

PROJECT EVIDENCE SYNTHESIS

Democratic Resilience in Ukraine

Expanded Evidence and Measurement Logic from the Build to Disrupt Project

Evidence period August 2025 to March 2026

Executive summary

Build to Disrupt worked on a practical component of Ukraine's democratic resilience: whether public institutions can identify AI-enabled cyber threats and foreign information manipulation and interference (FIMI), verify suspicious content, and use more systematic ways to respond. The eight-month project combined policy research, national and sub-national training, practical exercises, methodological recommendations, and pilot access to a Ukrainian information-monitoring tool.

The evidence shows substantial immediate changes in participants' reported knowledge and confidence, early use of the methods in routine work, and formal uptake of project-supported expertise in national cybersecurity guidance. In the online cohort, the share reporting high or very high knowledge increased from 13.5% to 52.3%; in the in-person cohort it increased from 29.8% to 75.3%. The share reporting confidence at the upper end of the scale increased from 15.0% to 54.5% online and from 33.0% to 82.0% in person.

The most concrete institutional evidence includes SSSCIP's approval of AI cybersecurity recommendations under Order No. 154; the Centre for Countering Disinformation's use of Let's Data alerts and reports in daily analytical products; and reported changes to the Ministry of Culture's disinformation-strategy work and the Ukrainian Startup Fund's AI Sandbox assessment. Ten tangible examples of positive change were documented by the project, although the source files do not provide a complete institution-by-institution denominator.

These are strong early contributions to institutional capability and practice. They are not direct measures of public trust, election integrity, citizens' exposure to accurate information, or national response speed. The project's own contribution assessment rated the evidence Level 4 of 5, meaning strong and triangulated evidence of plausible contribution with longer-term effects still to be demonstrated.

Key figures at a glance

- 729 participants in the two core cohorts: 635 regional and local officials trained online and 94 central-government representatives trained in person; a separate 14-person women's cohort is reported in addition.
- 58% women in the core cohort aggregate as recorded in the project's detailed training table; women were 62% of online participants and 30.9% of central-government participants.
- Knowledge at the high or very high level rose by 38.8 percentage points online and 45.5 points in person.
- Confident or very confident responses rose by 39.5 percentage points online and 49 points in person.
- 56 participants registered for Let's Data onboarding; the reported average alert open rate was 66%; 10 tangible practice-change examples were collected.
- 96.2% of regional and local respondents and 97.8% of central-government respondents requested further training.

Project scope and democratic resilience frame

The project's long-term impact statement set an ambition for Ukrainian institutions to detect, counter, and disrupt AI-driven FIMI and cyber threats in real time, while reinforcing public trust, protecting democratic processes, and aligning information and cybersecurity governance with EU and NATO standards. The report below assesses progress toward that ambition during the reporting period; it does not treat the ambition itself as an achieved result.

For this synthesis, democratic resilience means the capacity of institutions to keep public communication and decision-making secure, informed, and functional during manipulation or disruption, while responding transparently and respecting rights. This definition makes the chain of evidence explicit:

Research, training, guidance and tool access → stronger capability, opportunity and motivation → changed verification and response practices → more resilient public institutions and information flows

The project used a COM-B lens to assess early shifts in capability, opportunity, motivation, and behaviour. It also used a five-level contribution rubric and triangulated survey data, participant feedback, tool analytics, institutional communication, and documentary evidence. This is useful for assessing plausible contribution in a short programme, but it does not establish sole attribution.

Delivery and participant reach

The intervention reached people responsible for public communications, cybersecurity, strategic communications, information security, and policy at central, regional, and local levels. The online format created scale during a period of attacks and blackouts; the in-person format concentrated on national institutions with operational responsibility.

