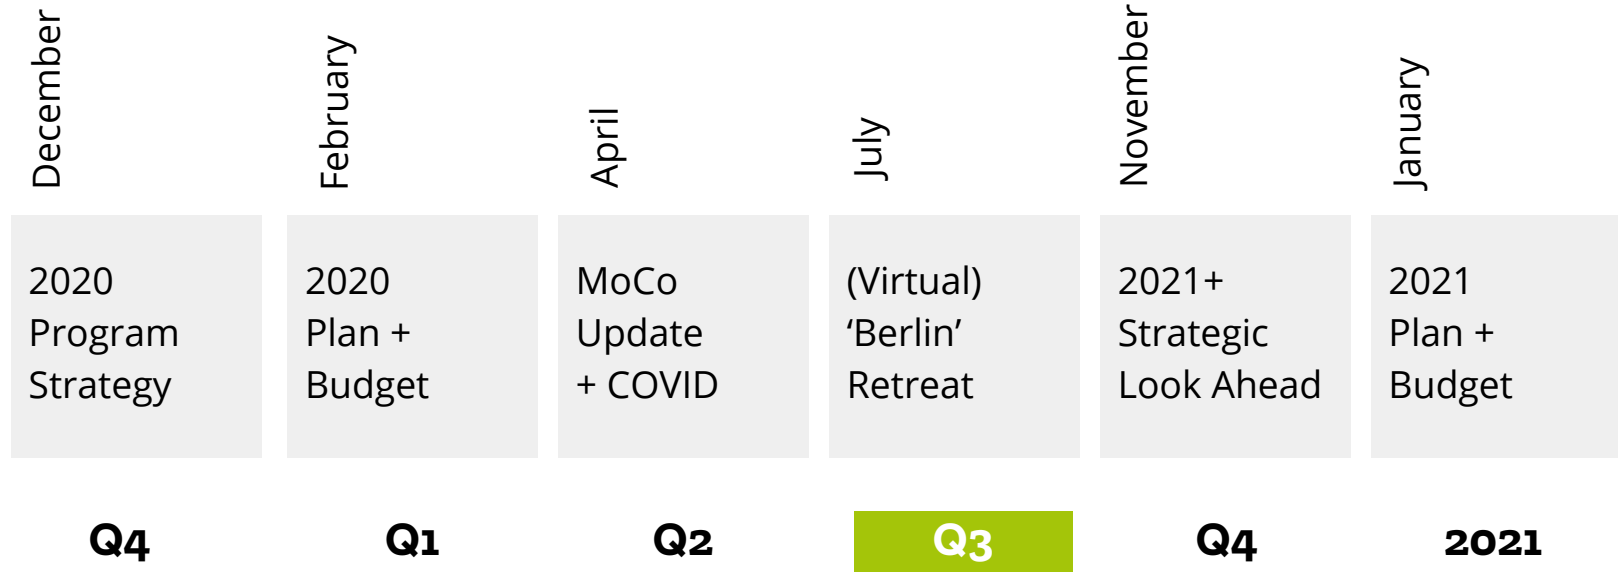


moz://a

Mozilla Foundation Board Meeting

July 13-15, 2020

2020 board workflow



Agenda

Monday July 13

2020 program + finance update

Tuesday July 14

2021+ scenario planning

Wednesday July 14

MoCo update



Monday,
July 13

Agenda

July 13

1. Board business
2. Racial justice commitments
3. MoFo program + finance update
4. Thunderbird update



Board Business

decisions

Board business

1. Approve minutes from May 2020 board meeting.
2. Approve resolution to open account for stock donations.
3. Update on racial justice commitments. Consider restart of board D+I committee?



