Election Misinformation Partnership

[Introductory Deck, WIP]

Goals for Introductory Meeting (7/9)

- Understanding the problem space: what are the needs in CISA's misinformation response, and what are the current state of affairs?
- Understanding SIO: what are the capabilities here, and how could they help?
- Understanding the proposal: what are some next steps?

Stanford Internet Observatory

The Stanford Internet Observatory (SIO) is a cross-disciplinary program of research, teaching, and policy engagement for the study and abuse in current information technologies, with a focus on social media.

Key Capabilities

- Experienced disinformation research team of analytical and technical talent
- Real-time narrative tracking capabilities for all major platforms (**Facebook**, **Instagram**, **Twitter**, **Reddit**, potential for **TikTok**)
- Additional API or historical access to 'fringe' platforms (**Gab**, **Parler**, **4Chan**).
- Established and collaborative node within the third-party misinformation research ecosystem.

Problem Statement

Voters in November will largely (if not exclusively) look for real time election information on social media. Election mis and disinformation on social media is therefore one of the largest obstacles to ensuring a safe and fair election.

There are more than 10,000 election administration jurisdictions in the U.S. **Currently, there is no central organization to support elections officials or CISA in identifying and responding to social-media based misinformation in their district in real time.**

Current Landscape

Who are some players that could potentially solve this problem? Why aren't they?

	CISA	Platforms	Academic/Research Institutions
Currently Offers	EI-ISAC collaboration to provide real-time monitoring tools such as the SOC as well as the classified and unclassified Situation Rooms	Direct contact with secretaries of state as well as some cross-platform communication on this front	Institutions have created theiro
Strengths	Direct communication with every election official, central node in the election infra ecosystem	Highest monitoring capacity into what is happening in the social landscape, lots of $$$ and resources	'Easiest' politically, transparent, existing institutions (SIO). Agile, lightweight teams.
Weaknesses	Most efforts focused on hardware, smaller misinformation workstream, govt entity, can't be seen as 'monitoring' the electorate, highly political.	Political, easily seen as partisan, don't have the direct communication/rapport with all election officials.	Don't have the direct communication or rapport with all election officials, need to raise $$$

Proposal

Leading up to, during, and after the presidential election this November, there should exist a coalition of third party research entities to support real-time informational exchange between the online mis and disinformation research community, state and local governments, and media platforms.

Key Features…

- **Open by default.** This will be a collaborative communication channel for spot research by academics on potential election-day info operations on American social media
- **Well defined.** The exact scenarios during which action is and is not required by such an entity will be decided upon beforehand.
- **Independent**. Main value add of third party research is transparency and independence, when compared to platform or government solutions.

Goals and Non-Goals of this Project

Goal: To detect and mitigate the impact of misinformation that could:

- Prevent people from voting
- Deter people from voting
- Delegitimize election results without factual basis

Ex: "Candidate X is no longer on the ballot!"

Non-Goal: Detecting and mitigating political misinformation not related to the mechanics of the election.

Ex: "Candidate X is a criminal!"