Cohort	Reach and composition	Training content and relevance
Regional and local online cohort 27–28 January 2026	635 participants. 62% women (394) and 38% men (241). Participants included regional administrations (48.1%), local councils (35.7%), and district authorities (5.8%). Leadership roles represented 44.4%; specialists and administrators 49.4%. About two-thirds worked in official communications and around one-third in IT or cybersecurity.	Two-day training on AI-enabled threats, FIMI, public-sector cybersecurity, verification, and response. Participants came from frontline regions including Luhansk, Kherson, Donetsk, and Sumy. Local and regional institutions are citizen-facing service providers and sources of trusted public information.
Central-government in-person cohort 10–13 February 2026	94 participants. 29 women (30.9%) and 65 men (69.1%). Functional roles: cybersecurity and technical specialists 31%; leadership and policy 25%; strategic communications and media 24%; academia and research 11%; other support and coordination 9%.	Reported as a three-day training in Kyiv, with experts from technology, academia, media monitoring, and cybersecurity. Sessions included cybersecurity fundamentals, synthetic voice, OSINT, and bot-network detection. Tools demonstrated included Pulse Media Scout, Talkwalker, Hootsuite, and TelepathyDB.
Additional early-career cohort	14 young women, all female, trained in partnership with UNDP Ukraine. The group included young professionals and interns aspiring to enter public service.	This supplementary activity was added in response to an emerging opportunity. It widened access to AI and information-security learning for women at an early-career stage; its reach is kept separate from the two core cohorts.

The project report describes near-national reach and notes lower participation from some western regions, which it interprets as a possible difference in perceived risk. That interpretation should be tested with outreach and demand data rather than treated as established explanation.

Delivery also adapted to severe winter blackouts. Online sessions were recorded and shared with materials, while the Kyiv venue was selected for backup power. The risk register had anticipated that electricity and connectivity failures could prevent attendance. The later report records that both online and in-person activities proceeded. This demonstrates continuity of the project's delivery; it does not by itself demonstrate continuity of public services during attacks.

Measurement logic

The table links each intervention to an observed result and then to its relevance for democracy and resilience. It separates outputs, self-reported outcomes, and early institutional effects.

Support or intervention	Measurable outcome and evidence	Democratic-resilience relevance and evidence status
Policy research and co-developed guidance	Two named products were shared with 14 governmental institutions: a policy paper on AI-enabled FIMI and cyber threats in the public sector, and methodological recommendations on cybersecurity for AI-enabled information and communication systems. Fourteen regional governmental institutions said the policy paper was useful. SSSCIP acknowledged IIG's contribution and approved the recommendations by Order No. 154 on 23 February 2026.	Moves expertise from a project document into national methodological guidance. This can support consistent public-sector practice. The evidence shows approval and reported usefulness, not uniform adoption or legal compulsion.
Online and in-person training	High or very high knowledge increased by 38.8 points online and 45.5 points in person. Confident or very confident responses increased by 39.5 and 49 points. Post-training response rates were 83.7% online and 94.6% in person.	Officials reported stronger ability to recognise manipulation and cyber risks. This supports safer institutional communications and decision-making. Measures are self-assessments, not independently scored tasks.
Practical sessions, OSINT methods, and tool onboarding	The MEL matrix records three Let's Data sessions against a target of three simulation exercises. Fifty-six participants registered for tool access and alerts; average alert open rate was 66%.	Provides a bridge from awareness to practice. Open rates measure engagement, not accuracy, response time, or the quality of decisions. The source files give limited detail on exercise scenarios and scoring.

Support or intervention	Measurable outcome and evidence	Democratic-resilience relevance and evidence status
Follow-up and workplace application	Among 44 online follow-up respondents, 32% reported using learning for information verification or threat prevention; 9% used it in incident response; 32% partly applied tools or approaches; 36% had not yet had an opportunity. Ten tangible change examples were documented across the project.	Shows early behaviour change, including phishing prevention and monitoring. The follow-up sample is small and self-selected, so it cannot establish how common application was among all trainees.
Adaptive, sub-national delivery	The online cohort reached 635 regional and local officials, including frontline areas; 62% were women. Recorded sessions and offline materials supported access during outages. Further training was requested by 96.2% of regional/local and 97.8% of central-government respondents.	Broadens the institutional base beyond Kyiv and helps sustain local capacity for citizen-facing services and public communication. Participation is evidence of access and demand, not proof of improved citizen trust.
UK–Ukraine exchange pathway	The project planned two knowledge-exchange events and seven participating institutions in a consortium or hub. The activity was not delivered in the reporting period; the report records no evidence for this outcome.	This pathway cannot yet support a resilience claim. SSSCIP identified a possible future agenda on electronic trust services, qualified e-signatures, e-delivery, and resilience to deepfakes and synthetic forgeries.