Racial justice commitments

At the June all hands meeting, Mitchell and Mark made a series commitments on Black representation and racial justice, covering three areas:

1. Who we are, with commitments to **grow Black and Latinx representation** across Mozilla, including at the board and management levels.
2. What we build, with commitments to **use Firefox and Pocket to help people educate themselves** on how to combat systemic racism.
3. Broader engagement, **directing 40% of MoFo grants in 2020 to Black-led projects or organizations** and promoting the work of partner orgs focused on anti-racism.

We'll have more to share on this work as part of our 2021 plans. In the meantime, one question for the board, should we restart the D+I committee we chartered in 2018?

2020 Update

update

Reminder:
what we set out
to do in 2020

2020 objectives

1. Test out our theory of change in ways that both give momentum to other orgs taking concrete action on trustworthy AI and **establish Mozilla as a credible thought leader**.
2. **Increase data stewardship innovations** that can accelerate the growth of trustworthy AI.
3. Mobilize an influential consumer audience and **use pivotal moments to pressure companies** to make 'consumer AI' more trustworthy.
4. **Partner with diverse movements** to deepen intersections between their primary issues and internet health, including trustworthy AI, so that we increase shared purpose.
5. **Update our organizational models and capabilities** so that our strategy and people can succeed, and our ambition can grow over multiple years.



AI Theory of Change

Working on in 2020

Could add to our work

Short term outcomes
(1-3 years)

Medium term outcomes
(3-5 years)

Long term outcomes
(5+ years)

Long term impact

Clear "Trustworthy AI" guidelines emerge, leading to new and widely accepted industry norms.

Engineers, product managers, and designers with trustworthy AI training and experience are in high demand across industry.

Diverse stakeholders — including communities and people historically shut out of tech — are involved in the design of AI.

There is increased investment in and procurement of trustworthy AI products, services and technologies.

Shifting industry norms

The people building AI increasingly use trustworthy AI guidelines and technologies in their work.

More foundational trustworthy AI technologies emerge as building blocks for developers.

Transparency is included as a feature in more AI enabled products, services, and technologies.

Entrepreneurs develop — and investors support — alternative business models for consumer tech.

The work of artists and journalists helps people understand, imagine, and critique what trustworthy AI looks like.

Building new tech and products

Trustworthy AI products and services are increasingly embraced by early adopters.

Agency

All AI is designed with personal agency in mind. Privacy, transparency, and human well-being are key considerations.

Trustworthy AI products and services emerge that serve the needs of people and markets previously ignored.

Consumers are increasingly willing and able to choose products critically based on information regarding AI trustworthiness.

Citizens are increasingly willing and able to pressure and hold companies accountable for the trustworthiness of their AI.

A growing number of civil society actors are promoting trustworthy AI as a key part of their work.

Generating demand

Consumers choose trustworthy products when available and demand them when they aren't.

Accountability

Companies are held to account when their AI systems make discriminatory decisions, abuse data, or make people unsafe.

In a world of AI, consumer technology enriches the lives of human beings.

Governments develop the vision, skills, and capacities needed to effectively regulate AI, relying on both new and existing laws.

Progress towards trustworthy AI is made through wider enforcement of existing rules like the GDPR.

Regulators have access to the data they need to scrutinize the trustworthiness of AI in consumer products and services.

Governments develop programs to invest in and incent trustworthy AI.

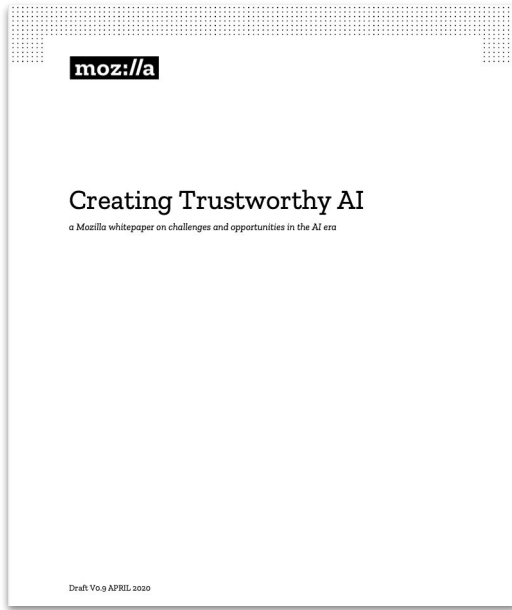
Creating regulations and incentives

New and existing laws are used to make the AI ecosystem more trustworthy.

