

The New Dynamics of Open Source: Relicensing, Forks, & Community Impact

OpenForum Academy Symposium 2024

Dr. Dawn M. Foster

Director of Data Science for CHAOSS

Paper and replication data: <https://github.com/chaoss/wg-data-science/tree/main/publications>

 <https://chaoss.community/>

 <https://github.com/chaoss>

 @chaoss@fosstodon.org

 <https://fastwonderblog.com/>

 <https://github.com/geekygirldawn>

 @geekygirldawn@hachyderm.io

Thank You!



**ALFRED P. SLOAN
FOUNDATION**



**FORD
FOUNDATION**

Whoami



- Geek, traveler, reader
- 20+ yr tech career focused on open source (VMware, Intel, Puppet, ...)
- CHAOSS Board and Data Science
- OpenUK Board
- CNCF TAG Contrib Strategy co-chair
- PhD on Linux kernel collaboration

Photos by Mom, [Josh Bancroft](#), [Don Park](#)

Agenda

- **Overview and Research Question**
- **Methods**
- **Case Study Results**
- **Implications and Future Research**
- **Conclusion**

Photo by Marco Verch - CC BY 2.0

Overview: Company Controlled OSS

A single company in control has a higher risk of relicensing, forks, and other disruptions.

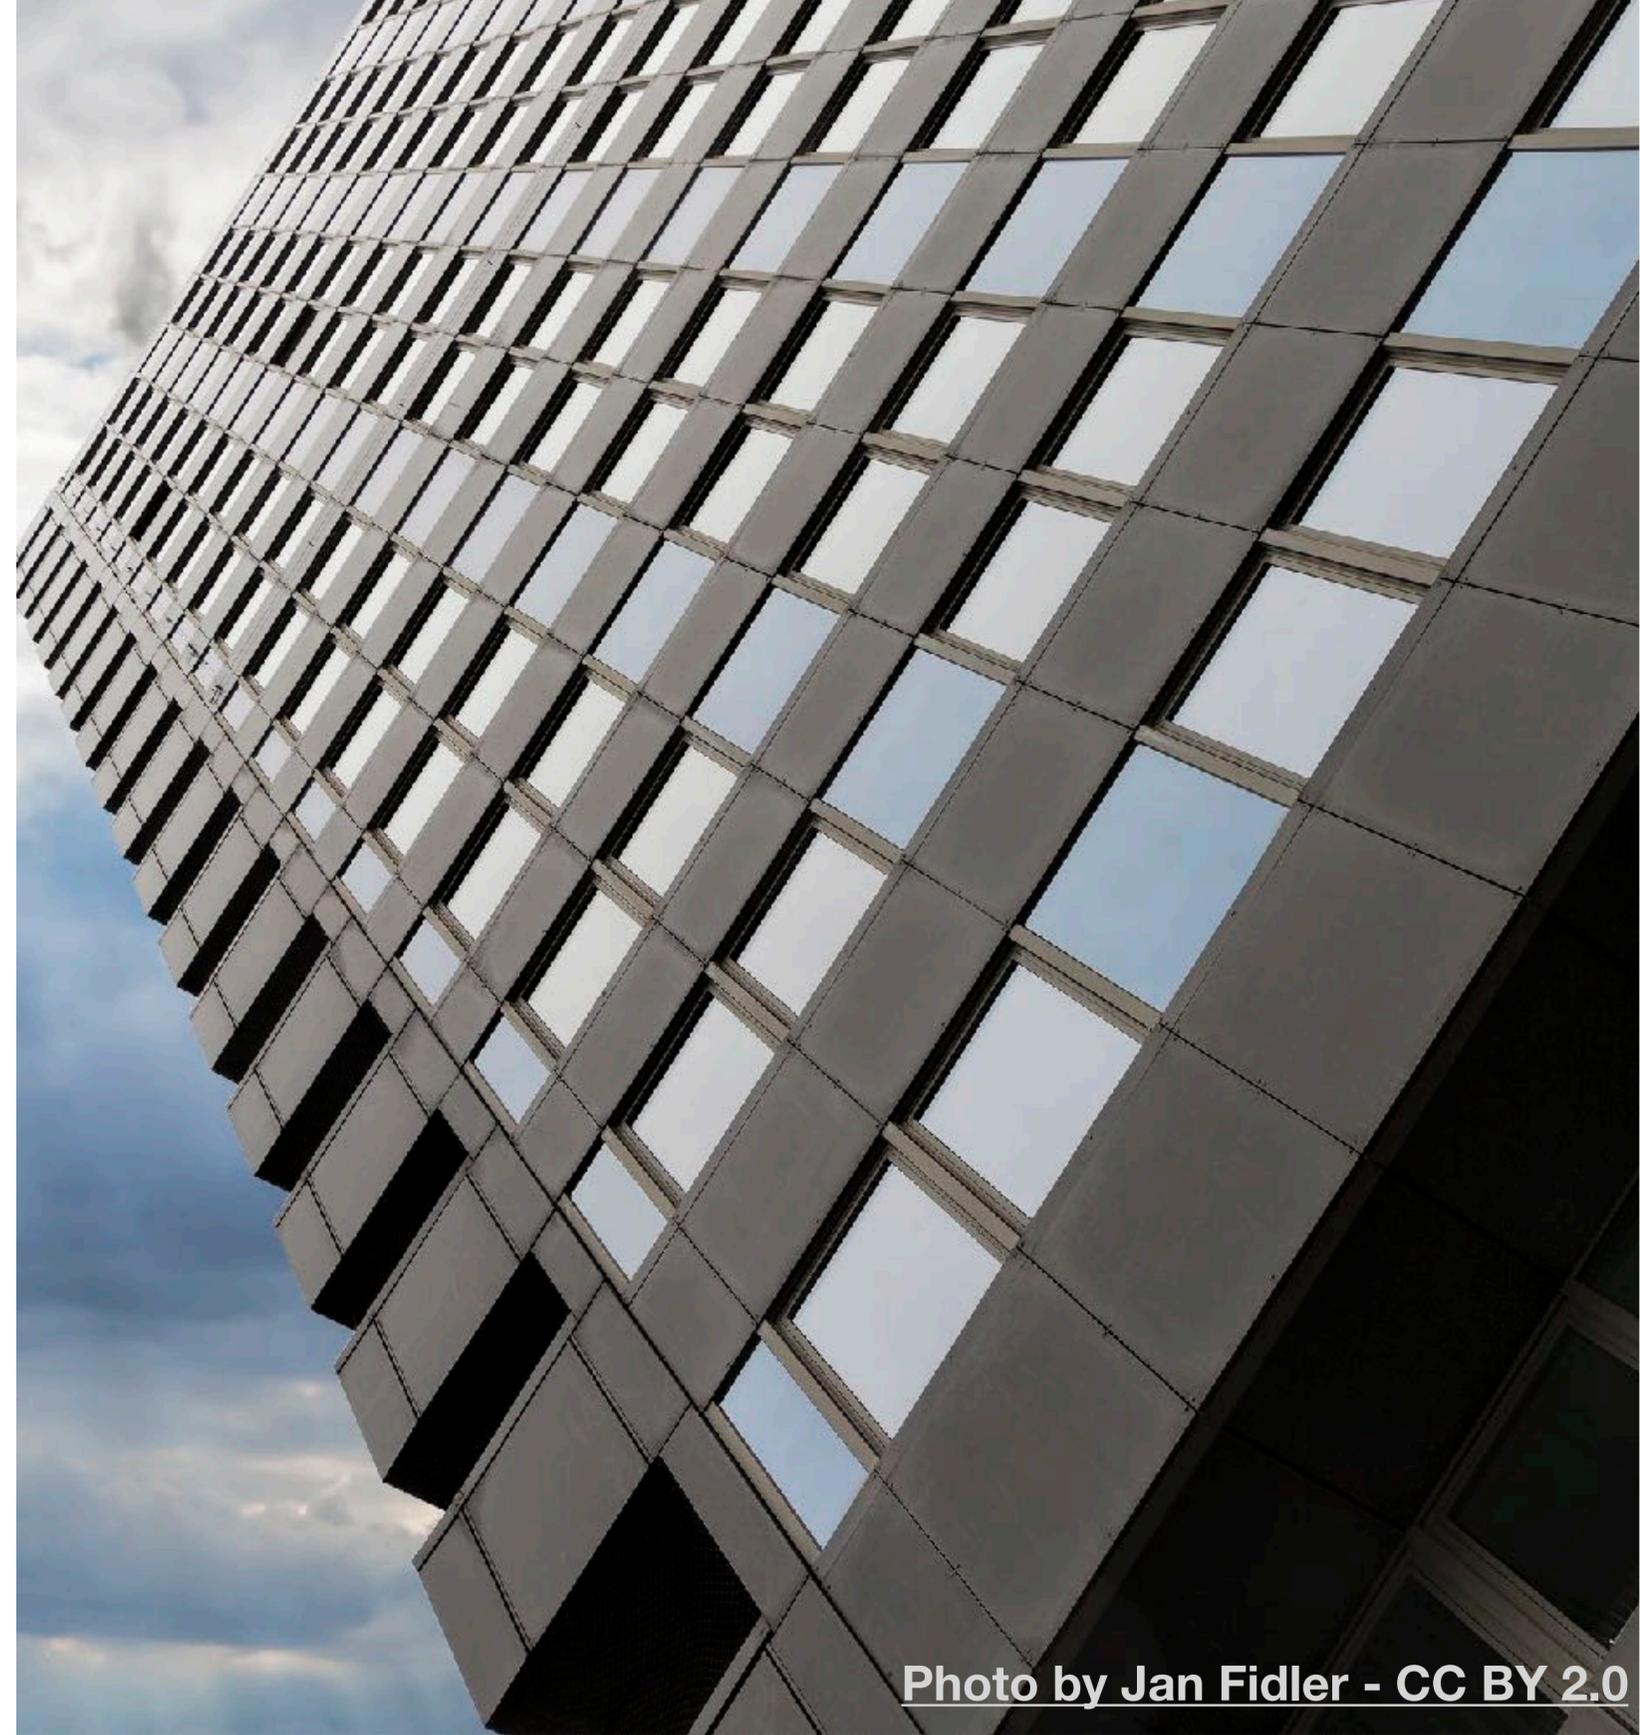


Photo by Jan Fidler - CC BY 2.0

Research Question

How do the organizational dynamics of an open source project evolve following a relicensing event, both within the original project and its resulting fork?



