

Update Translations Infrastructure

29 November 2022

OVERVIEW

The current translations infrastructure has become outdated and difficult to maintain. The i18n SIG has identified open-source Weblate as the best new platform to use, and we propose to use a hosted instance of Weblate at weblate.org, which is offering OpenStack a 50% discount. The OpenStack Technical Committee asks the Foundation to support this effort by contracting with weblate.org to pay for this service. Additionally, the TC asks the Foundation Staff to encourage member companies to volunteer a person with an engineering background to drive the effort to integrate Weblate into the current OpenStack translations infrastructure and to provide ongoing maintenance.

GOALS

1. Provide translators with a familiar, easy-to-use platform hosted by a service provider experienced in this area.
2. Provide engineering support to handle the plumbing that connects OpenStack projects to the localization infrastructure.
3. Allow the translations team to focus on actual translation by removing infrastructure maintenance as one of its duties.

DETAILS

The translations team currently uses a Zanata instance maintained by the OpenDev Infra team, but this is not a tenable situation given that the Zanata code has not been maintained since 2018, leading the OpenDev Collaboratory sysadmins to implement imperfect workarounds for known security vulnerabilities in Zanata. Additionally, Zanata uses an increasingly outdated toolchain and platform, which must be handled as an exception to normal systems deployment and maintenance by OpenDev sysadmins. See [0] for additional problems the OpenDev maintained Zanata instance is causing the translations team.

The translations team has been planning to migrate to the open-source Weblate platform since April 2019 [1], but this effort has stalled at the infrastructure provisioning and configuration stage.

We propose to address this by moving these duties to a hosting service and to a volunteer engineer specifically tasked with maintaining the OpenStack-to-Weblate interface.

With respect to the dedicated engineer, there are two primary integration points. The first is user accounts, which would ideally use OpenStackID to give us seamless integration and make it easier to track contributions. The second is the code that currently uploads .pot files to Zanata and downloads .po files that are proposed to projects via Gerrit. The code will need to be rewritten to use Weblate APIs instead of Zenata and will ideally preserve the current workflow. We estimate that the initial effort will take an engineer three weeks to a month, and then there would be an ongoing three or so days per month for tooling maintenance, dependency upgrades, and bug fixes.

Weblate [2] is copylefted libre software developed by weblate.org, which also provides hosting services [3]. Weblate.org does provide free hosting for libre projects at a level similar to their “Advanced” plan, but OpenStack exceeds the 10,000 source-string limit on that plan.¹

We contacted Benjamin Alan Jamie, the Weblate Community Manager, and weblate.org is willing to offer OpenStack a 50% discount on their “Enterprise” offering (with discount, this is currently €1500/year). This would give OpenStack a dedicated Weblate instance, with weblate.org handling maintenance (updates, backups, etc.) and providing support to the dedicated engineer tasked with the technical aspects of getting our infrastructure connected to theirs.

Several other open projects have migrated from Zanata to Weblate, including The Document Foundation (LibreOffice) and The Fedora Project [4] (which uses a dedicated instance at weblate.org, like what we’re proposing here).

Our initial contact with weblate.org discussed a year-to-year contract, but it’s possible that the OpenInfra Foundation could negotiate something longer-term. We don’t have reservations about a longer-term agreement given that the translations team did an analysis of various options back in 2018 and selected Weblate, and in the meantime, Weblate has become a popular localization platform, both for projects and translators. Additionally, Seongsoo Cho, an active member of the OpenStack i18n SIG, has been testing Weblate using a self-hosted instance on his laptop, and anticipates that it will be an easy transition to using a hosted Weblate.²

¹ This is despite projects reducing the amount of translated content several years ago to lessen the burden on the translation team, the point being that we’re already close to the bone in terms of source strings for localization.

² We mention this to emphasize that although the i18n SIG conducted its initial determination of the best post-Zanata platform in 2018, the team has continued to experiment with Weblate to the present day and continue to view Weblate as the best target for a migration away from Zanata.

ALTERNATIVES

No Localization

With respect to OpenStack individual projects' translations, the number of languages for which localization files are provided has been dropping (see [5] for an example), and we have not heard complaints from operators. This may be because many operators are using older releases for which there are a larger number of translations available, or because the decline in translations isn't visible in Horizon, which currently is translated into 29 languages (whereas the Zed release of Cinder, for example, has localization files for only 3 languages). The 2023 User Survey includes questions posed by the i18n SIG to get a sense whether operators use translations, and if so, which languages, but of course, that data isn't available yet. In any case, we anticipate that there is demand for at least the major non-Western languages. Additionally, it would be a bad look for an open source project that strives to serve a worldwide user base not to offer localization.

