




Does open source need its own Priority of Constituencies?





Tobie Langel (@tobie)
Principal, UnlockOpen

Why Open Source Failed



John Mark

Follow

Jul 30, 2018 · 10 min read ★

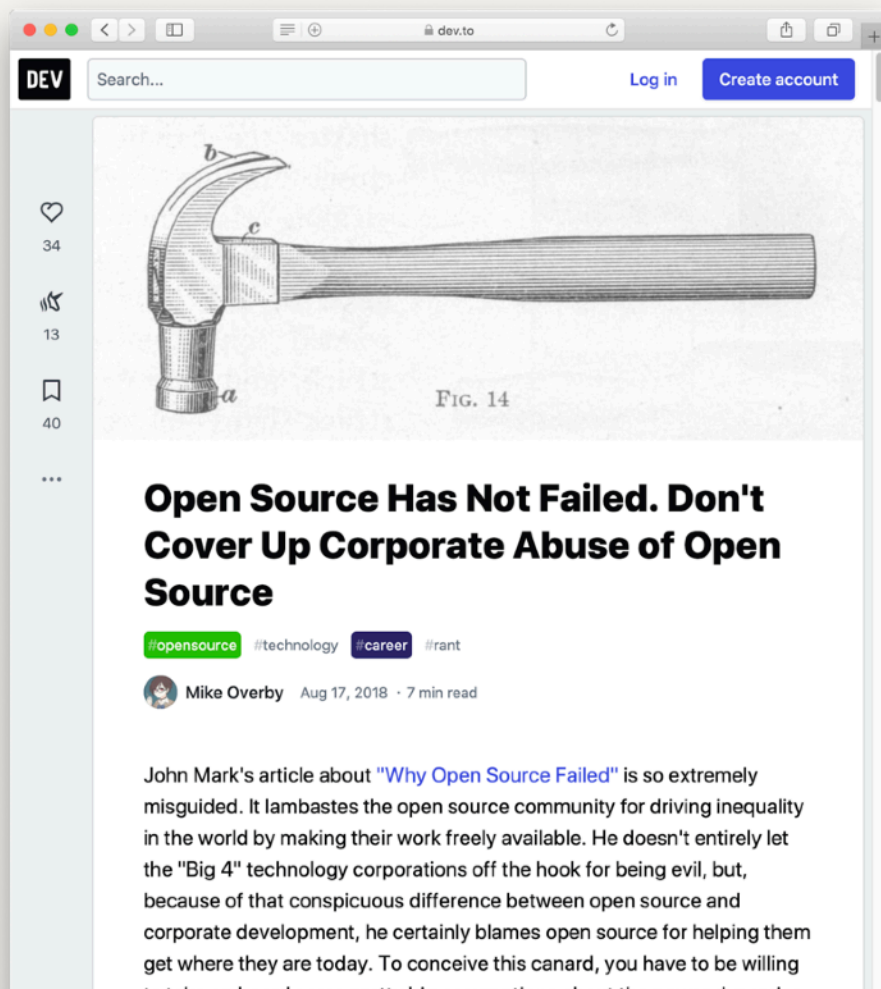


"Money" by Andrew Magill CC-BY

(If you like this essay, you may be interested in my follow-up, "[Save Open Source, Save the World](#)")

2018 is the 20th anniversary of the term "open source", and a good number

Tobie Langel (@tobie)
Principal, UnlockOpen



Featured Article

The crusade against open-source abuse



Cloud infrastructure providers threaten the viability of open source

Salil Deshpande @salil / 6:30 pm CET • November 29, 2018

💬 Comment



📷 Image Credits: Brad / Flickr under a CC BY 2.0 license.

aws

Contact SalesSupport▼My Account▼Create an AWS Account

ProductsSolutionsPricingDocumentationLearnPartner NetworkAWS Marketplace

>>Q


Blog HomeCategory ▼Edition ▼Follow ▼

Search BlogsQ

AWS Open Source Blog

Keeping Open Source Open – Open Distro for Elasticsearch

by Adrian Cockcroft | on 11 MAR 2019 | in Amazon Elasticsearch Service, Open Distro For Elasticsearch, Open Source | Permalink | Comments | Share



中文版 – At AWS, we focus on solving problems for customers. Over the years, customer usage and dependencies on open source technologies have been steadily increasing; this is why we've long been committed to open source, and our pace of contributions to open source projects – both our own and others' – continues to accelerate.

