.com/mozilla/bedrock/pull/5736/files>
- <https://github.com/mozilla/bedrock/pull/4645/files>
- <https://github.com/mozilla/bedrock/pull/5925/files>
- <https://github.com/mozilla/bedrock/pull/5443/files>

- <https://github.com/mozilla/bedrock/pull/5492/files>
- <https://github.com/mozilla/bedrock/pull/5499/files>

1.22 Analytics

1.22.1 Google Tag Manager (GTM)

Bedrock uses Google Tag Manager (GTM) to manage and organize its Google Analytics solution.

GTM is a tag management system that allows for easy implementation of Google Analytics tags and other 3rd party marketing tags in a nice GUI experience. Tags can be added, updated, or removed directly from the GUI. GTM allows for a *one source of truth* approach to managing an analytics solution in that all analytics tracking can be inside GTM.

Bedrock's GTM solution is CSP compliant and does not allow for the injection of custom HTML or JavaScript but all tags use built in templates to minimize any chance of introducing a bug into Bedrock.

1.22.2 The GTM DataLayer

How an application communicates with GTM is via the `dataLayer` object, which is a simple JavaScript array GTM instantiates on the page. Bedrock will send messages to the `dataLayer` object by means of pushing an object literal onto the `dataLayer`. GTM creates an abstract data model from these pushed objects that consists of the most recent value for all keys that have been pushed to the `dataLayer`.

The only reserved key in an object pushed to the `dataLayer` is `event` which will cause GTM to evaluate the firing conditions for all tag triggers.

1.22.3 DataLayer Push Example

If we wanted to track clicks on a carousel and capture what the image was that was clicked, we might write a `dataLayer` push like this:

```
dataLayer.push({
  'event': 'carousel-click',
  'image': 'house'
});
```

In the `dataLayer` push there is an `event` value to have GTM evaluate the firing conditions for tag triggers, making it possible to fire a tag off the `dataLayer` push. The `event` value is descriptive to the user action so it's clear to someone coming in later what the `dataLayer` push signifies. There is also an `image` property to capture the image that is clicked, in this example it's the house picture.

In GTM, a tag could be setup to fire when the event `carousel-click` is pushed to the `dataLayer` and could consume the `image` value to pass on what image was clicked.

1.22.4 The Core DataLayer Object

For the passing of contextual data on the user and page to GTM, we've created what we call the Core `DataLayer` Object. This object passes as soon as all required API calls for contextual data have completed. Unless there is a significant delay to when data will be available, please pass all contextual or meta data on the user or page here that you want to make available to GTM.

1.22.5 GTM Listeners & Data Attributes

GTM also uses click and form submit listeners to gather context on what is happening on the page. Listeners push to the dataLayer data on the specific element that triggered the event, along with the element object itself.

Since GTM listeners pass the interacted element object to the dataLayer, the use of data attributes works very well when trying to identify key elements that you want to be tracked and for storing data on that element to be passed into Google Analytics. We use data attributes to track clicks on all downloads, buttons elements, and nav, footer, and CTA/button link elements.

Important: When adding any new elements to a Bedrock page, please follow the below guidelines to ensure accurate analytics tracking.

For all nav, footer, and CTA/button link elements, add these data attributes:

Data Attribute	Expected Value
data-link-type	'nav', 'footer', or 'button'
data-link-name	name or text of the link

For all button elements, add this data attribute:

Data Attribute	Expected Value
data-button-name	name or text of the link

For all download buttons, add these data attributes:

Data Attribute	Expected Value
data-link-type	'Desktop', 'Android', or 'iOS'
data-download-os	name or text of the link
data-download-version	'standard', 'developer', 'beta'

For all conditional banners, add the following calls.

When a banner is shown:

```
dataLayer.push({
  'eLabel': 'Banner Impression',
  'data-banner-name': '<banner name>', //ex. Fb-Video-Compat
  'data-banner-impession': '1',
  'event': 'non-interaction'
});
```

When an element in the banner is clicked:

```
dataLayer.push({
  'eLabel': 'Banner Clickthrough',
  'data-banner-name': '<banner name>', //ex. Fb-Video-Compat
  'data-banner-click': '1',
  'event': 'in-page-interaction'
});
```

When a banner is dismissed:

```
dataLayer.push({
  'eLabel': 'Banner Dismissal',
  'data-banner-name': '<banner name>', //ex. Fb-Video-Compat
  'data-banner-dismissal': '1',
  'event': 'in-page-interaction'
});
```

When doing a/b tests configure something like the following.

```
if(href.indexOf('v=a') !== -1) {
  window.dataLayer.push({
    'data-ex-variant': 'de-page',
    'data-ex-name': 'Berlin-Campaign-Landing-Page'
  });
} else if (href.indexOf('v=b') !== -1) {
  window.dataLayer.push({
    'data-ex-variant': 'campaign-page',
    'data-ex-name': 'Berlin-Campaign-Landing-Page'
  });
}
```