Progress on 2020 objectives

OKR 1: thought leadership

Test out our theory of change in ways that both give momentum to other orgs taking concrete action on trustworthy AI and establish Mozilla as a credible thought leader.



Progress: **white paper v0.9 released**, community and partner RFC underway to generate interest and input.

Challenges: balancing exploration + focus. Where specifically should we collaborate on AI projects? Where should we simply offer an opinion?

Next steps: joining + developing **collaborations with** orgs focused on **developer, consumer, and policy audiences**. Also, exploring whether we can build standards or best practices for reviewing products.

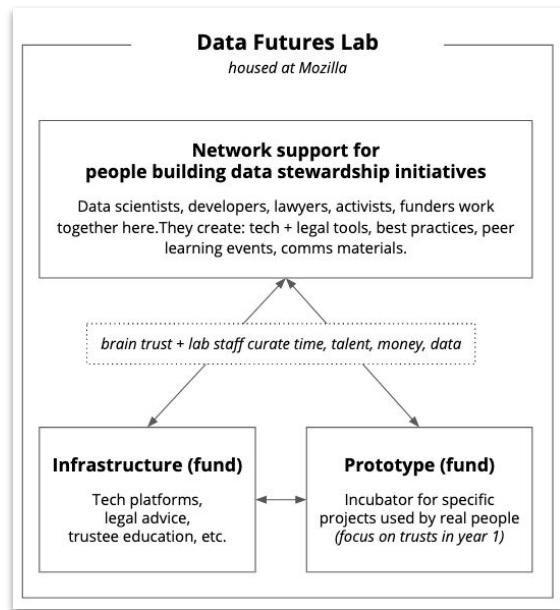
OKR 1: thought leadership

Test out our theory of change in ways that both give momentum to other orgs taking concrete action on trustworthy AI and establish Mozilla as a credible thought leader.

<i>KEY RESULT</i>	<i>BASELINE</i>	<i>PROGRESS</i>	<i>TARGET</i>
Publish a white paper theory of change	0	Feedback from 75+ parties so far, incl EU, funders + partners.	1
250 people and organizations participate in mapping to show who is working on key elements of trustworthy AI and offer feedback on the white paper	0	Mapping via interviews and research in Q3, will help ID new partners.	250
25 collaborations with partners working on concrete projects that align with short term outcomes in the theory of change	8	New partnerships established w/ IEEE, PAI, City of Amsterdam, etc.	25

OKR 2: data stewardship

Increase the number of data stewardship innovations that can accelerate the growth of trustworthy AI.



Progress: developed initial design for **Data Futures Lab** with Luminate. 'Protolab' phase underway, with initial fellows in place plus first grants by end of Q3 2020.

Challenges: major **gap between theory and practice**. Growing field talking about data governance, but few people actually building new things yet.

Next steps: securing first tranche of Luminate funding, finishing **data stewardship landscape analysis** and selecting projects for protolab collaboration / grants.

OKR 2: data stewardship

Increase the number of data stewardship innovations that can accelerate the growth of trustworthy AI.

<i>Key result</i>	<i>Baseline</i>	<i>Progress</i>	<i>Target</i>
\$3 million raised to support bold, multi-year, cross movement initiatives on data stewardship as an indicator of growing philanthropic support in this area.	0	\$1.5M close to closing from Luminare, seeking \$1.5M from others.	\$3M
10 awards or fellowships for prototypes or other concrete exploration re: data stewardship.	5	8 so far, working w/ fellows to identify prototypes and other concrete explorations.	10
4 concentric “networks of practice” utilize Mozilla-housed Data Futures Lab.	0	Landscape analysis work has generated interest from Aapti, Ada, WAAG, platform coops, etc.	4

OKR 3: consumer power

Mobilize an influential consumer audience using pivotal moments to pressure companies to make 'consumer AI' more trustworthy.