When AWS launches a service based on an open source project, we are making a long-term commitment to support our customers. We contribute bug fixes, security, scalability, performance, and feature enhancements back to the community. For example, we have been a significant contributor to Apache Lucene, which powers Amazon Elasticsearch Service. The Amazon EMR team has been making contributions to the Hadoop ecosystem for many years, and the Amazon Elastic Container Service for Kubernetes (EKS) team has been contributing to Kubernetes. We also invest in open source communities, [training developers and operators](#), and sponsor open source events and conferences such as [ApacheCon](#) and [KubeCon](#), and recently

Resources

Open Source at AWS
Projects on GitHub

Follow

☒ AWS Open Source

☒ AWS Cloud


☒ Facebook

☒ LinkedIn

☒ Twitch

☒ Open Source RSS

☒ Email Updates

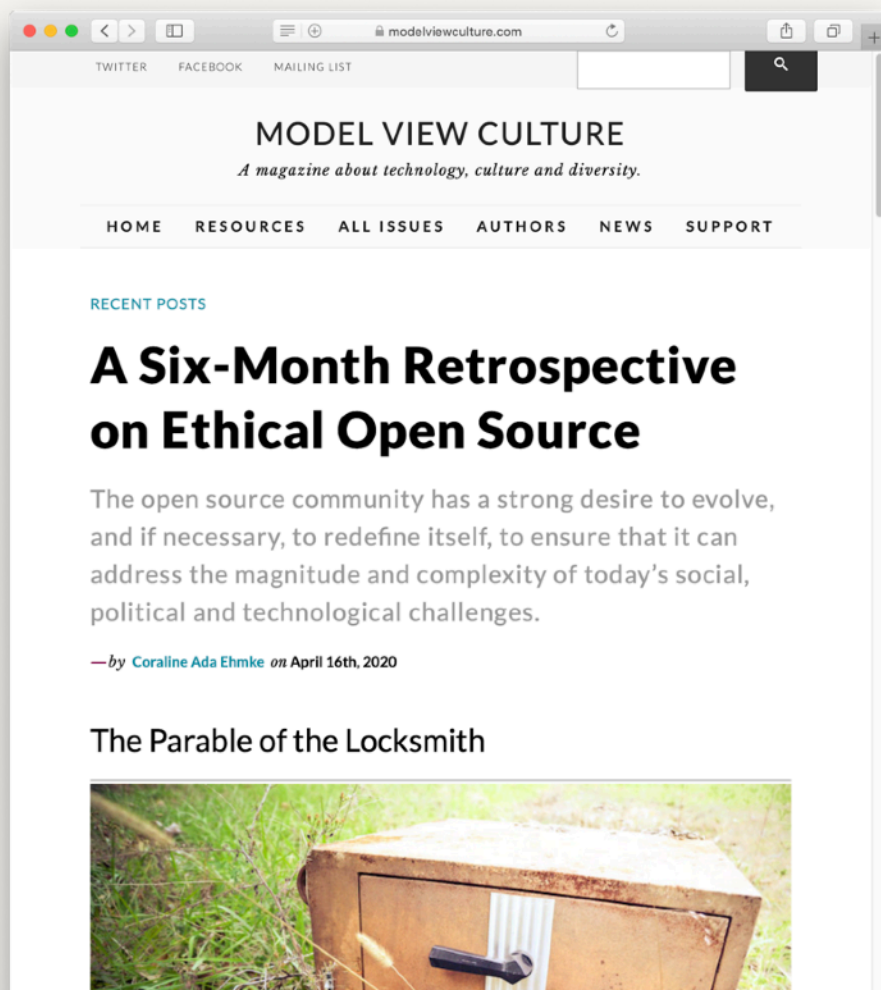


AWS Events

Discover the latest AWS events in your region

[Learn more >>](#)

Tobie Langel (@tobie)
Principal, UnlockOpen




youtube.com

Search

SIGN IN

What Comes After Open Source?

Bruce Perens
bruce@perens.com




0:10 / 59:44

Chat Replay is disabled for this Premiere.

What Comes After Open Source

658 views • Premiered Aug 24, 2020

29 2 SHARE SAVE ...

 **Bruce Perens**
34 subscribers

SUBSCRIBE

Free Software / Open Source have been going on for 37 years, and it's time to assess our performance and see if we can preserve our successes and improve upon our failures. I propose preserving Open Source as it exists and going forward with a new paradigm.

SHOW MORE

Tobie Langel (@tobie)
Principal, UnlockOpen

Open source crisis


- Maintainer burnt out.
- Cloud capturing too much value.
- Ethical concerns over impact on end-users.