Research Overview

Profitability concerns ->
re-licensing ->
forking

Case Studies:

- Scenario 1: Elasticsearch / OpenSearch
- Scenario 2: Terraform / OpenTofu
- Scenario 3: Redis / Valkey

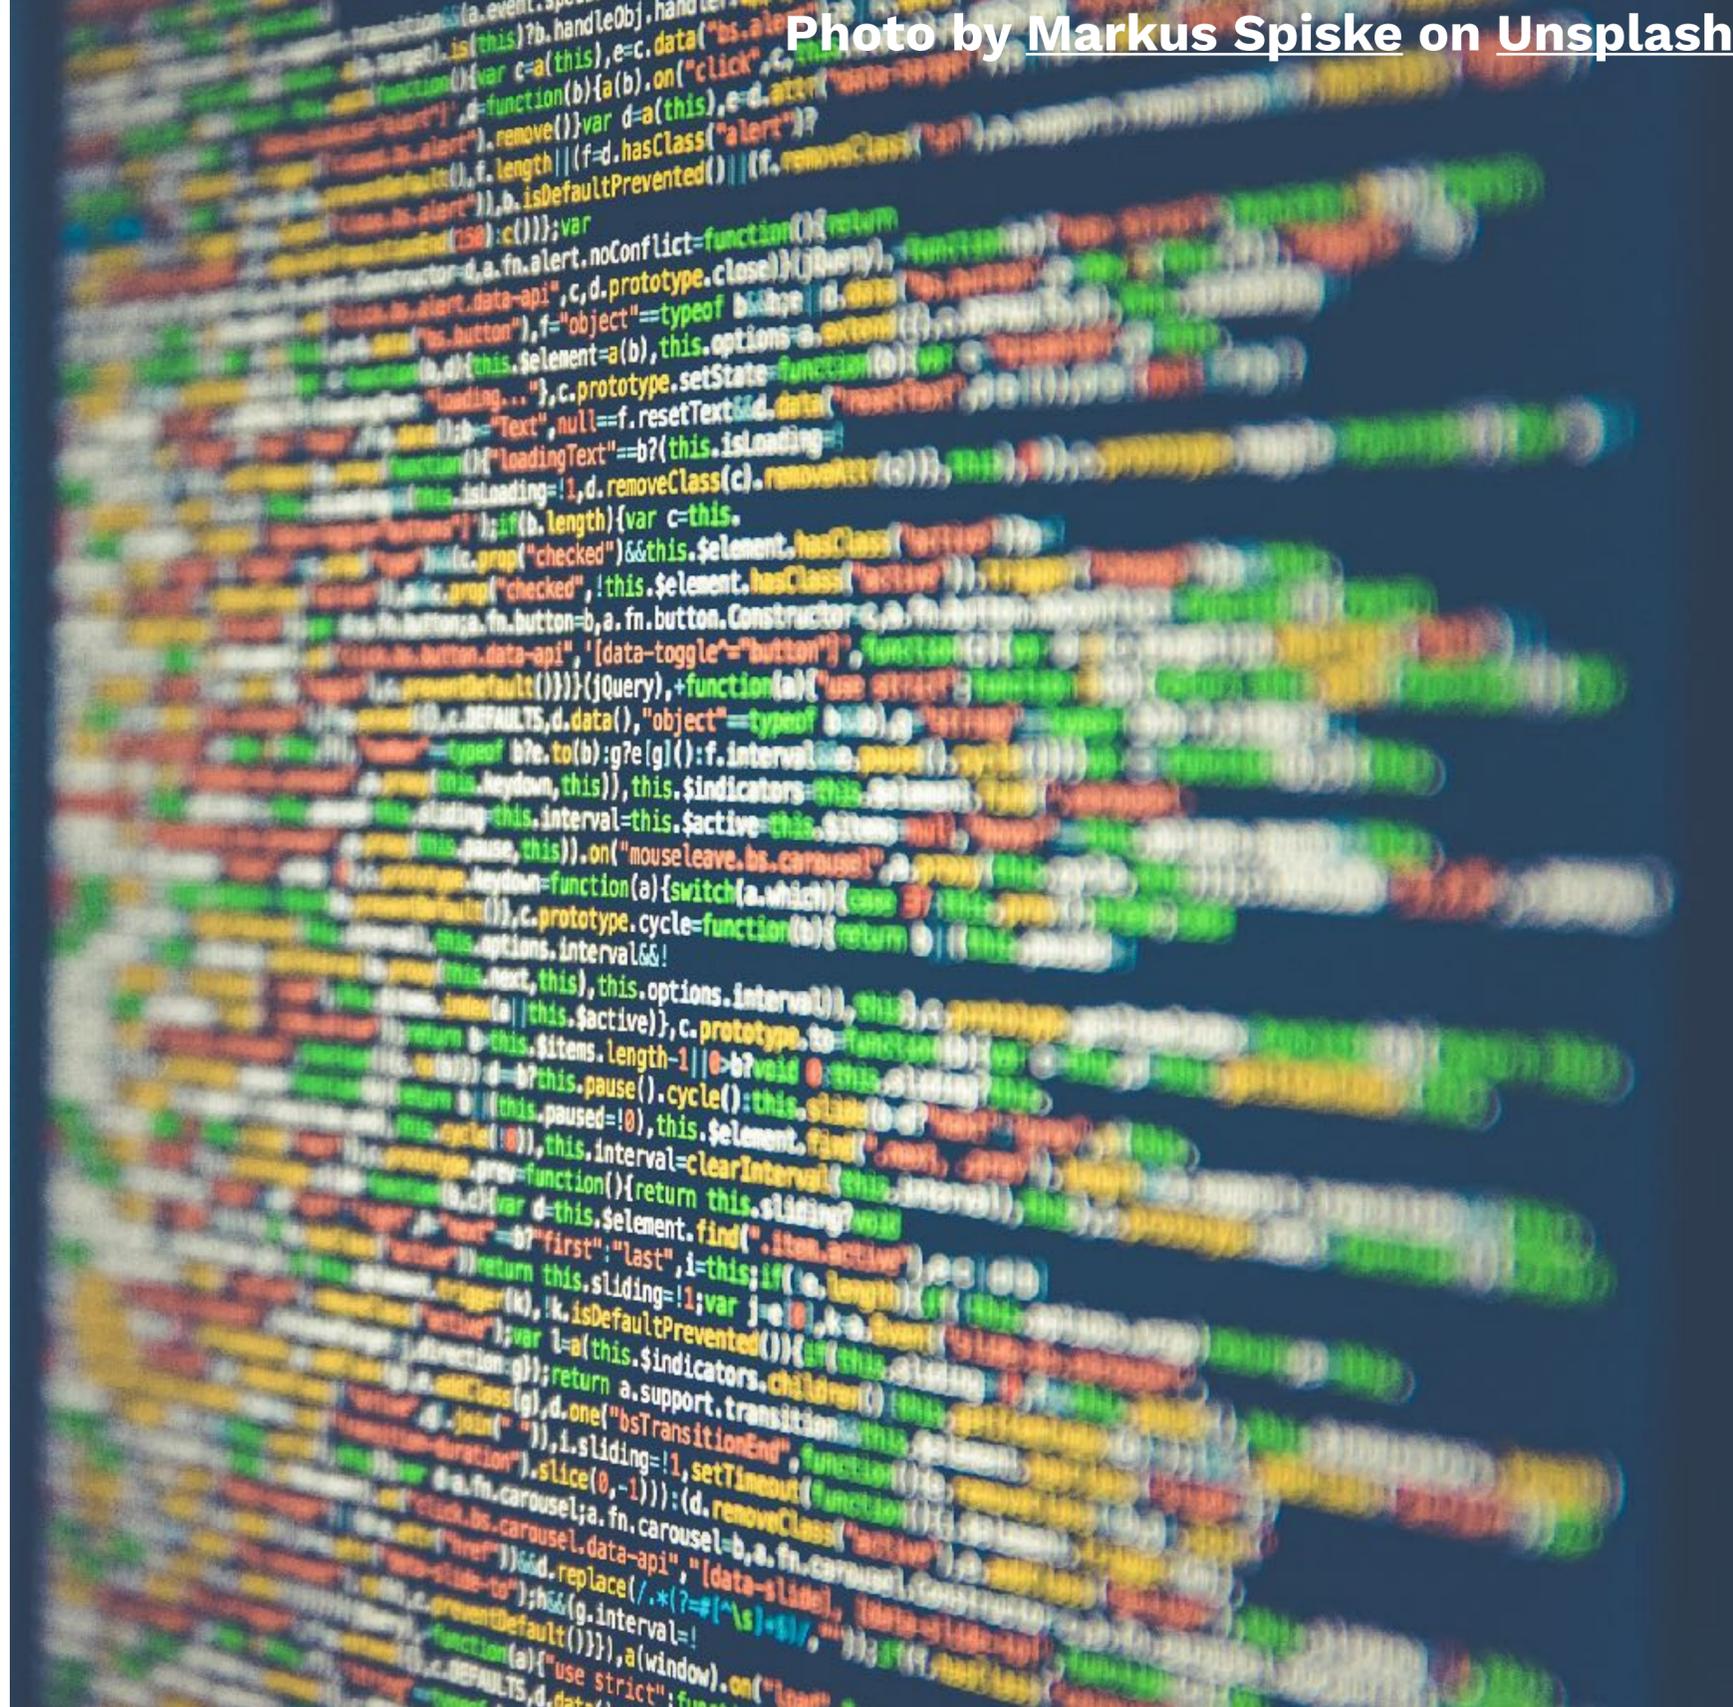


Methods

Case study approach

Commits and lines of code
added / deleted

Organizational affiliation
focus



Case Study Results



Scenario 1

**Contributions from vendor employees.
Fork created by new contributors and
owned by a company.**

Key Dates:

- 2021-02-03: Elastic relicenses to SSPL & Elastic
- 2021-04-12: OpenSearch fork owned by AWS
- 2024-08-29: Elastic adds AGPL
- 2024-09-16: OpenSearch moves under LF



elasticsearch



Elasticsearch repo (<https://github.com/elastic/elasticsearch>)

| Timeframe | Org Affiliation | People | Commits | Additions | Deletions |
|---|--|--------|-------------|-----------------|-----------------|
| 1 Year Before the Relicense (2020-02-03 - 2021-02-03) | Elastic employees with 10+ commits | 67 | 6,477 (92%) | 1,377,558 (96%) | 623,561 (97%) |
| | Non-Elastic employees with 10+ commits | 3 | 94 (1%) | 3,855 (<1%) | 740 (<1%) |