Major Stakeholders *(link)*

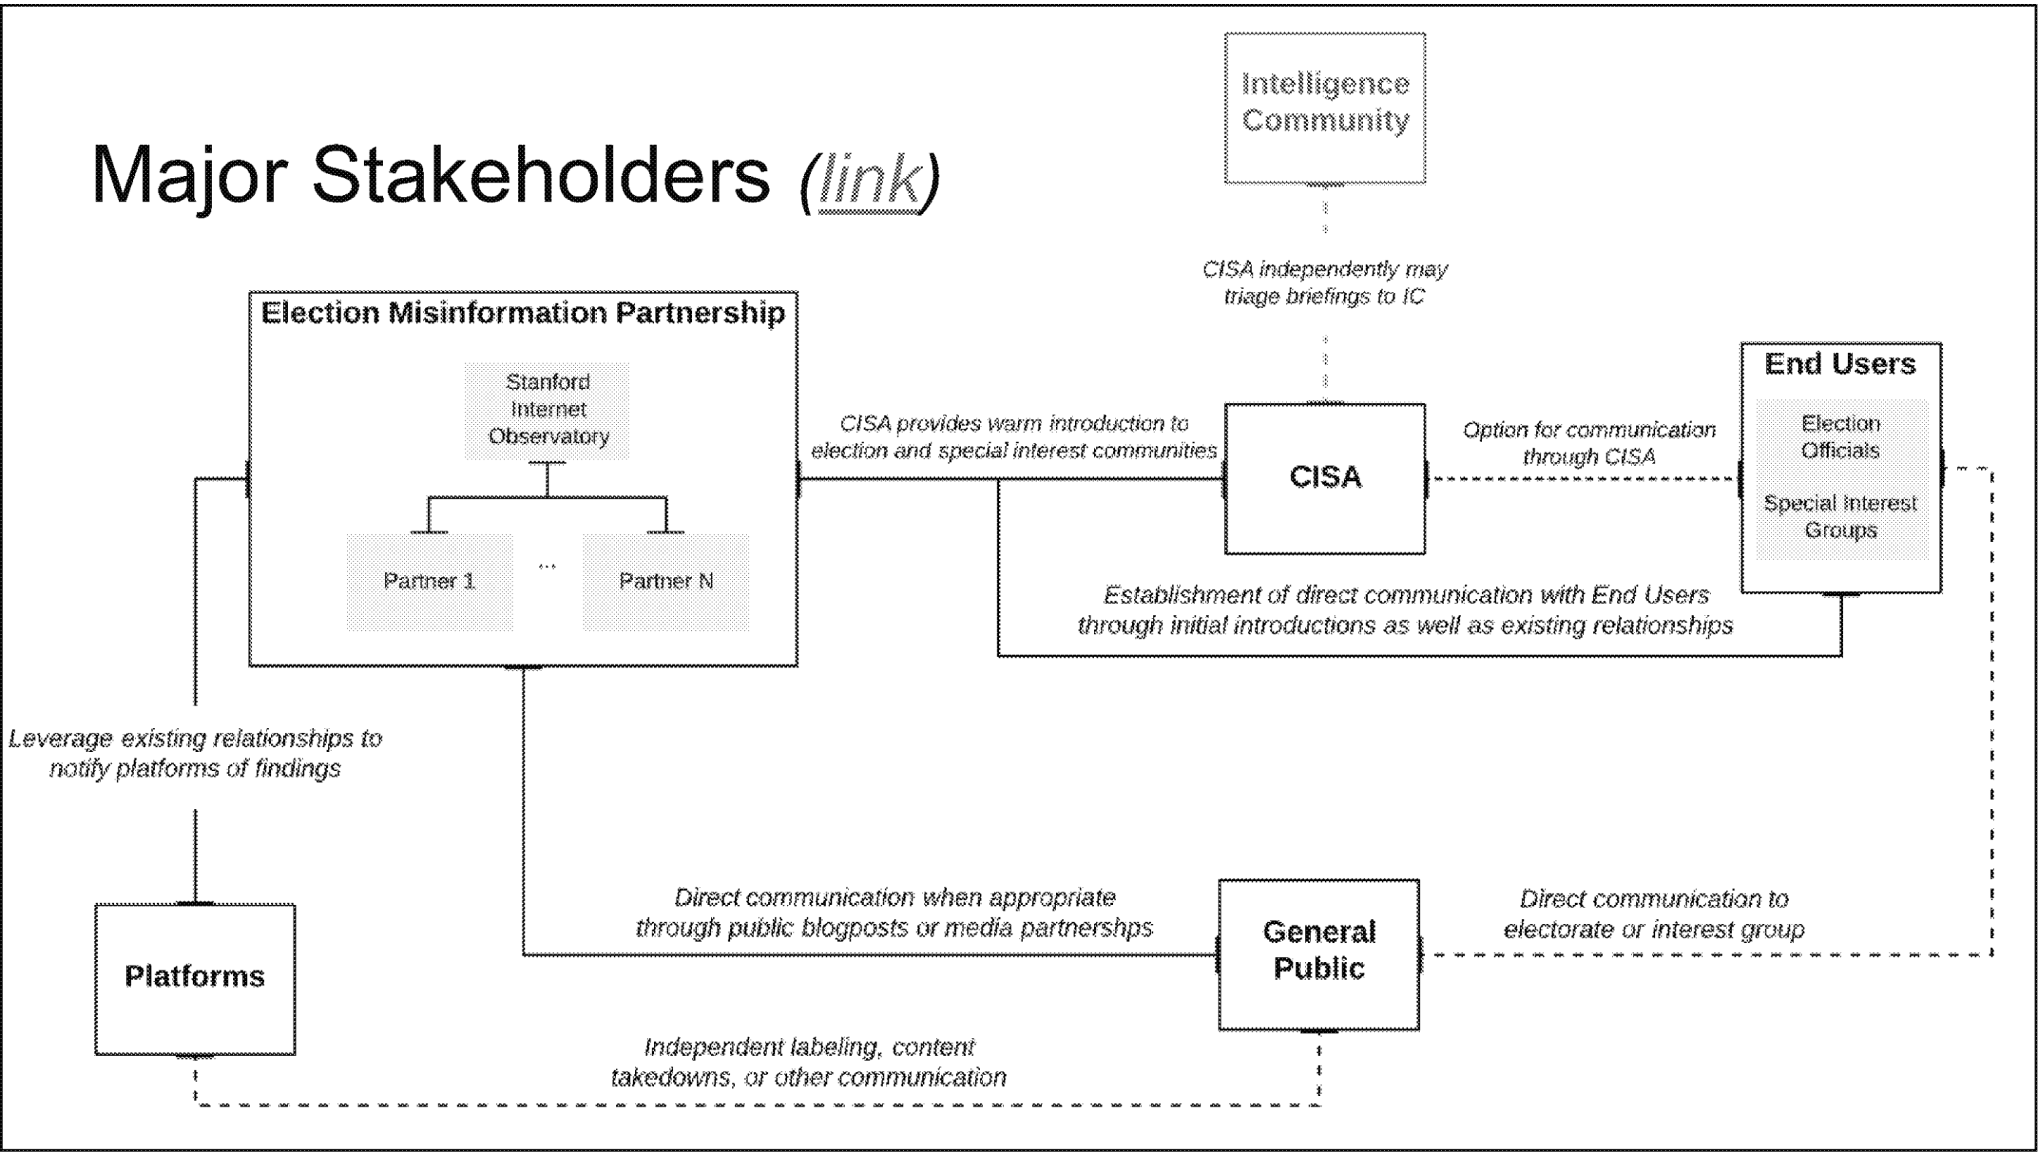

REL0000008097.0002

Qualifications for Involvement

What is meant by the ***third party research community***?

The main qualification is any academic, nonprofit or for profit entity dedicated to the monitoring and research of mis and disinformation in the social media space.

Such an entity must also be:

- US-based
  - EX: Sputnik's disinfo research team is not eligible
- A non-government entity
  - EX: the DOD is not eligible
- Nonpartisan, or not explicitly working for one political ideology
  - EX: Joe Biden's tech team is not eligible

Key Assumptions / Existential Hypotheses

	Assumption	Risk
1	**Mis and Disinformation on social media platforms is a problem which can negatively impact the integrity of the 2020 Election**	LOW RISK: if this was false, then CFI wouldn't exist.
2	One can **reasonably identify** this problem by monitoring the social media space	LOW RISK: this is proven by SIO's previous work
3	Once identified, a **trusted party exists** which can **reasonably mitigate** its impacts. (EX: counter narratives, tempering voter expectations, etc.)	MEDIUM RISK: it is unclear whether mitigation works, and if so, **who the right source of trust is**.
4	There are **many potential stakeholders here**. Some are better at **identifying** the problem, others are better at **mitigating** it. **These are not currently the same entity**. *EX: platforms can detect and mitigate, but not most trusted. Election officials or local news sources may be more trusted, but cannot detect.*	MEDIUM RISK: this is only true if there is some mitigation strategy, and the right trustworthy source has access to the electorate.
5	We [CISA + SIO] **can bridge this gap and improve election integrity by building trust and opening communication** between the third party misinformation research community and trusted entities that are better equipped to mitigate mis/disinformation during the Election, **IN THE ALLOTTED TIME** (~2 months) before the Election.	**HIGH RISK:** this means all the above assumptions are correct and Stanford and CISA are able to establish trust, communication lines, and effective processes in the alloted time frame.