What's going on?

We're witnessing the emergence of new constituencies.

The 4 freedoms & the OSD

The 4 freedoms & the OSD



How do we handle this increased diversity of constituencies?

W3C's Priority of Constituencies

[/irc-logs / w3c / #html-wg / 2007-03-27 / end](#)

Options: ☐ principles ☒ Hide Join/Parts/Quits

```
# [01:09] <gnedders> haivonen: "blank in only is assumed" -- black ink, surely?
# [06:04] * Disconnected
# [06:07] * Attempting to rejoin channel #html-wg
# [06:07] * Rejoined channel #html-wg
# [06:07] * Topic is 'W3C HTML WG http://www.w3.org/html/wg/ - http://krijinhoetmer.nl/irc-logs/ (logged)'
# [06:07] * Set by anne on Fri Mar 21 16:12:15
# [06:45] <achy> > it's ironic that the most vocal people on the list are those complaining about the amount of traffic :-)
# [06:52] <ixie> statistically it's most likely that the people sending the most comments are the ones going to complain :-)
# [07:04] <achy> > Hey Ixie, we really need to do something to improve image maps.
# [07:09] <achy> > when you get up to writing that section, of course
# [07:09] * achy_ just had to implement an image map by positioning <a> with CSS cause image maps are too limited
# [07:11] <mje> the problem is, what you really want instead of image maps is hit testing of arbitrary vector shapes
# [07:11] <mje> including possibly transparent ones
# [07:12] <mje> and once you start going there, you kind of end up with svg
# [07:15] <mje> or with something like d3.js's CSS-VG idea
# [07:16] <ixie> I thought I'd finished image maps
# [07:29] <anne> This Mike Schinkel complains about too much e-mail and then he sends another six e-mails to the list...
# [07:31] <haivonen> I don't like the wiki helping, but I try to shut up for now about it on the list
# [07:32] <mje> discussing real issues may reduce discussion of bullshit non-issues
# [07:32] <haivonen> I don't *use* the wiki helping...
# [07:32] <mje> the way <ixie> discussion has somewhat reduced background chatter on whatwg list
# [07:32] <mje> but it probably won't reduce total mail volume
# [07:32] <mje> s
# [10:33] <mjs> I really would like to note down some general design principles to use when evaluating new
# features
# [10:33] <mje> I think that ontologically precedes test cases and user stories
# [10:33] <mjs> I think there is rough consensus around a core set of principles but it is good to have
# something to point to
# [10:34] <mjs> especially for the few who disagree with core principles and therefore might not be satisfied
# with any result from the group
# [10:34] <haivonen> real vs. be pretty much blocks on MS lawyers. I wonder if Chris Wilson will join the telecon. If not, I'd expect the telecon to
# be similarly blocked
# [10:36] <mje> yes, most of the people with real issues are waiting for him
# [10:37] <anne> well, the reason this WG started is largely because of MS I think
# [10:37] <mje> that and to have a reasonable framework to handle patents
# [10:38] <mje> I would support continuing it even if MS flakes out
# [10:38] <mje> on the other hand, a lot of W3C members with large patent portfolios have also not joined yet
# [10:39] <haivonen> It would be great if the IBM AC rep allowed Sam Ruby to join.
# [10:39] <mje> Sam Ruby wants to join but doesn't have permission?
# [10:39] <mje> or is that a guess?
# [10:40] <haivonen> mje: http://www.intertwingly.net/blog/2007/03/27/open-invitation#1173393798
# [10:41] <haivonen> the way I read this is that he wanted to join, but the W3C did not let him join as an individual (I wasn't allowed, either) and
# the IBM AC rep has not given permission
# [10:42] <mje> it's unclear if he has asked the IBM AC rep
# [10:42] <anne> haivonen, you were not allowed to join as individual?
# [10:42] <mje> it is clear that he wasn't allowed to join as an individual
# [10:42] <haivonen> mje: well, yeah
# [10:42] <mje> IBM is officially pro-W3C
# [10:42] <haivonen> anne: I have a financial relationship with a Member, so yeah
# [10:42] <mje> but I don't know how strong the internal politics on that are
# [10:42] <anne> I think they're pretty strong XForms-focused and XML web like
# [10:42] <haivonen> anne: except the Atom guys
# [10:42] <anne> Atom is XML focused
# [10:42] <haivonen> anne: markp even left to Google over this
# [10:43] <anne> yup
# [10:44] * anne was going to say that
# [10:44] <mje> was markp formerly an IBM employee?
# [10:44] <haivonen> I wonder if Google will be paying him for Gecko WF 2.0 accessibility work that IBM didn't let him be involved with
# [10:44] <haivonen> mje: yes
# [10:45] <mje> I hope Google lets markp join the FTHWG
# [10:45] <haivonen> mje: see wiki history about him withdrawing from Gecko WF 2.0 accessibility
# [10:45] <mje> I did not realize the extent of IBM's partisanship in this matter
# [10:53] <marcos_> mjs: why don't we bash out the design principles now?
# [10:54] <mje> marcos_ I wouldn't mind, I could post more about then, but I feel bad undermining Dan's attempt to set agenda and steer the
# conversation
# [10:54] <mje> I think besides DFWN the other big one is Degrade Gracefully
# [10:55] <mje> which is basically the opposite: let content using the new stuff work reasonably in old browsers
# [10:55] <mje> others would be general feel-goodism without as many obvious consequences.
# [10:55] <marcos_> I don't think it is a problem for us to do that here. At least we can have a document to discuss.
# [10:55] <mje> Like Avoid Needless Complexity
# [10:56] <mje> then there is the hierarchy of cost, which is difficulty for end user matters more than difficulty for content author matters more
# than difficulty for browser maker matters more than difficulty for spec editor
# [10:57] <mje> (although these aren't always in conflict, as making it easy for the author to do the right thing often helps users, and making rules
# in the spec simple so that browsers can implement them readily helps authors through improving interoperability, etc)
# [10:58] <mje> this will all seem like obvious stuff to whatwg people
```