| 1 Year after the relicense (2021-02-03 - 2022-02-03) | Elastic employees with 10+ commits | 65 | 5,668 (91%) | 1,597,988 (96%) | 1,061,154 (98%) |
| | Non-Elastic employees with 10+ commits | 2 | 47 (1%) | 7,283 (<1%) | 2,178 (<1%) |
| 1 Year Before Adding AGPL (2023-08-29 - 2024-08-29) | Elastic employees with 10+ commits | 99 | 7,616 (95%) | 2,621,830 (95%) | 1,123,628 (97%) |
| | Non-Elastic employees with 10+ commits | 1 | 11 (<1%) | 326 (<1%) | 326 (<1%) |

OpenSearch Repo (<https://github.com/opensearch-project/OpenSearch>)

| Timeframe | Org Affiliation | People | Commits | Additions | Deletions |
|---|---------------------------------------|--------|-----------|---------------|---------------|
| 1 Year After the Fork (2021-04-12 to 2022-04-12) | Amazon employees with 10+ commits | 7 | 246 (34%) | 296,720 (80%) | 224,179 (91%) |
| | Non-Amazon employees with 10+ commits | 2 | 110 (15%) | 26,995 (7%) | 10,799 (4%) |
| 1 Year before LF (2023-09-16 to 2024-09-16) | Amazon employees with 10+ commits | 40 | 923 (49%) | 237,781 (63%) | 48,894 (65%) |
| | Non-Amazon employees with 10+ commits | 6 | 242 (13%) | 42,863 (11%) | 9,936 (13%) |

Scenario 2

Contributions from vendor employees. Fork created by new contributors in a foundation.