Information Workstreams

What are the responsibilities of this coalition? Under what circumstances will it take action? What will be the actions?

Information Workstreams : Relevant Components

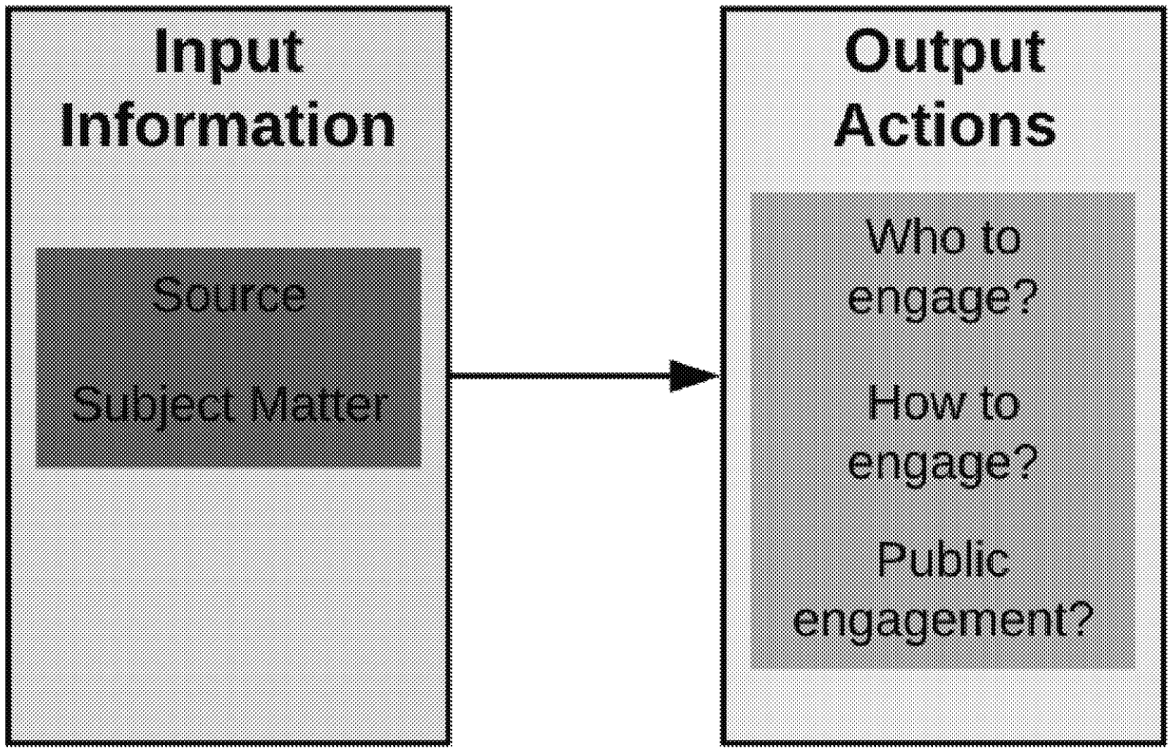

REL0000008097.0002

Input Information

Information Source

Where is the triggering information coming from? There are four cases here...

- **End Consumers**
  - Election officials
  - Special Interest Groups (EX: NAACP, VA, etc.)
- **CISA/CFITF**
- **Platforms**
  - FB, Twitter, TikTok, others?
- **War Room**'s own research