Self-Hosted Weblate

The OpenDev Infra team can supply a server, but the translations-using community would need to supply specific volunteers to add the deployment to OpenDev. The Infra team has indicated that they would be happy to help, but they don't have bandwidth to drive the effort. The primary problem with this approach is that there is a skills mismatch between people interested in doing translations and people interested in configuring, maintaining, and deploying infrastructure. For example, Ian Choi's call for volunteers to help with Weblate migration on the openstack-discuss mailing list [0] only received one response [6]. It would be better to shift the infrastructure maintenance to weblate.org, and the OpenStack-Weblate integration to an engineer focused primarily on that task, and then the i18n SIG can focus on localization tooling and translations.

NEXT STEPS

The first milestone of the 2023.1 ("Antelope") OpenStack release passed on 17 November 2022, so it's important to move on this as soon as possible. While it's not clear that this late in the development cycle the translations team will be able to complete the Weblate migration, this proposal will remove the infrastructure setup and maintenance duties from the team and allow them to focus on identifying an engineer who can concentrate on rewriting/adapting the localization export/import automation.

Our contact person at weblate.org is Benjamin Alan Jamie, who is Weblate Community Manager and also the head of sales. Benjamin indicated that the best way to contact him is to make an appointment on his web calendar, <https://benjamin.weblate.org/>.

A copy of our email correspondence and a report of our initial meeting with Benjamin is attached as an appendix.

REFERENCES

[0] <https://lists.openstack.org/pipermail/openstack-discuss/2022-October/030713.html>

[1] <https://etherpad.opendev.org/p/l18n-weblate-migration>

[2] <https://github.com/WeblateOrg/weblate>

[3] <https://weblate.org/en/hosting/>

[4] https://fedoraproject.org/wiki/L10N_Move_to_Weblate

[5] <https://etherpad.opendev.org/p/translation-situation>

[6] <https://lists.openstack.org/pipermail/openstack-discuss/2022-October/030723.html>

[7] <https://releases.openstack.org/antelope/schedule.html#a-1>

APPENDIX

Correspondence with weblate.org

Initial Inquiry, 2022-11-01

Date: Tue, 1 Nov 2022 11:55:15 -0400
Cc: Ghanshyam Mann <gmann@ghanshyammann.com>
From: Brian Rosmaita <rosmaita.fossdev@gmail.com>
To: sales@weblate.org
Subject: inquiry into gratis hosting for OpenStack

Hello Weblate sales org,

Please forward this on to the Weblate Community Manager, Benjamin Alan Jamie.

OpenStack (<https://openstack.org>), a large open source project, has been looking for a translations solution to replace our current self-hosted Zenata instance (<https://translate.openstack.org/>) and our i18n team is very interested in Weblate. Additionally, our infrastructure team was very excited to learn that Weblate offers gratis hosting to libre projects.

I'm a member of the OpenStack Technical Committee and am making an initial inquiry about the feasibility of gratis hosting for OpenStack. (Ghanshyam Mann, the chair of the OpenStack Technical Committee, is cc'd on this email.)

OpenStack meets your basic requirements:

- Translated content has to be released under a license approved by OSI or recognized as libre by FSF.
 - * The OpenStack project is provided under the Apache 2.0 license, which is an OSI approved license.

- All the source code has to be publicly available in a supported version control system.
 - * Code is hosted in git at opendev.org/openstack and mirrored at [github](https://github.com)

- Translations have to be in the same repository as the source code, or under the same project/organization as the source code repository.
 - * OpenStack code repositories have the same structure; translations are in `project/project/locale/` , e.g., for the horizon project, translations are committed to opendev.org/openstack/horizon/horizon/locale

-
- Using Weblate has been approved by the upstream project.
 - * the OpenStack SIG i18n has discussed migrating to Weblate with the OpenStack Technical Committee; no objections have been raised
 - Mention you use Weblate for translations in the README. Itâ€™ll attract new translators and also help Weblate to be free for more projects.
 - * we can do this
 - The project has had active development for at least three months.
 - * OpenStack has been actively developed since 2010
 - You will have 14 days to set up your project.
 - * OK
 - Please request project approval once you finish the setup.
 - * OpenStack is a large project, so we thought we should ask for approval before doing any setup
 - Unapproved projects will be removed.
 - * OK
 - There is no guarantee of service availability, nor quality for Libre hosting.
 - * Understood