3.2. Priority of Constituencies

In case of conflict, consider users over authors over implementors over specifiers over theoretical purity. In other words costs or difficulties to the user should be given more weight than costs to authors; which in turn should be given more weight than costs to implementors; which should be given more weight than costs to authors of the spec itself, which should be given more weight than those proposing changes for theoretical reasons alone. Of course, it is preferred to make things better for multiple constituencies at once.

3.3. Secure By Design

Ensure that features work with the security model of the web. Preferably address security considerations directly in the specification.

Communicating between documents from different sites is useful, but an unrestricted version could put user data at risk. Cross-document messaging is designed to allow this without violating security constraints.

3.4. Separation of Concerns

HTML should allow separation of content and presentation. For this reason, markup that expresses structure is usually preferred to purely presentational markup. However, structural markup is a means to an end such as [media independence](#). Profound and detailed semantic encoding is not necessary if the end can be reached otherwise. Defining reasonable default presentation for different media may be sufficient. HTML strikes a balance between semantic expressiveness and practical usefulness. Names of elements and attributes in the markup may be pragmatic (for brevity, history, simplicity) rather than completely accurate.

The `article` element defines an individual article, but not the details of how it is displayed. A journal article may be the only article on a page, formatted in multiple columns, while a blog post may share a page with multiple other articles and be presented in a box with a border.

The `b` and `i` elements are widely used — it is better to give them good default rendering for various media including aural than to try to ban them.

3.5. DOM Consistency

The two serializations should be designed in such a way that the DOM trees produced by the respective parsers appear as consistently as feasible to scripts and other program code operating on the document trees. Discrepancies can be allowed for compatibility with legacy implementations, but the differences should be minimized.

Also, unless required for compatibility with legacy implementations and deployed content, gratuitous difference in syntactic appearance should be avoided as well.

The HTML (`text/html`) parser puts elements in the `http://www.w3.org/1999/xhtml` namespace in the DOM for compatibility with the XML syntax of HTML 5.

§ 1.1. Put user needs first (Priority of Constituencies)

If a trade-off needs to be made, always put user needs above all.

The internet is for end users: any change made to the web platform has the potential to affect vast numbers of people, and may have a profound impact on any person's life.

User needs come before the needs of web page authors, which come before than the needs of user agent implementors, which come before than the needs of specification writers, which come before theoretical purity.

Similarly, when beginning to design an API, be sure to understand and document the user need that the API aims to address.

Like all principles, this is not absolute. Ease of authoring affects how content reaches users. User agents have to prioritize finite engineering resources, which affects how features reach authors. Specification writers also have finite resources, and theoretical concerns reflect underlying needs of all of these groups.

See also:

- [The web should not cause harm to society](#)
- [The Internet is for End Users](#)

§ 1.2. It should be safe to visit a web page

Hyperlinks, links from one page to another, are one of the foundations of the Web.