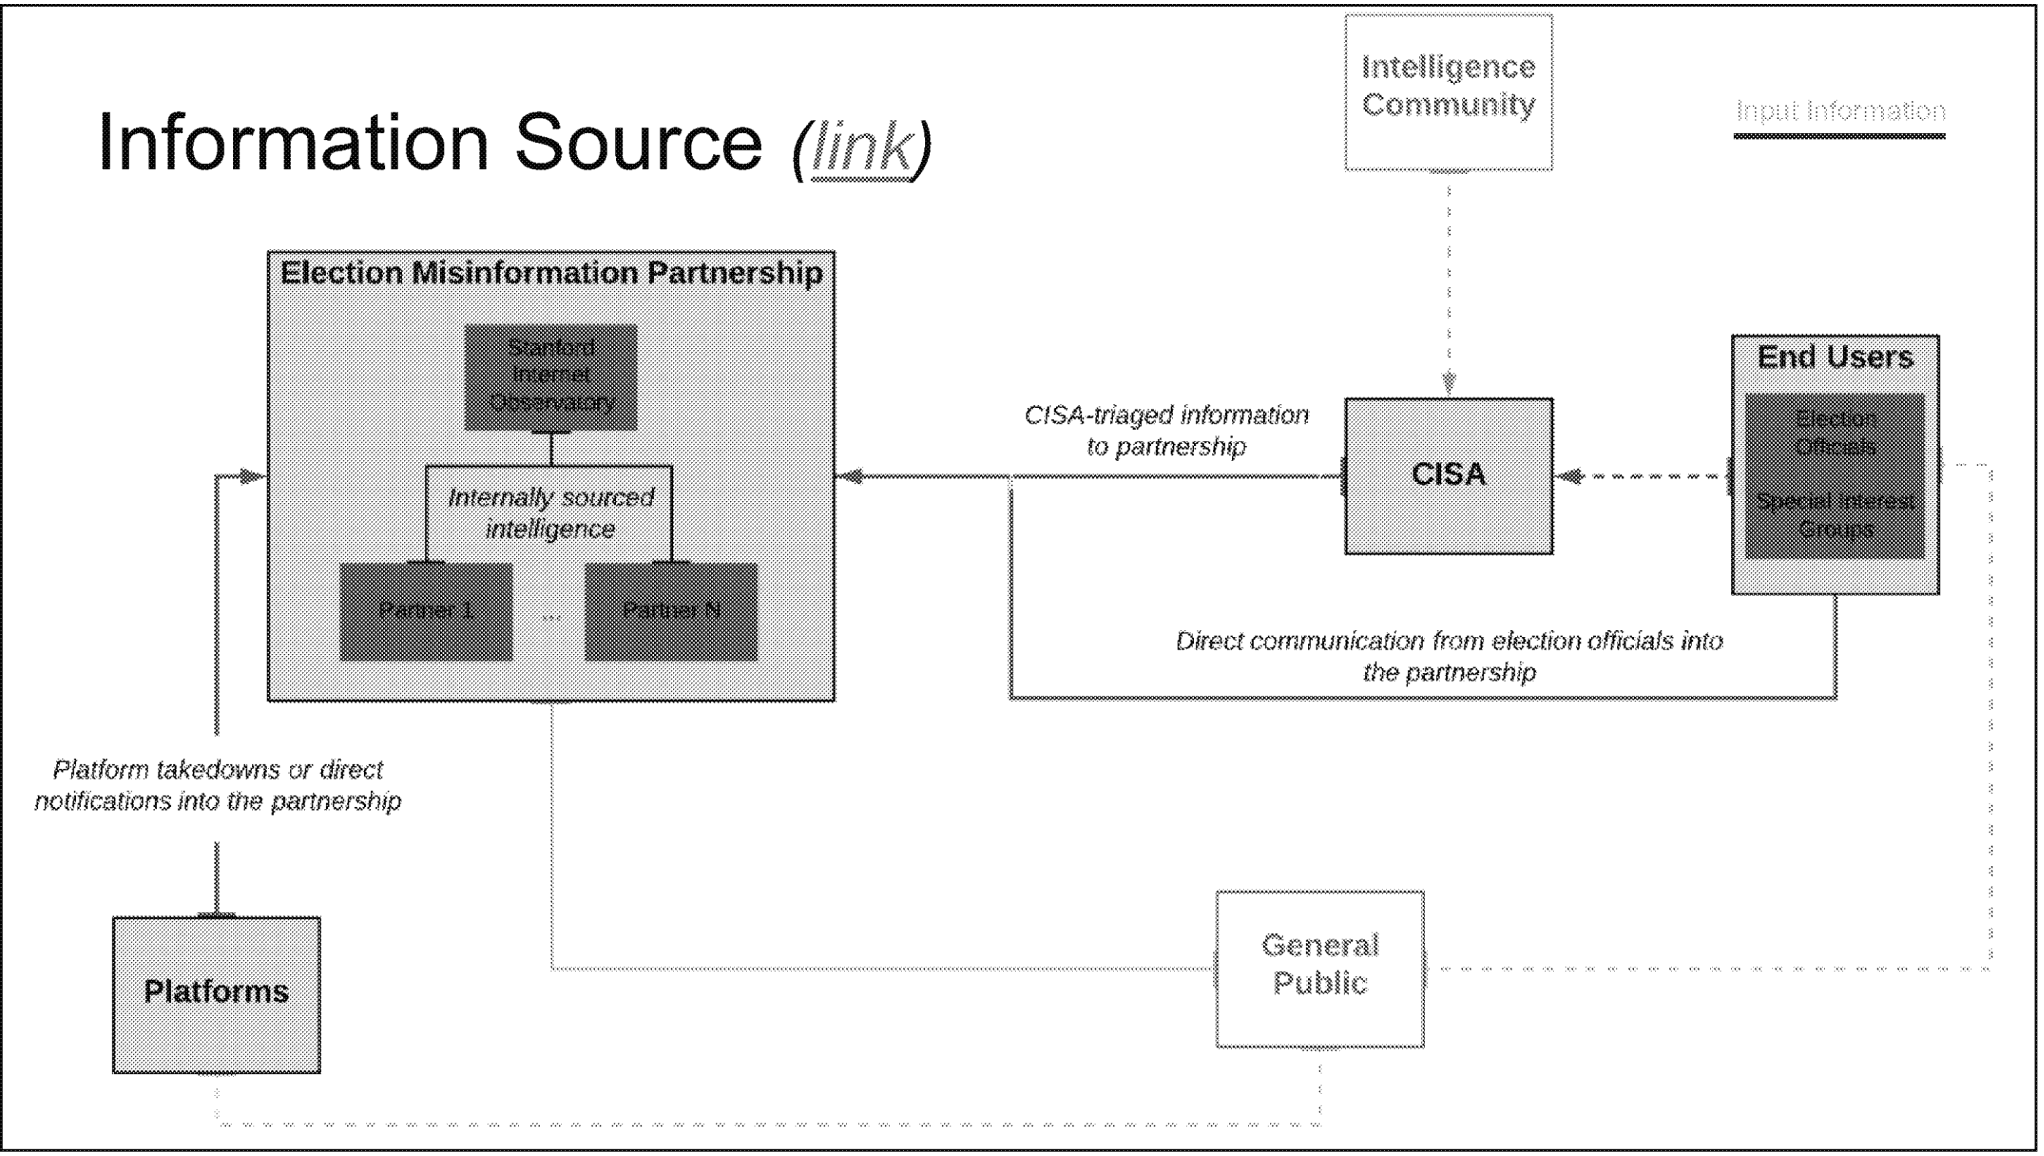

REL0000008097.0002

Information Prioritization and Scope

What is the subject matter priority?

- **HIGH**
  - Incorrect information that will **likely to certainly** change an individual's likelihood to go vote
  - Potential to instigate violence
- MEDIUM
  - Incorrect information that **may or may not** change an individual's likelihood to go vote
  - Polling information after election day
- **LOW**
  - Incorrect information that **will not likely** change an individual's likelihood to go vote
- **OUT OF SCOPE**
  - Generic political speech

Additional network risks or scale?

- Engagement
  - **HIGH** / MEDIUM / **LOW**
- Virality
  - **HIGH** / MEDIUM / **LOW**
- Coordinated Behavior or Inauthentic Activity
  - **YES** / UNCERTAIN / **NO**

Output Actions

Who to engage?