Your website states that the Libre plan has the same limits as the Advanced plan, but only for public projects. So that would be:

source strings: 10,000
target languages: 60
translation projects: unlimited
translation components: unlimited
number of translators: unlimited

As of today, OpenStack would have a total of 672 translation components spread over 48 translation projects. We have translations to 29 different languages, so the 60 target language limit is not an issue. We already have close to 10,000 source strings, however, so we would exceed that limit.

Thanks for reading all this background info. Here are some questions:

1. Would it be possible for you to gratis host OpenStack translations, but with a higher limit on the number of source strings?
2. If not, it looks like OpenStack would need the "Enterprise Cloud" hosting plan, which your website lists at â‚¬299 monthly. Would you be interested in offering a discounted price for a large open source

project?

3. Do you have any other suggestions for how you could help us out?

Thanking you in advance for your attention to this request!

On behalf of the OpenStack Technical Committee and SIG i18n,
Brian Rosmaita

weblate.org response, 2022-11-03

Date: Thu, 03 Nov 2022 12:16:00 +0100

From: Benjamin Alan Jamie via Weblate Care <care@weblate.org>

To: rosmaita.fossdev@gmail.com

Cc: gmann@ghanshyammann.com

Subject: RE: inquiry into gratis hosting for OpenStack [Ticket#206762]

Hello Brian,

Thank you for your message; we appreciate OpenStack's interest in Weblate. There shouldn't be any issue with hosting the translations. Would you mind connecting for a short coordination call? You can reserve the time at <https://benjamin.weblate.org/>.

Kind regards,
Benjamin

--

Weblate

--

report on meeting with Benjamin Jamie, 2022-11-08

Date: Tue, 8 Nov 2022 10:51:55 -0500

Subject: Re: [i18n] Weblate hosting update - followup

From: Brian Rosmaita <rosmaita.fossdev@gmail.com>

To: Seongsoo Cho <ppiyakk2@printf.kr>, Ian Y Choi <ianyrchoi@gmail.com>, Jeremy Stanley <fungi@yuggoth.org>, Clark Boylan <cboylan@sapwetik.org>

Cc: Ghanshyam Mann <gmann@ghanshyammann.com>

tl;dr - Weblate is willing to give us a 50% discount, but not free hosting

summary

I met with Benjamin Alan Jamie (Weblate community manager) earlier today. Weblate is very excited to work with OpenStack, but given our size, they can't offer free hosting, but are willing to give us a 50% discount, which will be 1500 euros/year.

There are two options (both cost the same):

1. Hosted plan on Hosted Weblate

- downside: we're all mixed in with other opensource hosted projects, so it can be difficult for our translators to find our projects.
- upside: more visibility; people on the platform may notice that we're missing Norwegian translations and contribute them

2. Dedicated instance for OpenStack

- downside: only "our" translators, so we won't get drive-by translations
- upside: we can fully customize the server, set up LDAP or SAML or whatever

Benjamin suggested the Dedicated option (#2) to start. Weblate is working on a feature to add something like a GitHub "organization" that will allow us to group all the OpenStack projects together, which will make it easier to find OpenStack if we move to Hosted (option #1) after the feature's been implemented (sounded like it's at least 6 months away). It will be easy to migrate from Dedicated to Hosted.

Fedora has migrated from Zenata to Weblate, and is using option #2:

<https://translate.fedoraproject.org/>

The person who organized the move to Weblate is Jean-Baptiste Holcroft <jean-baptiste@holcroft.fr>, <https://opensource.com/users/jibec> (may be a useful contact for us).

OpenSUSE hosts their own Weblate instance with a support contract (I forgot to ask how much that would cost if we wanted to go that route).

As far as more information goes: Benjamin is booked for the rest of this week, and then will be away until 21 November. I told him that some of

our infrastructure and translations people will want to talk to him directly; you can schedule time with him at:
<https://benjamin.weblate.org/>

The Weblate team will be at FOSDEM in February if anyone wants to meet in person.

next steps

We need to see if the Foundation will spring for €1500/year for Weblate.
I'm not sure how we do that.