Learning and confidence results

The training surveys measured perceived knowledge on a four-point scale (low, average, high, very high) and confidence on a four-point scale (not confident, moderately confident, confident, very confident). The project compared the distribution of responses before and after training. Because the survey data were aggregated rather than matched person by person, the changes below describe cohort-level shifts, not the exact percentage of individuals who improved.

Cohort and coverage	High or very high knowledge	Confident or very confident	Baseline exposure and response
Online regional and local officials 635 pre-surveys; 532 post-surveys	13.5% before → 52.3% after +38.8 percentage points	15.0% before → 54.5% after +39.5 percentage points	Post-survey response rate 83.7%. Before training, 50% reported prior professional exposure to AI or FIMI threats.
In-person central-government officials 94 participants; 89 post-surveys	29.8% before → 75.3% after +45.5 percentage points	33.0% before → 82.0% after +49.0 percentage points	Post-survey response rate 94.6%. Before training, 88.3% reported encountering AI-powered cyber threats or disinformation at work. Very confident responses rose from 2.1% to 24.7%.

The online knowledge result is a near fourfold increase in the share selecting high or very high; it should not be described as a fourfold increase in tested competence. For the in-person group, the high/very high share rose despite substantial pre-existing exposure to threats, suggesting the course added structure and confidence to experience participants already had.

The project report records three out of four immediate outcomes as achieved: usefulness and awareness of policy products, capability, and confidence. The fourth outcome—access to established cross-border networking channels—was not achieved because the UK–Ukraine exchange did not take place.

Behaviour and institutional uptake

Online follow-up

A follow-up questionnaire was conducted a few weeks after the online training. Forty-four participants responded. The project recorded four types of self-reported response: 32% applied learning to verification or threat prevention; 9% used it in incident response; 32% partly applied tools or approaches; and 36% had not yet encountered an opportunity to apply the learning. The categories may overlap, and the sample is too small to estimate prevalence across the full cohort.

Tool use after in-person training

Following the central-government training, selected participants received onboarding and pilot access to Let's Data, a Ukrainian monitoring product. Fifty-six participants registered and received alerts on selected topics. The reported average open rate was 66%. Users particularly valued video transcription for TikTok and YouTube and asked for wider language

and geographic coverage. The open-rate data confirm interaction with alerts; effectiveness would require measures such as verified narratives, false positives, time to confirmation, and use in decisions.

The project also records that testing of commercial AI-detection products raised reliability and security concerns, including opaque ownership, data-sharing arrangements, and dependence on infrastructure that Ukrainian institutions may not trust during wartime. The project therefore prioritised a Ukrainian-developed tool in its practical curriculum. This is a project learning finding, not a market-wide independent audit.

Policy and operational uptake

The strongest formal policy evidence concerns SSSCIP. On 4 March 2026, an official communication from the acting director of its Cybersecurity Department stated that IIG comments and proposals had been taken into account in preparing the recommendations. SSSCIP approved the document under Order No. 154 dated 23 February 2026 as part of the national AI Development Concept implementation plan (Cabinet of Ministers Order No. 457-r, 9 May 2025).

The policy paper on countering AI-enabled FIMI and cyber threats was shared in March 2026. Fourteen regional governmental institutions provided feedback that it was relevant and useful to their work. The project report cautions that more time is needed for internal review and adoption into strategies, operational guidance, and routine practice.

The project's intermediate-outcome assessment marked evidence-based policy integration and proactive detection/verification as achieved. It marked regular UK–Ukraine collaboration as not achieved. Across the whole project, ten tangible examples of positive changes were collected; these examples are not necessarily ten separate institutions.