Of all the War Room stakeholders, who should be engaged with in an Ouput action from the EMP?

- **End Consumers**
  - Election officials
  - Special Interest Groups (EX: NAACP, VA, etc.)
- **CISA/CFITF**
- **Platforms**
- **General Public**
  - Independent blog posts, Tweets, etc.
  - Media partnerships

Who to engage? *(link)*

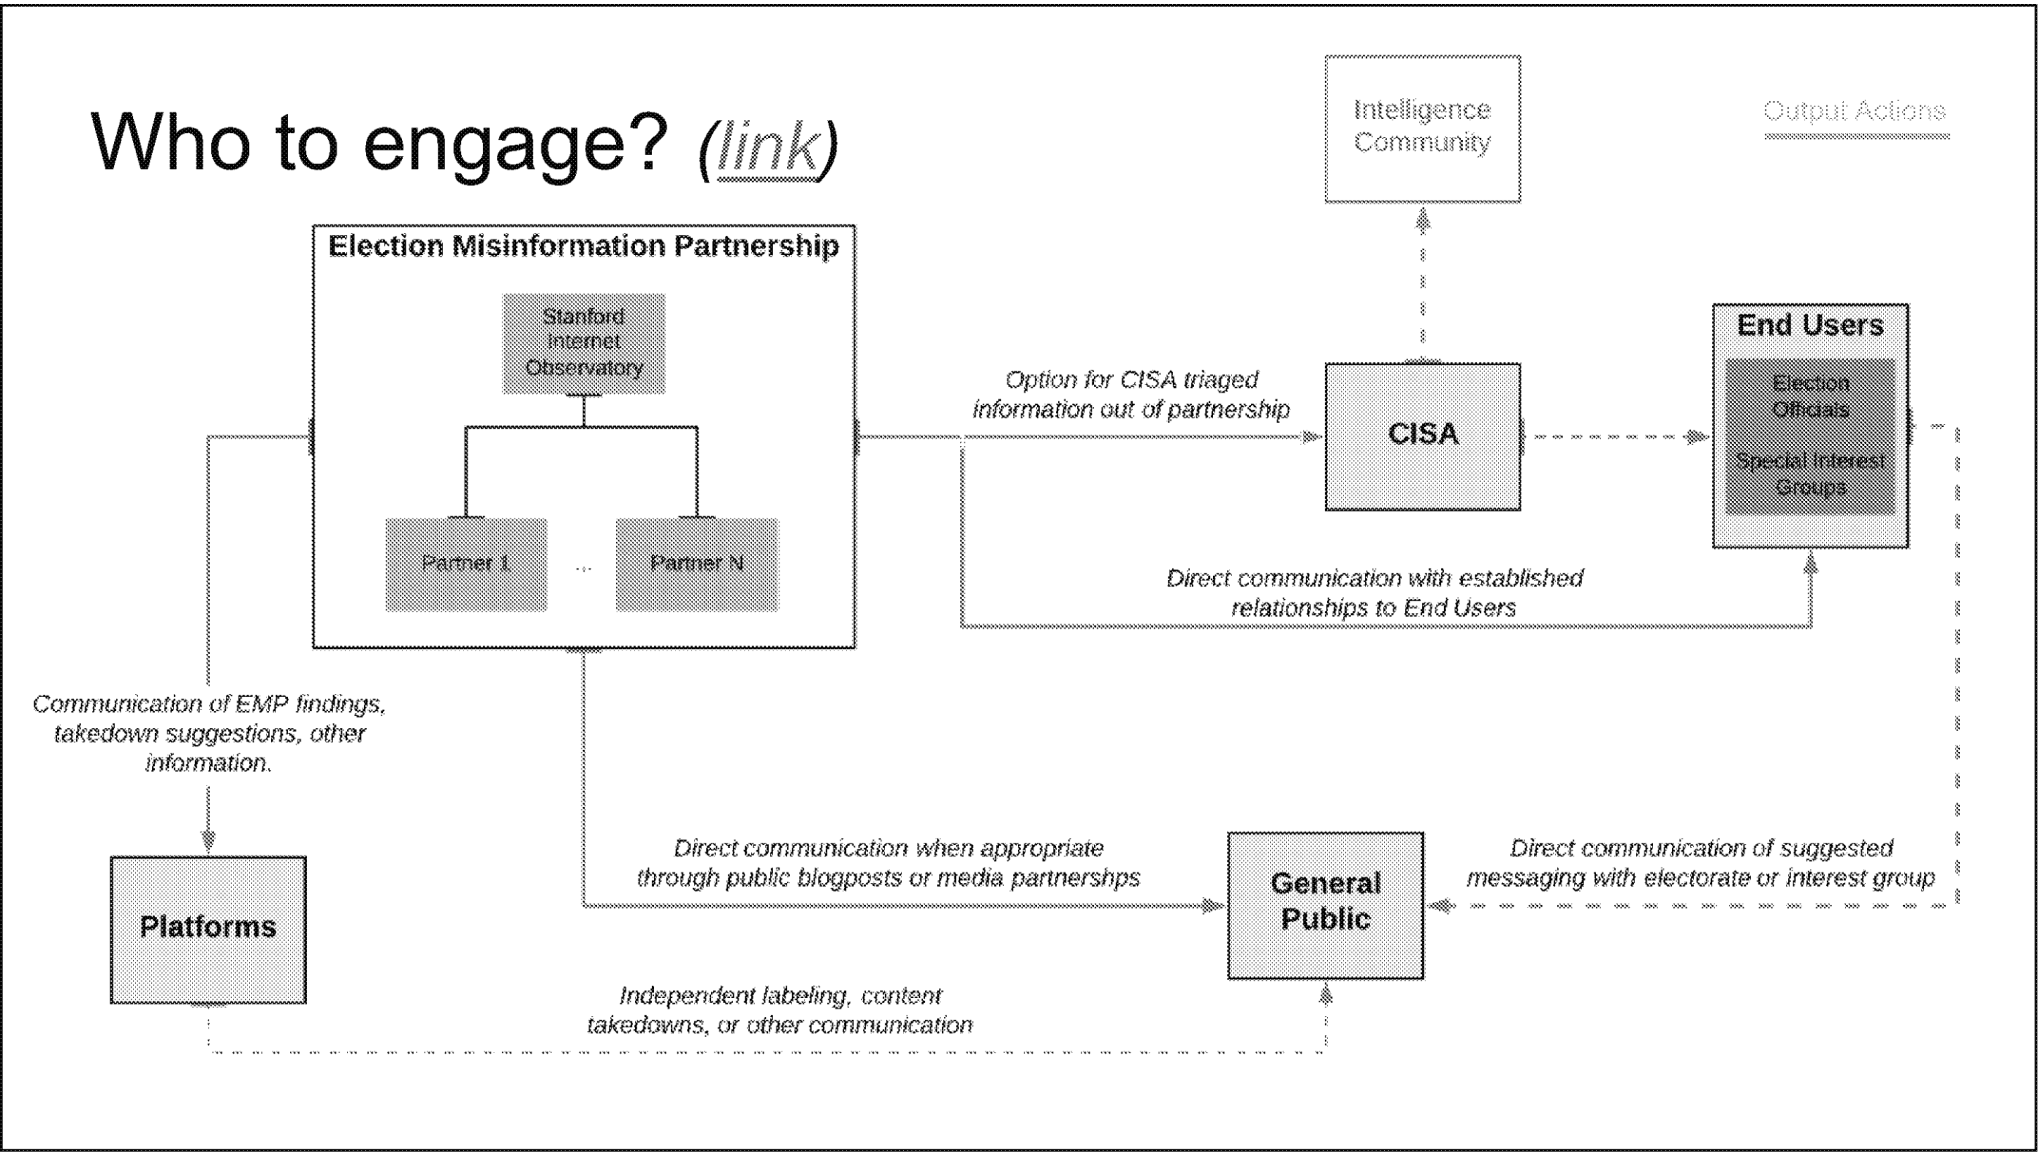

REL0000008097.0002

How to engage?

What level of engagement will the input conditions call for?

- **LOW** engagement
  - ~250 word summary of findings and suggestions for counter narrative or response
- MEDIUM engagement:
  - 1 page memo of findings and suggestions for counter narrative or response
- **HIGH** engagement
  - MEDIUM engagement + continued engagement, team formation to resolve the input
- **EXTREME** engagement
  - HIGH engagement + potential for Zoom call with multiple stakeholders, other escalations

Overall DecisionTable

Are there other variables to consider? What are the most ‘representative’ case studies?

Example Flows

Example Flow 1 : *Bread and Butter*

EMP researchers have identified a network of Pakistani pages, all named some variation of 'USA Live'. These pages are uploading recycled FB Live videos showing extremely long lines at NY polling locations. They are within the top 10 results when the hashtag #NYVote is searched, and have between 2K-100K views each.

Input Information			Output Actions		
Source	**Subject Matter**		**Who to engage?**	**Engagement Level?**	**Public engagement?**
Election Misinformation Partnership	*Subject Matter Priority?*	**HIGH**	Election Officials in the affected area, IC, Platforms	**HIGH** : dedicated team to continue engagement	**Yes** - public blog post
	Impressions?	**HIGH**			
	Virality?	**HIGH**			
	Coordinated/Inauthentic ?	**YES**			

Notes: this case is essentially just augmenting the current audience of third party research community with a direct pipeline to election officials and IC in real time. Think regular SIO activity, but one step beyond publishing on Twitter.

REL0000008097.0002

Example Flow 2 : *Dummy / No Engagement*

Candidate A tweets misinformation about Candidate B on their personal account. An election official reaches out to the War Room, concerned this is misinformation affecting the voting decision of their electorate. The post has gone viral on Twitter, but there is no immediate evidence of inauthentic activity.