Progress: established **MoFo channels deployed quickly to respond to pandemic and protests**. E.g. Zoom security campaign. Driving 3x web traffic over this time last year.

Challenges: not all timely topics relate to AI. Leveraging **pivotal moments = less control over content focus**.

Next steps: focusing on **election related AI topics** (e.g. misinfo and recommendation systems), adding AI ratings to holiday PNI + AI topics in campaign pipeline.


OKR 3: consumer power

Mobilize an influential consumer audience using pivotal moments to pressure companies to make 'consumer AI' more trustworthy.

<i>Key result</i>	<i>Baseline</i>	<i>Progress</i>	<i>Target</i>
3m visits to 'trustworthy AI' content on Mozilla channels (website, social media, YouTube, etc.).	2M	Timely content = 1.7m page views YTD (3x YOY).	3M
50k new subscribers drawn from sources (partnerships, contextual advertising, etc.) oriented towards people ages 18-35.	263k	32k new followers, w/ instagram push to get younger demographic.	313k
25k people share information with us (stories, browsing data, etc.) in order to gather evidence about how AI currently works and what changes are needed.	YouTube Regrets = 1.5k for useable stories	Several data gathering projects ready to launch Dependency w/MoCo has delayed 3 projects.	25k

OKR 4: movement building

Partner with diverse movements to deepen intersections between their primary issues and internet health, including trustworthy AI, so that we increase shared purpose.



Progress: a number of **high profile orgs using shared language w/ Mozilla**, including City of Amsterdam AI procurement policy and House of Lords political ads report.

Challenges: it's not just number of partnerships that matter, but also **depth of partnership** (e.g. number of teams across Mozilla working w/ the same partner).

Next steps: movement building landscape analysis plus public launch of **Ford partnership in global south**.


OKR 4: movement building

Partner with diverse movements to deepen intersections between their primary issues and internet health, including trustworthy AI, so that we increase shared purpose.

<i>Key result</i>	<i>Baseline</i>	<i>Progress</i>	<i>Target</i>
30% increase in partners with whom we (have both) published, launched, or hosted something that includes shared approaches to their issues and internet health (e.g. language, methodologies, resources, or events).	332	Reviewing baseline, depth of partnership may be more important than volume.	432
75% of partners from these diverse movements report deepening intersection between their issues and internet health/AI.	0%	47% of partners.	75%
4 new partnerships in the Global South report deepened intersection between their work and ours.	0	16 host orgs in the global south selected. Next step: survey.	4

OKR 5: organizational effectiveness

Update our organizational models and capabilities so that our strategy and people can succeed, and our ambition can grow over multiple years.



Progress: developed **new financial model** that we can use to study historical trends and develop future scenarios.

Challenges: tricky **balancing long range program vs financial planning**. There are chicken and egg questions. Board input can help here.

Next steps: developing **multi year 2021+ plans** based on scenario planning. Also, new staff survey tool based on non-profit best practices, will help us sharpen execution.

OKR 5: organizational effectiveness

Update our organizational models and capabilities so that our strategy and people can succeed, and our ambition can grow over multiple years.

<i>Key result</i>	<i>Baseline</i>	<i>Progress</i>	<i>Target</i>
70% overall favorable score for staff perception of organizational effectiveness.	New factor	Rolling out more robust survey tool (instead of just adding 'new factor').	70%
A long-term funding model that matches our ambitions exists.	0	Financial model delivered, scenario planning underway.	1
100% of staff strengthen skills (through PD) needed to accomplish our objectives.	55%	31% of staff requested PD through end of June; prioritizing org-wide job description audit for H2.	100%

2020 Finances

To do: 2021 reporting + budget resolution

At finance committee, we discussed **changing the way we formulate the annual budget resolution** starting in 2021.