HashiCorp

Terraform

OpenTofu 

Key Dates:

- 2023-08-10: Terraform relicenses to BSL
- 2023-09-05: OpenTofu launched as fork under the LF

Terraform Repo (<https://github.com/hashicorp/terraform>)

| Timeframe | Org Affiliation | People | Commits | Additions | Deletions |
|--|---|--------|-------------|---------------|---------------|
| 1 Year before relicense (2022-08-10 - 2023-08-10) | HashiCorp employees with 5+ commits | 21 | 971 (82%) | 202,612 (93%) | 81,019 (95%) |
| | Non-HashiCorp employees with 5+ commits | 2 | 13 (1%) | 84 (<1%) | 33 (<1%) |
| 1 Year after relicense (2023-08-10 to 2024-08-10) | HashiCorp employees with 5+ commits | 24 | 1,620 (91%) | 672,393 (90%) | 242,052 (93%) |
| | Non-HashiCorp employees with 5+ commits | 2 | 18 (1%) | 353 (<1%) | 354 (<1%) |

OpenTofu Repo (<https://github.com/opentofu/opentofu>)

After 2023-09-05 fork until 2024-09-05:

- 31 people at 11 companies with ≥ 5 commits
- None of these people previously contributed to Terraform

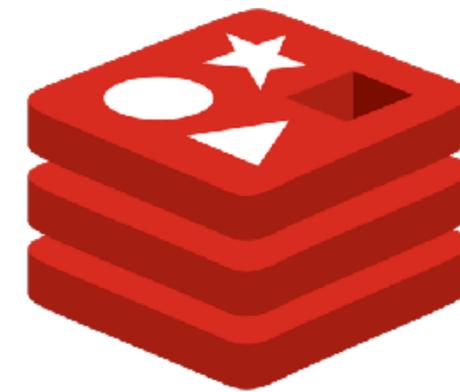
| People | Organization | Commits | Additions | Deletions |
|--------|----------------------------|---------|----------------|----------------|
| 10 | Spacelift | 328 | 88121 (55.21%) | 63992 (69.15%) |
| 6 | Env0 | 99 | 26507 (16.61%) | 12248 (13.23%) |
| 3 | Scalr | 47 | 12516 (7.84%) | 3374 (3.65%) |
| 3 | Harness | 17 | 2948 (1.85%) | 366 (0.40%) |
| 3 | Red Hat | 15 | 1605 (1.01%) | 159 (0.17%) |
| 1 | Hangzhou Dianzi University | 6 | 891 (0.56%) | 242 (0.26%) |
| 1 | Chainguard | 6 | 266 (0.17%) | 93 (0.10%) |
| 1 | lessops | 6 | 2017 (1.26%) | 226 (0.24%) |
| 1 | claranet | 6 | 118 (0.07%) | 20 (0.02%) |
| 1 | Cooby-inc | 5 | 72 (0.05%) | 69 (0.07%) |
| 1 | <u>nvdnc</u> | 5 | 68 (0.04%) | 11 (0.01%) |

Scenario 3

**Significant contributors not employed by the vendor.
Fork created by existing contributors in a foundation.**

Key Dates:

- 2024-03-20: Redis relicenses to SSPL & RSAL
- 2024-03-28: Valkey fork launched as an LF project



redis



Valkey

Redis Repo (<https://github.com/redis/redis>)

| Timeframe | Org Affiliation | People | Commits | Additions | Deletions |
|---|-------------------------------------|--------|-----------|---------------|--------------|
| 1 year before relicense (2023-03-20 - 2024-03-20) | Redis employees with 5+ commits | 6 | 164 (28%) | 189,656 (80%) | 83,122 (74%) |
| | Non-Redis employees with 5+ commits | 12 | 319 (54%) | 28,334 (12%) | 16,684 (15%) |
| 6 Months after relicense (2024-03-20 - 2024-09-20) | Redis employees with 5+ commits | 7 | 154 (74%) | 38,270 (75%) | 10,464 (72%) |
| | Non-Redis employees with 5+ commits | 0 | 0 | 0 | 0 |

Valkey Repo (<https://github.com/valkey-io/valkey>)

After 2024-03-28 Fork (until 2024-08-20):

- 29 people at 10 companies with ≥ 5 commits
- 18 of these people previously contributed to Redis

| People | Organization | Commits | Additions | Deletions |
|--------|---------------|---------|----------------|----------------|
| 13 | Amazon | 149 | 18232 (18.31%) | 6288 (8.66%) |
| 1 | Tencent Cloud | 92 | 4859 (4.88%) | 2429 (3.35%) |
| 4 | Huawei | 76 | 3561 (3.58%) | 3016 (4.16%) |
| 2 | Ericsson | 45 | 5867 (5.89%) | 1954 (2.69%) |
| 2 | Google | 39 | 40698 (40.86%) | 38643 (53.24%) |
| 1 | Intel | 12 | 632 (0.63%) | 464 (0.64%) |
| 1 | Alibaba | 8 | 415 (0.42%) | 71 (0.10%) |
| 1 | @gnet-io | 8 | 104 (0.10%) | 73 (0.10%) |
| 2 | ByteDance | 7 | 3952 (3.97%) | 572 (0.79%) |
| 2 | Samsung | 5 | 48 (0.05%) | 48 (0.07%) |