1.22.6 Some notes on how this looks in GA

`data-link-type="button"` and `data-link-name=""` trigger a generic link click with the following structure:

- Event Category: {{page ID}} Interactions
- Event Action: {{data-link-type}} click
- Event Label: {{data-link-name}}

Any element that has a `data-button-name=""` triggers an event with this structure:

- Event Category: {{page ID}} Interactions
- Event Action: button click
- Event Label: {{data-button-name}}

1.23 Stub Attribution

Stub Attribution is a process that enables the construction and transmission of marketing attribution code on www.mozilla.org. When a user visits a mozilla.org download page, the website constructs an attribution code, which is then passed to download.mozilla.org via a URL parameter. This attribution code is then passed to the Windows stub installer, where it finally gets passed to Firefox for use in Telemetry. This attribution funnel enables Mozilla to better understand how different marketing campaigns effect overall retention in Firefox.

Note: The [AMO](#) team also relies on Stub Attribution on the `/new` page to inform their [Return to AMO](#) feature.

1.23.1 How does it work in bedrock?

The base template in bedrock contains a stub attribution script that runs on every page of the website. The script only runs on Windows (since the stub installer is a Windows feature) and if cookies are enabled. The script also respects Do Not Track, so it will not run if the user agent has this flag set.

The stub attribution script looks for *utm_** params on the page URL as well as the referrer, and then creates a hash of that data by passing it to an authentication service in bedrock. This hash is then accompanied by a signed, encrypted signature to prove that the data came from a trusted source. Both pieces of authenticated data are then stored in a cookie.

Once the user reaches the download page, bedrock then checks if the cookie exists, and if so appends the authenticated data to the download URL. The rest of the process is handled by the download service and stub installer.

1.23.2 Local testing

For stub attribution to work locally or on a demo instance, a value for the HMAC key that is used to sign the attribution code must be set via an environment variable e.g.

```
STUB_ATTRIBUTION_HMAC_KEY='thedude'
```

Note: This value can be anything if all you need to do is test the bedrock functionality. It only needs to match the value used to verify data passed to the stub installer for full end-to-end testing via Telemetry.

1.23.3 Measuring campaigns and experiments

Stub Attribution was originally designed for measuring the effectiveness of marketing campaigns where the top of the funnel was outside the remit of *www.mozilla.org*. For these types of campaigns, stub attribution requires zero configuration. It just works (as configured in */media/js/base/stub-attribution.js*) in the background and passes along any attribution data that exists.

For campaigns or experiments that originate on *www.mozilla.org*, Stub Attribution requires some extra configuration. We want to avoid using first-party utm parameters on *mozilla.org* pages, since those create new sessions in Google Analytics and can confound data. To work around this, there is a custom stub attribution script that can be used to pass custom data to the authentication service, avoiding the need for any first-party utm parameters.

To use the custom attribution script, override the regular stub attribution block in the page template where an experiment exists:

```
{% block stub_attribution %}
  {% if settings.STUB_ATTRIBUTION_RATE %}
    <!--[if !lte IE 8]><!-->
    {{ js_bundle('stub-attribution-custom') }}
    {{ js_bundle('my-page-experiment-script') }}
    <!--<![endif]-->
  {% endif %}
{% endblock %}
```

The *stub-attribution-custom* bundle contains the code that enables passing custom data to the authentication service. You can then use that code in *my-page-experiment-script*, or whatever you decide to call the experiment bundle:

```
var params = {
  utm_source: 'www.mozilla.org',
  utm_medium: 'download_button',
  utm_campaign: 'download_thanks_page',
  utm_content: 'download_thanks_install_experiment-v1'
};

// an optional callback for tracking when the experiment is seen in GA.
```

(continues on next page)

(continued from previous page)

```
function callback() {
  window.dataLayer.push({
    'data-ex-name': 'download_thanks_install_experiment',
    'data-ex-variant': 'v1'
  });
}

Mozilla.CustomStubAttribution.init(params, callback);
```

Authenticating custom data this way will override any existing stub attribution cookie a visitor may have.

Note: There are some exceptions in the way that the custom attribution script behaves. Existing utm data will not be overwritten if there is already a *utm_content* parameter in the page URL, or if the *utm_source* is from *addons.mozilla.org*. The custom script will follow the regular attribution flow in this scenario and will not modify anything. If any other utm parameters already exist on the page URL, those will be passed through to the custom attribution script and will take precedence over any custom values. This is to try and preserve as much existing information as possible, whilst still retaining the *utm_content* value that is essential to attributing an experiment.

1.24 Architectural Decision Records

We record major architectural decisions for bedrock in Architecture Decision Records (ADR), as described by [Michael Nygard](#). Below is the list of our current ADRs.

1.24.1 1. Record architecture decisions

Date: 2019-01-07

Status

Accepted

Context

We need to record the architectural decisions made on this project.

Decision

We will use Architecture Decision Records, as described by [Michael Nygard](#).

Consequences

See [Michael Nygard's article](#), linked above. For a lightweight ADR toolset, see [Nat Pryce's adr-tools](#).