Following a link, or visiting a web page, should be safe: users doing this should not have to fear for the security of their computer or for essential aspects of their privacy. (But it's not completely safe, in the sense that users may face consequences if their use of the Web is harming others.) Furthermore, users should understand that it is safe (and how it isn't) so they can make informed decisions between use of the Web versus other technologies.

Saying "essential aspects" here skips over quite a bit of detail. The Web today is far from being perfectly private. One aspect of privacy problems is when reality doesn't meet expectations. For example, a person walking down the street generally expects to be recognized by their friends, but (depending on the country) may not expect that they walked down that street at that time to be recorded in a permanent government database. Online, people might have less understanding of what is or isn't possible, and their

Abstract

Status of This Memo

This document is a product of the Internet Architecture Board (IAB) and represents information that the IAB has deemed valuable to provide for permanent record. It represents the consensus of the Internet Architecture Board (IAB). Documents approved for publication by the IAB are not candidates for any level of Internet Standard; see Section 2 of RFC 7841.

Copyright Notice

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

1. Introduction
2. Who Are "End Users"?
3. Why the IETF Should Prioritize End Users
4. How the IETF Can Prioritize End Users
 - 4.1. Engaging the Internet Community
 - 4.2. Creating User-Focused Systems
 - 4.3. Identifying Negative End-User Impact
 - 4.4. Handling Conflicting End-User Needs
 - 4.5. Deprioritizing Internal Needs
5. IANA Considerations
6. Security Considerations
7. Informative References

IAB Members at the Time of Approval

Acknowledgements

W3C priority of constituencies

End users > Authors > Implementors > Spec editors > Theoretical purity

Size of each constituency

End users > Authors > Implementors > Spec editors > Theoretical purity

Billions

Millions

Thousands

Tens

Cost of 1 hour of spec work

End users > Authors > Implementors > Spec editors > Theoretical purity

<i>1,000,000,000 h</i>	<i>1,000,000 h</i>	<i>100 h</i>	<i>1 h</i>	
------------------------	--------------------	--------------	------------	--

Apache Software Foundation mantra

End users > Authors > Implementors > Spec editors > Theoretical purity

Apache Software Foundation mantra

End users > Authors > Implementors > Spec editors > **Theoretical purity**

Apache Software Foundation mantra

End users > Theoretical purity

Apache Software Foundation mantra

People > Code

What would an open source
priority of constituencies look like?

Open source priority of constituencies

People

Contributors

End users

App devs

Maintainers

*(Individuals &
corporations)*

Cloud infra

Open source priority of constituencies

People > End users > App devs > Cloud infra > Contributors > Maintainers > Theor. pur.

But is W3C's priority of constituencies a silver bullet?


Economic situation

End users > Authors > Implementors > Spec editors > Theoretical purity




Economic situation

End users > Authors > Implementors > Spec editors > Theoretical purity





Economic situation

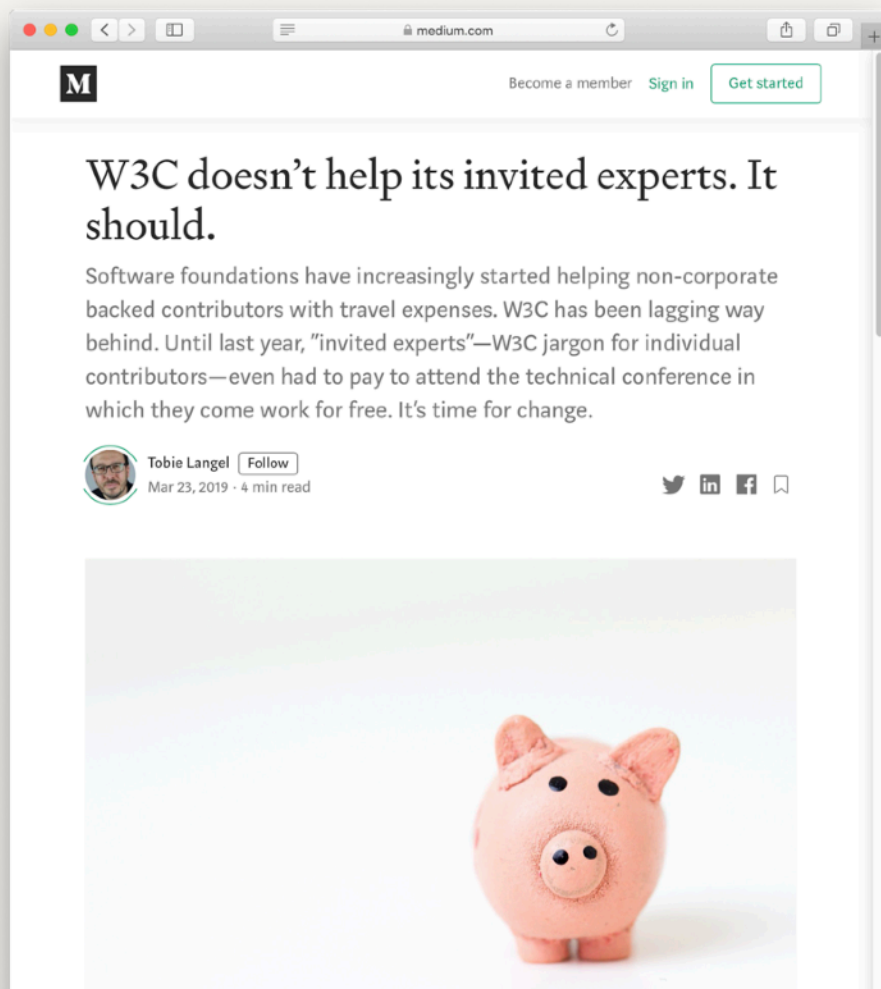
End users > Authors > **Implementors** > Spec editors > Theoretical purity