Summary of Results

| | |
|--|---|
| Scenario 1. Almost all contributions to the original project came from employees of the original vendor and the fork was created by new contributors and owned by a single company. | |
| Elasticsearch: Contributors are mostly Elastic employees both before and after the relicense. | OpenSearch: Contributors are mostly from Amazon, but organizational diversity is gradually improving. |
| Scenario 2. Almost all contributions to the original project came from employees of the original vendor and the fork was created by new contributors as a foundation project. | |
| Terraform: Contributors are mostly HashiCorp employees both before and after the relicense. | OpenTofu: 31 people employed at 11 companies, but none previously contributed to Terraform. |
| Scenario 3. The original project had significant contributors who were not employed by the original vendor and the fork was created by those existing contributors as a foundation project. | |
| Redis: Strong organizational diversity before the relicense, but only Redis employees after. | Valkey: 29 people employed at 10 companies have contributed, and 18 of them moved from Redis. |

Implications, Future Research, and Conclusion

Implications

Practice.

Vendors should carefully consider the decision to relicense. For adoption, projects that are dominated by a single vendor are at risk of relicensing / forking.

Research.

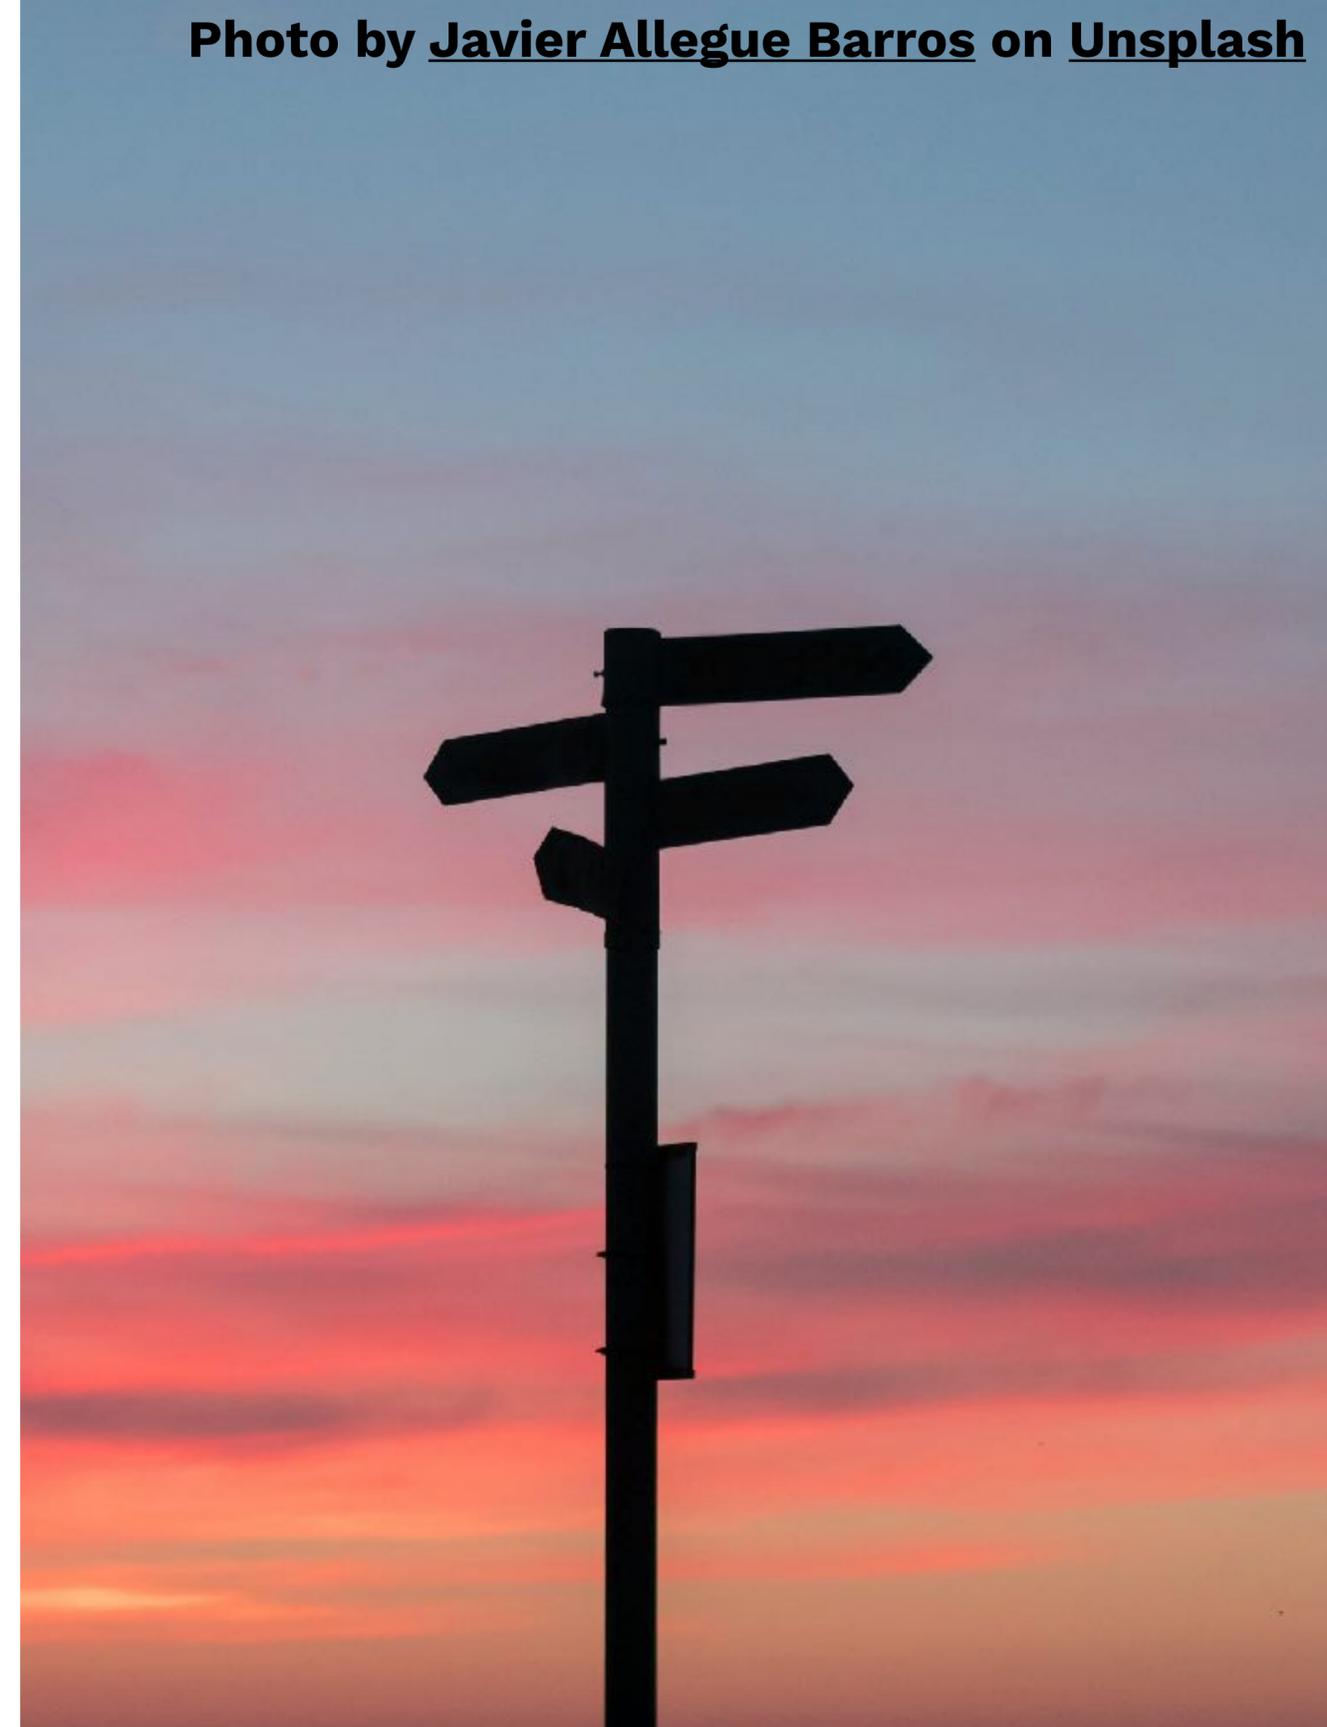
There is more to study beyond organizational affiliation. Highlights importance of considering organizational dynamics when performing OSS research.