Right now, the resolution is based around total unrestricted assets at the end of the year, which will fluctuate w/ our investments. Our aim is to **modify the format of the resolution to focus on real capital spending / growth / break even within a range.**

We will work with the finance committee and come back to the board with a proposal for both the resolution and for some changes to how we report on our investment results accordingly.



Thunderbird update *(Brian)*

Thunderbird update

Revenue from **donations continued to be strong** even during early COVID months.

Team is focused on product improvements in two areas: a) **modernizing the user experience** and b) adding **end to end encryption as a core feature** (vs. an add-on).

As the MZLA board, our **main concern is business leadership of the project.**



Tuesday,
July 14

Agenda

July 14

1. 2021+ scenarios
2. Trustworthy AI: what's next?
3. Discussion and questions



Our focus *(enduring)*

moz://a

Internet Health Report **2019**

How healthy is the internet?

Our 2019 compilation of research and stories explains what's key to a healthier internet across five issues, from personal experience to global concerns.

MoFo's overall focus remains creating a healthier internet by **supporting a movement of people and allies around the world** who share our cause ...

Our focus *(right now)*

moz://a

Internet Health Report **2019**

Spotlight: Let's ask more of AI



... and the **AI work we started in 2019 adds focus and momentum** to our overall internet health agenda.

As a next step ...

At the beginning of 2020, we agreed to **map out a rough, three year trajectory for MoFo's movement building programs**, including next steps for our trustworthy AI work.

Since then, **the world has changed**: the pandemic may send us into an **economic downturn** *and* the wave of protests in the US has increased **attention on racial justice** and inclusion.

Given the setting, we have developed initial **program scenarios that consider both growth and decline**. We looked at four variables: program ambition; revenue; capabilities; public resonance.

AI strategy scenarios

In addition to running resource and capability scenarios, we are asking ourselves: **where should we will take our trustworthy AI work next?** This will feed our three year plan.

Assumption: we'll continue to pursue the threads in our 2020 OKRs: industry norms; data stewardship; consumer power; and working w/ other movements. *See example on next slide.*

However, we could consider whether we want to expand, narrow or change the focus of our trustworthy AI work if the scale of our programs changes. *See subsequent slides.*

Example of a plan *(ie. what we might do on data stewardship over three years)*

2020

Initial cohort of people with data stewardship innovations help us to prototype the Data Futures Lab.

Feasibility study for Mozilla data trust / coop.

2021

Tools and approaches in data stewardship landscape well documented.

Prototypes from initial cohort are working with real users.

Initial infrastructure projects funded and under development.

Mozilla data trust / coop built if deemed feasible.

2022

Infrastructure investments show up as working legal or technical models, are deployed by others.

Number of working prototypes grows, some succeed and some fail.

We begin to really see what works.

Mozilla data trust / coop generating revenue.

2023

How consent works begins to shift -- norms begin to change.

Many users know how to pick products and services that give value back for their data.

Data stewardship innovations start to accelerate the growth of trustworthy AI.

Mozilla data trust / coop driving product change across industry.

AI Theory of Change

Working on in 2020

Could add to our work

Short term outcomes
(1-3 years)

Medium term outcomes
(3-5 years)

Long term outcomes
(5+ years)

Long term impact

Clear "Trustworthy AI" guidelines emerge, leading to new and widely accepted industry norms.

Engineers, product managers, and designers with trustworthy AI training and experience are in high demand across industry.

Diverse stakeholders — including communities and people historically shut out of tech — are involved in the design of AI.

There is increased investment in and procurement of trustworthy AI products, services and technologies.

Shifting industry norms

The people building AI increasingly use trustworthy AI guidelines and technologies in their work.