Economic situation

End users > Authors > Implementors > Spec editors > Theoretical purity





Let's apply this to open source.



The canary in the coal mine

Open source priority of constituencies


People > End users > App devs > Cloud infra > Contributors > Maintainers > Theor. pur.

Open source priority of constituencies


People > End users > App devs > Cloud infra > Contributors > Maintainers > Theor. pur.



—




—




Open source priority of constituencies

People > End users > App devs > Cloud infra > Contributors > Maintainers > Theor. pur.

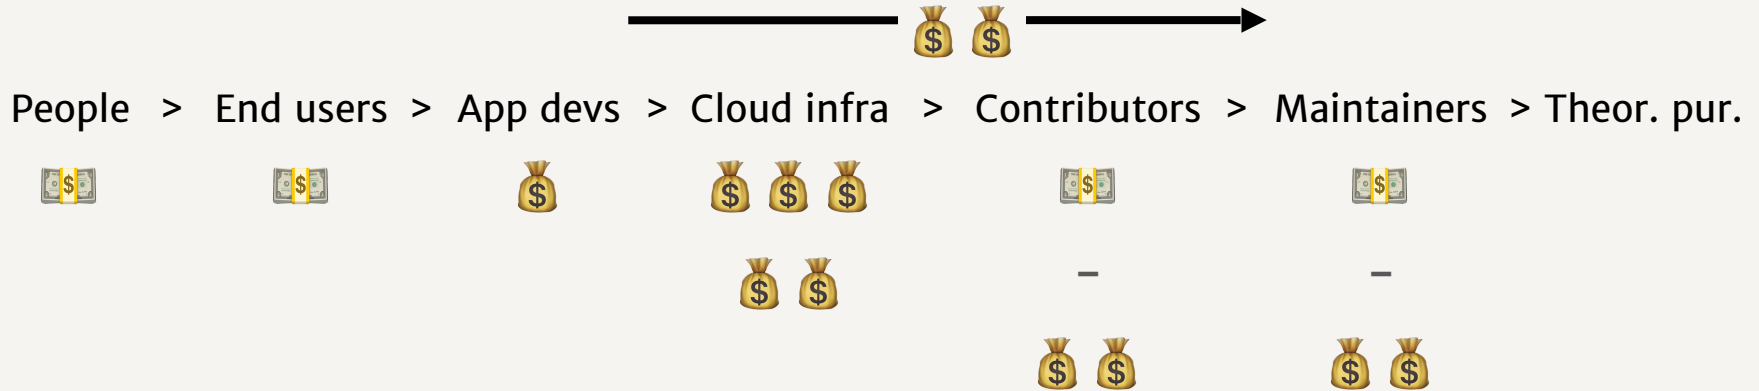


—

—



Open source priority of constituencies



Benefits of a priority of constituencies

- Keeps focus on people & downstream impact.
- Maximizes benefits to the commons by upstreaming work.
- Surfaces discrepancies between economic situation and work expectations.



FOSDEM'21

Tobie Langel (@tobie)
Principal, UnlockOpen

Thank you.