Policy.

Consider organizational dynamics when making policy decisions related to adoption of OSS. Need to consider that projects may not be OSS forever. Funding programs should consider the risk of funding projects that are controlled by a single vendor.

Future Research

- Go beyond organizational affiliation with additional metrics and data about these projects.
- Understand impact of recent changes (e.g., Elasticsearch adding AGPL & OpenSearch moving to LF).
- Expand to additional cases.



Conclusion

**The forks coming from
relicensed projects have more
organizational diversity
than the original projects.**

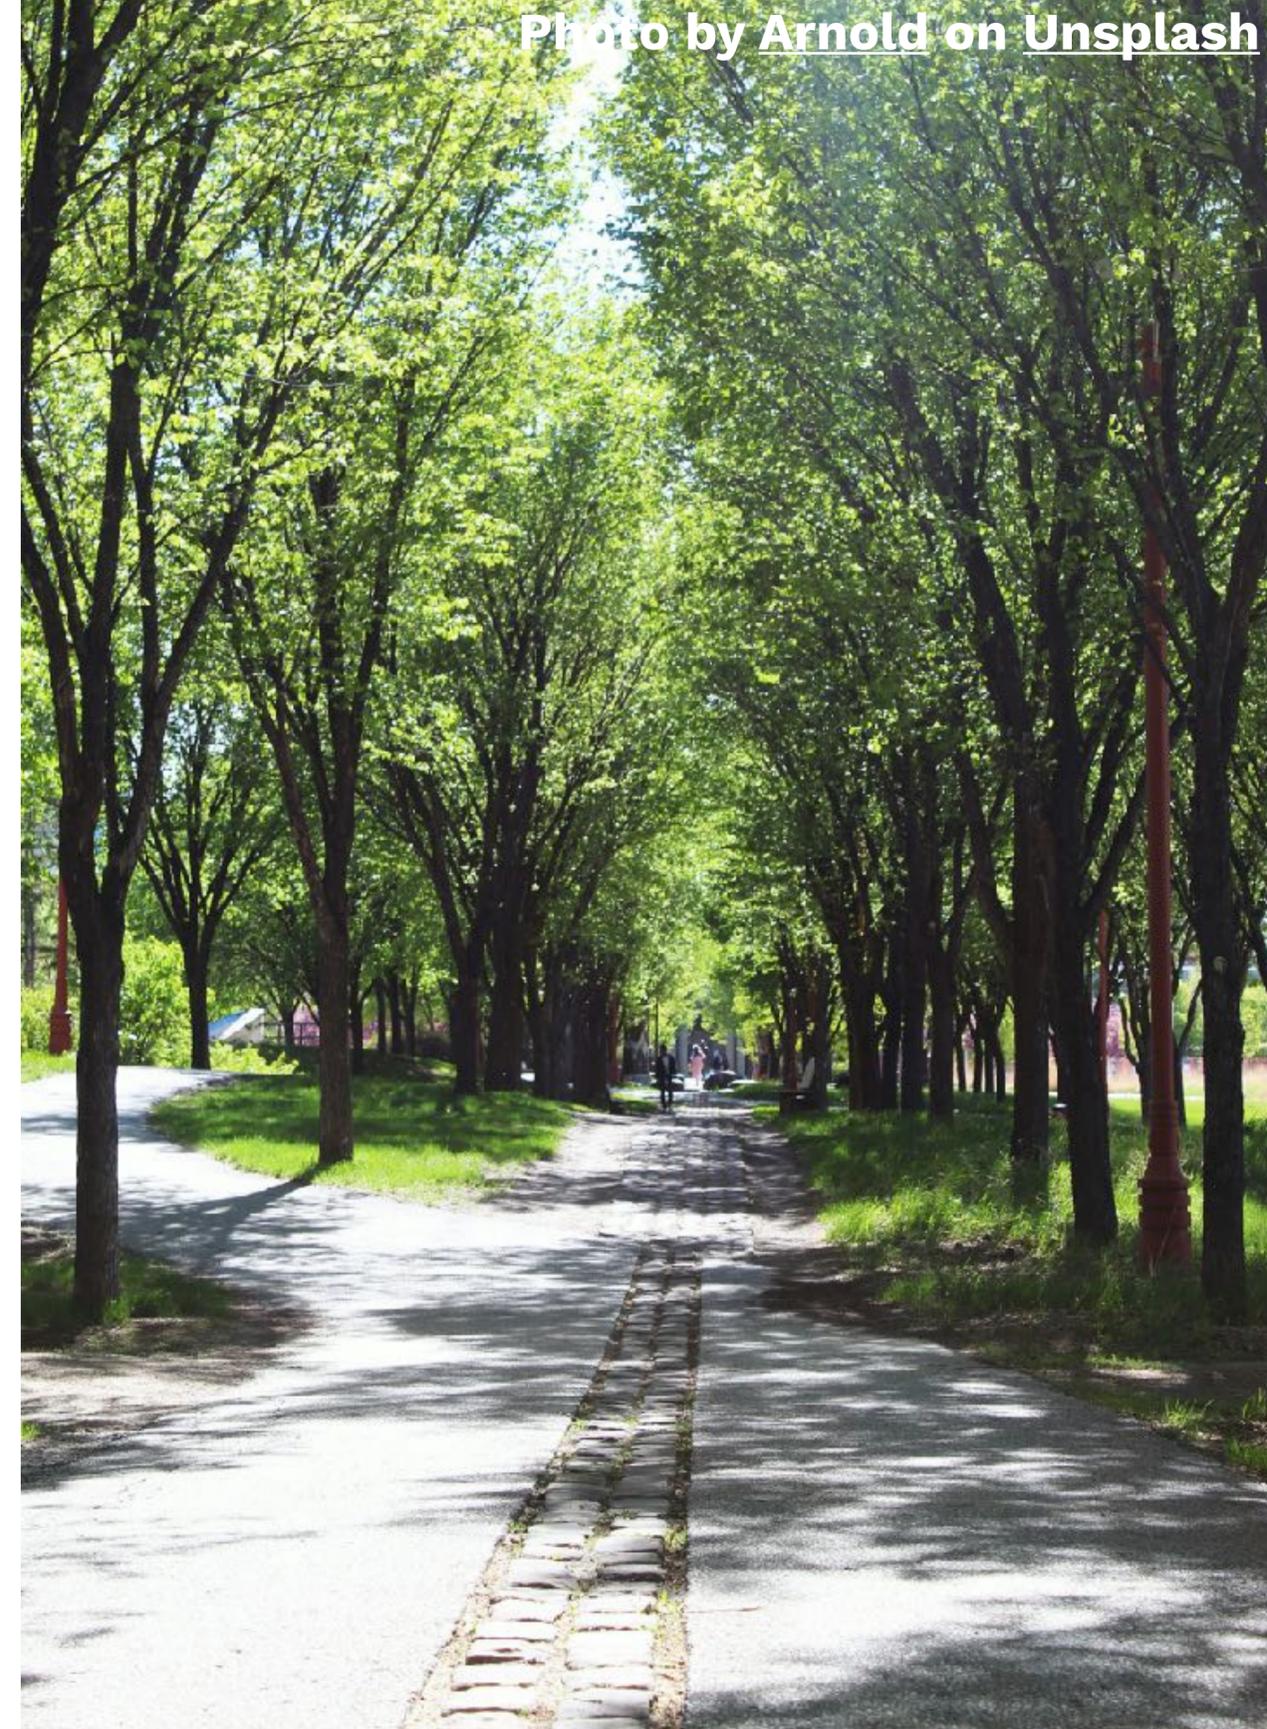


Photo by [Arnold](#) on [Unsplash](#)

THANK YOU!

Any Questions or Feedback?

 <https://chaoss.community/>
 <https://github.com/chaoss>
 @chaoss@fosstodon.org

 <https://fastwonderblog.com/>
 <https://github.com/geekygirldawn>
 @geekygirldawn@hachyderm.io