Input Information			Output Actions		
Source	**Subject Matter**		**Who to engage?**	**Engagement Level?**	**Public engagement?**
Election Official	*Subject Matter Priority?*	**OUT OF SCOPE**	Election Official	LOW to **NONE**	**No**
	Impressions?	**HIGH**			
	Virality?	**HIGH**			
	Coordinated/Inauthentic?	NO			

Notes: no engagement with general political speech, especially when there is no evidence of coordinated or inauthentic activity. Only potential response here is to respond to the direct request with this policy.

Example Flow 3 : *Stickier*

#BidenStoleMichigan is trending on Twitter on election day. Groups of seemingly-local accounts tweet @MISecOfState to demand the Michigan election results be declared invalid, citing a fresh Epoch Times article alleging shady connections between Michigan's SoS, Bill Gates, and Joe Biden. Ther tweets are relatively few, but see high engagement shortly after posting and spread around right-leaning Twitter. Researchers trace the origin of the article to posts on 4chan and Parler encouraging Michiganders to confront @MISecOfStatea on Twitter over the story and calling for the Michigan results to be declared invalid.

Input Information			Output Actions		
Source	**Subject Matter**		**Who to engage?**	**Engagement Level?**	**Public engagement?**
Election Partnership (SIO)	*Subject Matter Priority?*	MEDIUM	Election Official, IC, Platforms	**LOW**, perhaps suggestions for containment and counter messaging	None until verification
	Impressions?	MEDIUM			
	Virality?	**HIGH**			
	Coordinated/Inauthentic?	UNCERTAIN			

Notes: This scenario has a geographical component, but seems targeted to ideological groups online. While particular election officials are targeted, the political nature of the content makes counter-messaging difficult. A government-only response would be even stickier however.

REL0000008097.0002

Example Flow 4 : *Sourced from Local Official*

A local election official notifies EMP researchers that voters have been calling their county's elections hotline to enquire if the election has really been canceled, citing a news story that is clearly fake that quotes a county official canceling the election due to an "unprecedented COVID-19 surge" in the state. Preliminary research shows just a few Facebook pages and Twitter accounts tweeting the story, but the origin of the article is unknown.