More foundational trustworthy AI technologies emerge as building blocks for developers.

Transparency is included as a feature in more AI enabled products, services, and technologies.

Entrepreneurs develop — and investors support — alternative business models for consumer tech.

The work of artists and journalists helps people understand, imagine, and critique what trustworthy AI looks like.

Building new tech and products

Trustworthy AI products and services are increasingly embraced by early adopters.

Agency

All AI is designed with personal agency in mind. Privacy, transparency, and human well-being are key considerations.

Trustworthy AI products and services emerge that serve the needs of people and markets previously ignored.

Consumers are increasingly willing and able to choose products critically based on information regarding AI trustworthiness.

Citizens are increasingly willing and able to pressure and hold companies accountable for the trustworthiness of their AI.

A growing number of civil society actors are promoting trustworthy AI as a key part of their work.

Generating demand

Consumers choose trustworthy products when available and demand them when they aren't.

Accountability

Companies are held to account when their AI systems make discriminatory decisions, abuse data, or make people unsafe.

In a world of AI, consumer technology enriches the lives of human beings.

Governments develop the vision, skills, and capacities needed to effectively regulate AI, relying on both new and existing laws.

Progress towards trustworthy AI is made through wider enforcement of existing rules like the GDPR.

Regulators have access to the data they need to scrutinize the trustworthiness of AI in consumer products and services.

Governments develop programs to invest in and incent trustworthy AI.

Creating regulations and incentives

New and existing laws are used to make the AI ecosystem more trustworthy.

AI Theory of Change

Working on in 2020

Could add to our work

Short term outcomes
(1-3 years)

Medium term outcomes
(3-5 years)

Long term outcomes
(5+ years)

Long term impact

Clear "Trustworthy AI" guidelines emerge, leading to new and widely accepted industry norms.

Engineers, product managers, and designers with trustworthy AI training and experience are in high demand across industry.

Diverse stakeholders — including communities and people historically shut out of tech — are involved in the design of AI.

There is increased investment in and procurement of trustworthy AI products, services and technologies.

Shifting industry norms
The people building AI increasingly use trustworthy AI guidelines and technologies in their work.

More foundational trustworthy AI technologies emerge as building blocks for developers.

Transparency is included as a feature in more AI enabled products, services, and technologies.

Entrepreneurs develop — and investors support — alternative business models for consumer tech.

The work of artists and journalists helps people understand, imagine, and critique what trustworthy AI looks like.

Building new tech and products
Trustworthy AI products and services are increasingly embraced by early adopters.

Trustworthy AI products and services emerge that serve the needs of people and markets previously ignored.

Consumers are increasingly willing and able to choose products critically based on information regarding AI trustworthiness.

Citizens are increasingly willing and able to pressure and hold companies accountable for the trustworthiness of their AI.

A growing number of civil society actors are promoting trustworthy AI as a key part of their work.

Generating demand
Consumers choose trustworthy products when available and demand them when they aren't.

Governments develop the vision, skills, and capacities needed to effectively regulate AI, relying on both new and existing laws.

Progress towards trustworthy AI is made through wider enforcement of existing rules like the GDPR.

Regulators have access to the data they need to scrutinize the trustworthiness of AI in consumer products and services.

Governments develop programs to invest in and incent trustworthy AI.

Creating regulations and incentives
New and existing laws are used to make the AI ecosystem more trustworthy.

Agency

All AI is designed with personal agency in mind. Privacy, transparency, and human well-being are key considerations.

In a world of AI, consumer technology enriches the lives of human beings.

Accountability

Companies are held to account when their AI systems make discriminatory decisions, abuse data, or make people unsafe.

Questions to discuss

In the growth scenario, **where do we want to invest in order to drive growth?** Deep technical horsepower? Big bilateral funder coalitions? A data trust or membership program?