Documented change stories

Phishing recognition in frontline and local authorities

A representative of Shchastia District State Administration, operating in exile because the area was occupied, reported recognising phishing signs in a suspicious email and avoiding opening its file. A participant from Velyka Oleksandrivka village council described checking the sender, language errors, and file type in an email with an archived attachment that could contain malware. A Vinnytsia City Council participant reported deleting a suspicious link without opening it. These examples connect training to protective decisions in real work, including a displaced administration.

Centre for Countering Disinformation uses monitoring alerts

Before the project, the Centre for Countering Disinformation (CCD) had not used Let's Data as part of routine FIMI monitoring. After the in-person training and pilot access, CCD staff used alerts and custom reports to identify narratives for daily analytical products. Daryna Harib described the tool as useful for monitoring regional dynamics and Ukrainian- and Russian-language content. CCD was exploring funding for continued access, so the example shows operational uptake but not yet sustained access.

National cybersecurity recommendations receive formal approval

SSSCIP's approval of the AI cybersecurity recommendations is the clearest documentary shift from expert input to official guidance. The evidence consists of the order approving the recommendations and the department's written acknowledgement that IIG's comments were taken into account. The practical effect will depend on whether public institutions translate the guidance into procedures, procurement checks, incident playbooks, and staff onboarding.

Training informs strategy and AI assessment

Oleksandr Makobrii, Deputy Director of the Department of Information Policy and Information Security at the Ministry of Culture and Information Policy, said he had begun work on restructuring the national Strategy for Countering Disinformation after the training. The report describes proposed attention to synthetic media, coordinated manipulation, early detection, and verification workflows; revisions remained in progress.

Iryna Terzi of the Ukrainian Startup Fund said the training helped her consider AI as both a product-development tool and a means of manipulation or cyberattack. She reported that the approaches could be integrated into assessment of AI and blockchain solutions in the government Sandbox, particularly for security, reliability, and information risks.

Specialists apply methods in their roles

Sofiia Petkova, an AI specialist at the Ministry of Digital Transformation, reported that the synthetic-voice session helped her understand how voice models differ in accent and delivery, which she used when tracking AI-generated video and digital communications. Yelyzaveta Rumiantseva of the Armed Forces' Main Operational Communications and Information Systems Centre reported using image-search and audio-related tools shown in the OSINT session. Nataliia Muzyka of

SSSCIP's Cyber Protection Department said the OSINT methods had become relevant to work with international partners and were already in active use.

These accounts are evidence of use reported by named stakeholders in project interviews or follow-up messages. They have not been independently audited in this synthesis.

Demand for further support

High requests for additional training indicate that participants saw unresolved operational needs. The next-step priorities also reveal what participants understood as the practical resilience problem: staff need methods that work during outages, tools for verification, and ways to respond to synthetic content.

Respondent group	Requested further training	Most requested topics
Regional and local officials	96.2%	Crisis protection such as blackouts 68.8%; verification without stable internet 62.8%; advanced OSINT 58.8%; deepfake threats to government 57.1%; countering AI-driven disinformation 54.3%.
Central-government officials	97.8%	Advanced OSINT tools 77.5%; countering AI-driven disinformation 61.8%; deepfake threats to government 46.1%; crisis response 34.8%.

These results support a follow-on model that combines scalable online refreshers with practical, role-specific exercises. Regional and local stakeholders reported a gap in hands-on simulation-based training, while central participants asked for deeper operational capability. The project report also recommends more inter-agency escalation protocols, clearer coordination with platforms, and practical standard operating procedures tailored to institutional roles.

Delivery against planned outputs

Planned output or target	Evidence reported	Assessment and qualification
At least three research reports, briefs, or policy recommendations	Two named products appear in the project report: the public-sector policy paper and the AI cybersecurity recommendations.	The materials provided do not identify a third product. The target of three is therefore not evidenced as met in the supplied record.
At least 250 trainees, with 150 government and 100 civil-society participants	729 participants across the two core government cohorts, plus the separate 14-person UNDP early-career group. Core reach exceeded the overall headcount target.	The supplied breakdown documents government and aspiring public servants, but does not show whether the planned 100 civil-society places were reached. Overall headcount alone does not verify the stakeholder-type target.
Three simulation exercises	The MEL matrix records three Let's Data sessions against the target of three simulation exercises.	The activity count is reported. The supplied evidence does not describe scenario design, participant performance scoring, or learning retention.
Two UK-Ukraine knowledge-exchange events and seven institutions in a hub or consortium	The project report says the exchange was not delivered and records no participating institutions for this pathway.	Output not delivered; related immediate and intermediate outcomes were not achieved. SSSCIP's request on electronic trust services is a concrete future entry point, not a completed result.