Input Information			Output Actions		
Source	**Subject Matter**		**Who to engage?**	**Engagement Level?**	**Public engagement?**
Local Election Official	*Subject Matter Priority?*	**HIGH**	Election Official, IC, Platforms	**HIGH,** dedicated team to research origin of article, track its spread, develop counter-messaging, coordinate with platforms if needed	None until verification
	Impressions?	MEDIUM			
	Virality?	MEDIUM			
	Coordinated/Inauthentic?	UNCERTAIN			

Notes:

REL0000008097.0002

Example Flow 5 : *Sourced from Platform*

Days after 11/03, Facebook notifies EMP of an impending takedown of a group of pages exhibiting coordinated inauthentic behavior. Since the election, these pages have consistently pushed a narrative encouraging Americans in key states to call for invalidation of election results. Facebook will take these pages down in one hour, and is already briefing relevant state and local election officials.

Input Information			Output Actions		
Source	**Subject Matter**		**Who to engage?**	**Engagement Level?**	**Public engagement?**
Platform	*Subject Matter Priority?*	**HIGH**	Election Officials	LOW, given platform involvement. Higher if requested.	None, pending follow-up analysis, potential collaboration with platform
	Impressions?	MEDIUM			
	Virality?	MEDIUM			
	Coordinated/Inauthentic?	YES			

Notes: Given that information is platform-verified, and Facebook has a direct relationship with local election officials, EMP's involvement can be smaller with the initial dump. EMP should follow up with election officials and the platform in case either stakeholder wants for further research.

Open Questions

Collaboration Surfaces

What will be the actual surface which SIO, CISA, and trusted 'End Users' such as election officials will communicate and collaborate on?

Some potential options…

- **Slack**: channels, direct messages
- **Jira**: ticketing workflow
- **Email**: listserv, direct emails
- **Adobe Connect**: leverage EI-ISAAC touchpoints

‘Product Market Fit’ : who is the right End User?

One Key Assumption is that Election Officials, special interest group leaders, or other CFI partners are 1) in a position to take actions that can mitigate the impact of misinformation in realtime and 2) willing to take such actions.

Open Questions

- Who exactly do we mean when we say ‘Election Officials’? Who is the best contact here?
- Can we find a representative subset of this group and run some sort of initial test with them? (EX: during a primary?)
- Are there other interest groups we haven’t thought of that could be better End Users? (EX: from the Director Krebs meeting, coalition of local reporters?)

Further Open Questions

- Customer research
  - What is the frequency of mis/disinfo events which would benefit from such an entity?
  - How to establish credibility? Communication pipeline only through CISA or open directly?
  - How best to interface with state officials? Ask each state to identify someone?
- Response Thresholds
  - What research should be publicly released, and when? How is that decided?
  - What level of certainty do researchers need to escalate response?
- War room capabilities
  - What are the exact entities who would like to be involved?
  - What tooling exists now? Does anything need to be further developed/acquired?
- Platform relationship: will this be driven by war room, or by CISA? Under what circumstances?
- Others?

Operational Details

What might the actual coalition look like? What are the Timelines and MVP? Key delivery dates?

[This Section is Under Construction]

Stanford Internet Observatory Calendar

July	August	September	October
Organizational Setup • Establish responsibilities and expectations with CISA collaboration • Explore partnership options within third party research community • Define investigatory toolset • Call for students • Acquire funding	**Onboarding** • Train students in open source investigation. Begin staffing on a part-time basis • Election official buy-in, relationship development • MVP of Just SIO capabilities by Republican Convention: **August 17th**	**Part Time Intake** • Begin intake and monitoring with student staff • Initial product lines for election officials • Begin to onboard outside research teams into full flow	**Full Time, Pre-Election** • Coalition solidified, public facing • Students working part time, nearing full time hours • Beginning of 'pre-election period' monitoring

REL0000008097.0002

Stanford Internet Observatory Calendar

November	December	January
Full Time, Election • 24/7 monitoring in shifts • Heightened monitoring during voting times • Emphasis on voter suppression tactics • Election **November 3, 2020**	**Full Time, Post-Election** • Full time monitoring continues, but not 24/7 • Emphasis on narratives around election legitimacy (EX: mail in ballot theories) • Release brief post-mortem	**Part Time, Inauguration** • Part time, wrapping up coalition, outtakes. • Release full report. • Inauguration **January 20, 2021**

Research and Hiring Needs

Headcount will be determined by the geographic and topical breakdown we determine is necessary to get a full picture through the extended Election Period.

Example Breakdown: 40-person research team

- **5 Research Leads:** experienced SIO members who oversee work done by Geographic and Topical researchers, write the final blog posts, etc.
- **20 Geographic Coverage Researchers:** each will be assigned one or more states, for which they will determine key electoral districts, voting dates, search terms.
- **10 Topical Researchers:** each will be assigned one topical area (EX: QAnon, Russian sponsored media, others) which they will watch in more depth.
- **5 Tech/Data Support**: floaters to assist in data visualization or real time scripts which might be needed

Appendix / Graveyard

Customer Research

- What is the frequency of mis/disinfo events which would benefit from such an entity?
- How to establish credibility? Communication pipeline only through CISA or open directly?
- How best to interface with state officials? Ask each state to identify someone?
- What role do interest groups best play? Counter-messaging?