For both scenarios, are there particular **trustworthy AI outcomes that we should double-down on?** And, do we add a stronger racial justice lens to our AI work?

Also, where can we tune MoFo's focus to help MoCo unlock innovation? And, vice versa, what could MoCo do to unlock MoFo growth? **Where can Pan Mozilla alignment help most?**

Appendix B

Notional 3-Year AI Plans

Inputs and outputs

Inputs	Outputs
<ul style="list-style-type: none">● Our movement building pillars: shape the agenda, rally citizens, connect leaders● Theory of change + white paper● 2020 org-wide OKRs● Team plans, expertise, learnings	<ul style="list-style-type: none">● 3-year org-wide objectives (2021-2023)● MoFo strategy update, including updated org structure● 3-year resource plan and fundraising strategy

Thought leadership -> industry norms *(notional)*

2020

Test out our theory of change in ways that both give momentum to other orgs taking concrete action on trustworthy AI and establish Mozilla as a credible thought leader.

2021

Minimum 'best practices' for Trustworthy AI product development are published in collaboration with developers & developer orgs (i.e, here's how you should build) and begin to be used for evaluation.

2022

More organizations and journalists promote the 'best practices' components when building and evaluating products.

Professional orgs include 'best practices' elements in trainings and certifications.

2023

Trustworthy AI 'best practices' become a standard element of evaluating and discussing AI-enabled consumer products by journalists and organizations.

Professional engineering, design and product management entities develop formal training pathways to teach and evolve the 'best practices' framework.

Data stewardship *(notional)*

2020

Initial cohort of people with data stewardship innovations help us to prototype the Data Futures Lab.

Feasibility study for Mozilla data trust / coop.

2021

Tools and approaches in data stewardship landscape well documented.

Prototypes from initial cohort are working with real users.

Initial infrastructure projects funded and under development.

Mozilla data trust / coop built if deemed feasible.

2022

Infrastructure investments show up as working legal or technical models, are deployed by others.

Number of working prototypes grows, some succeed and some fail.

We begin to really see what works.

Mozilla data trust / coop generating revenue.

2023

How consent works begins to shift -- norms begin to change.

Many users know how to pick products and services that give value back for their data.

Data stewardship innovations start to accelerate the growth of trustworthy AI.

Mozilla data trust / coop driving product change across industry.

Consumer power *(notional)*

2020

Mobilize an influential consumer audience using pivotal moments to pressure companies to make 'consumer AI' more trustworthy.

2021

Consumer Index of Trustworthy AI tool developed.

Terminology of 'trustworthy AI' begins to appear in press.

Best practices begin to serve as basis for Mozilla and others to evaluate consumer tech products.

Consumer organizations begin to develop interest and expertise in Trustworthy AI.

2022

Consumer Index of Trustworthy AI tool used as a transparency & accountability lever by Mozilla and global consumer & civil society orgs.

Trustworthy AI terminology and best practices are tested and appear regularly in mainstream press.

A number of products changed or removed based on concerns about trustworthiness.

2023

Consumer Index of Trustworthy AI becomes standard tool to evaluate consumer tech + AI.

Global consumer organizations regularly cite trustworthy AI best practices.

Consumer tech companies report on their practices to produce trustworthy AI products.

Third party sellers and major retailers use Consumer Index to guide product decisions.

Movement building *(notional)*

2020

Partner with diverse movements deepen intersections between their primary issues and internet health, including trustworthy AI, so that we increase shared purpose.

2021

Increase in technical ability and/or our ability to partner alongside social movements, geographically diverse civil society actors.

2022

Develop a more interdependent relationship with other movements with which we intersect.

Philanthropic advocacy builds our credibility and trust in the field, boosting more \$ into the field and more unrestricted investment in MoFo work.

2023

Policy that supports the functional development of technologies that are meeting the needs of African civil society - led by those on the ground.

moz://a

Mozilla Foundation Board Meeting

July 13-15, 2020