War Room Capabilities

- What groups will be involved?
- What will the role of SIO be? Explicit partnership with SIO, nonprofit around multiple research entities, other?
- Enumeration of all resources
  - General 'manpower' (# of employees, hrs/week, language coverage, % tech vs. RA, etc.)
  - Technical tooling alongside platform coverage
  - Funding constraints

All Inputs / 'Clearinghouse Process'

All reasonable* inputs will receive

- 15 min research time from team of X researchers
- Classification by subject matter risk/priority, network risk/priority
- 120 character response

A reasonable input will be defined by:

- //need to define//

Summarized notes (Thank you to (b)(6))

Overview: CISA has limited capabilities to identify, track disinfo narratives + attempts to undermine confidence in elections

- SIO does = good partnership
- Major goal: prevent a crisis of confidence in 2020 elections
    - E.g., where Russia doesn't change any votes (or changes just a few), but claims they changed many more and hysteria is blown out of proportion

Scope: Keep scope narrow: focus on election-related disinfo that has the potential to impact the public's voting patterns

Partnerships and Relationships: SIO would be the coordinators, working w/ Graphika, DFRLab, and (b)(6) (b)(6) s team at UW

- Mutual trust is key: don't want to need NDAs, legal red tape
- Need to build out workflow management system: JIRA/Slack/other communications channels, shared processes and definitions, etc.
    - (b)(6) envisions Tier 1 and Tier 2 partners
    - Tier 1 is intake (of tips, disinfo reports, etc.): consisting of people either digging for narratives, or processing info received from other partners
        - Think students, election officials, etc. who are looking for disinfo
        - Workflow: check that info against protocols, do some initial data aggregation, triage it into the workflow management system
    - Tier 2 is the 4 orgs: (b)(6) team at Stanford, Graphika, etc.
        - Workflow: take stuff off the workflow management queue, process it
        - Need to sketch that out
        - SLA for different times of the calendar based on the level of severity obtained by triage
            - E.g., a report from the general public will have less priority than a report from an on-the-ground election official; a report for disinfo that is not popular will have less priority than disinfo that is going viral
            - General public = more turnaround time, but election officials = less turnaround time: need to get back to them fast
- SIO has good relationship w/ platforms who already care
    - See the Secondary Infektion (Russian disinfo op) report
        - Think through all the platforms that might have been useful there (e.g., communicating with Twitter at stage *x* would have stopped the spread)
- Meanwhile, CISA has strong relationships w/ election officials
    - CISA is happy to introduce SIO to them, do outreach
    - Just keep CISA in the info-sharing pipeline
- SIO normally won't share with intelligence community (IC) or law enforcement
    - No gov't agencies other than CISA
    - CISA is free to converse, triage with IC

    - SIO won't publish or share private info: goal is to find disinfo, not expose suspects behind it
- Regarding the general public: create a workflow to 1) understand, 2) mitigate, 3) communicate
    - Understand a narrative and its theme; ensure it's relevant; identify accounts behind it
    - Then, mitigate it: send details to the platforms and CISA
    - Finally, hopefully it's mitigated by the time we communicate the details to public
    - Need a trustworthy place for general public to go
        - Could be social media, a live blog, or something more complex
        - It'll probably be a live blog with social media: ~2 paragraph statements of what we've found, the accounts affected, note we've forwarded it to gov't + platforms
        - Goal: all organizations involved approve a central statement saying *x* narrative/piece of disinfo is not true and why

Timeline from here: SIO publicly announces the partnership in 2-3 weeks

- Public post about the partnership
- Hold a webinar to talk through everything with election officials + stakeholders
    - Tell them how they can submit tips, give feedback, stay in sync

CISA's concern starts 45 days out operationally, when military/overseas voters start mailing

- Start hunting, messaging at beginning of September
    - Lower SLA (higher turnaround time/less priority), but start looking for search terms and taking tips
- The days leading up to/right after Election Day will be much more intense
- It'll be an effective SOC, maybe a physical one, but in a much larger space

Example scenarios

What if we find a disinfo op in a state SIO hasn't yet built a relationship with?

- CISA can make the introduction, preferably through the ISAC
    - ISAC has a reporting structure where election officials can report tips to platforms
    - CISA can provide SIO those channels
- SIO will publicly disclose the op: it serves as counterspeech, and it reaches relevant officials outside the info-sharing pipeline
    - CISA can decide if it looks foreign, whether to forward it on to IC/other parties

What if official sources post disinfo (e.g. candidate's Twitter)?

- That scenario is out of scope, and even if it's misinfo, it'll still be widely received
- SIO *could* add value if a candidate tweets misinfo about the electoral process
    - Need to discuss: how much do we want to unanimously say this is untrue? Emphasize respected election lawyers' voices? This could get political fast

How would we notify political parties (along with their own stakeholder/notification groups)?

- We'd do it in a completely fair manner and offer services to both of them

- Like above, still need to consider SLA challenges
    - We'll take a tip from any reasonable group, but different prioritizations
    - E.g. if the tip comes from an election official, we need to get back quickly, effectively, transparently
    - If it comes from a political group, must be extremely careful and don't disclose anything not intended for public disclosure

Best way to collaborate

What's the best way to collaborate?

- CISA can't create their own Slack channels, but *can* participate in others'
    - Listservs are bad (public records requirements)
- Jira is fine
- CISA has privacy concerns: can't monitor people's individual accounts; ensure CISA doesn't participate in discussions or notes concerning U.S. persons
- Setup:
    - SIO will have dedicated Slack, something like Jira or Salesforce (will ask for donation), separate from Stanford and destroyed once over
    - We'll intake info by email, but direct people to private forms SIO and CISA have distributed
    - Info from there will go into queue -> be triaged, assigned SLA

Next steps and points of contact

EI-ISAC: get ahold of **(b)(6)** he's leading their social media reporting project

Brian (CISA): discuss process to integrate into CISA's ops center and send tips back to SIO

(b)(6) will remain liaison b/t CISA and SIO

All reporting to CISA will be fully unclassified

Calendar from now till finish line:

- Next few weeks are critical
    - Onboard partners w/ next week's coordination call
    - Hire students
    - Make a public announcement, but keep CISA separate for now
- After the public announcement:
    - Call for hire, funding, etc.
    - CISA can make direct introductions, or forward the SIO webinar link to contacts so people can join at will
        - One webinar for election officials, one for NGOs + other stakeholders
- Define our intakes, communicate endpoints
    - By August, hold meetings to start discussing searches, alerts, etc.
    - Starting point: rank cities on the Cook political report, combine city and county names, etc. -> these will go to the engineers who will be doing the physical monitoring

- Before and after the election is most intense, dangerous
  - Up 18 hours a day
  - Most activity will be through November, into December
- Once it's all over, wrap with report, collection of data/findings, final thoughts
  - Group celebration

CISA will intro this to ISAC people

- (b)(6) to give short presentation