The project report states that core activities were completed within the planned budget and timeframe and assesses value for money as strong. It attributes efficiency to online scale for 635 participants, in-person focus for central-government teams, and existing institutional partnerships. The supplied materials do not provide the actual budget totals or a cost-per-participant figure, so this synthesis does not calculate a financial efficiency ratio.

Risk, sustainability and democratic safeguards

Risk in the supplied register	Evidence and mitigation during delivery	What must be measured next
War, blackouts, connectivity failures and staff safety	The report describes prolonged winter outages. Sessions were recorded, materials were shared, and the Kyiv venue had backup power. Both online and in-person training proceeded.	Track attendance recovery, offline access, availability of official communications and citizen-facing services during future disruptions.

Risk in the supplied register	Evidence and mitigation during delivery	What must be measured next
Cyberattack and disinformation targeting partners or the project	The register rates these risks high. It records no change in the cyber risk and no evidence in the supplied report that a disinformation campaign caused reduced participation or disrupted delivery.	Record incidents, time to verify, corrective action, data loss, and whether responses remain transparent and proportionate.
Staff turnover and institutional reorganisation	The project reports reshuffles and rotations affecting the Ministry of Digital Transformation, Ministry of Defence, Ministry of Culture and security institutions. It highlights repeated re-onboarding and loss of institutional memory.	Measure trained staff retention, onboarding of replacements, handover quality, and adoption of role-based guidance beyond individual focal points.
Tool licences and soft-law governance	Continued access to Let's Data was not secured in the evidence supplied; CCD was exploring funding. The recommendations provide methodological guidance but are not shown to be mandatory across institutions.	Track licence renewal, active monthly users, agency procedures incorporating the guidance, and implementation checks at 3, 6, and 12 months.
Gender balance and language access	Women were 62% of online regional/local participants but 30.9% of central in-person participants. The register also notes English proficiency as a possible barrier to using tools and materials.	Disaggregate advanced training and tool use by gender, role, institution, and region; document use of translated materials and participation in decision-making roles.

The report states that no formal transition plan was implemented within the project period. Sustainability is instead supported by early integration into SSSCIP guidance, continued institutional demand, and observed day-to-day application. These are promising foundations, but they do not resolve licensing, turnover, or the difference between advisory guidance and routine compliance.

A further democratic safeguard is to measure how institutions counter manipulation, not only whether they detect it. The next phase should assess transparency, proportionality, correction practices, and respect for freedom of expression alongside speed and reach.

What the evidence supports and what remains open

- Supported: strong immediate shifts in self-reported knowledge and confidence; practical examples of phishing prevention and information verification; one formal national guidance uptake; and early use of monitoring methods in several institutions.
- Plausible but not yet established: wider or sustained improvement in inter-agency response, national detection speed, public confidence, citizen access to reliable information, or protection of democratic processes.
- Not supported in this reporting period: a completed UK–Ukraine knowledge exchange, systematic cross-border cooperation, uniform implementation of SSSCIP guidance, or durable access to the monitoring tool.
- Attribution: project evidence is triangulated and the project rates its contribution Level 4 of 5. The short implementation period, self-reported surveys, small follow-up sample, and absence of a comparison group prevent a causal claim about national democratic outcomes.
- Equity: overall female participation was strong in the regional online cohort, but central security-sector training remained male-dominated. The reporting period does not show whether women's participation translated into sustained influence over institutional decisions.

Priority measurement plan for a follow-on phase

A next phase should follow the same institutions over time, pair participant reports with administrative records, and define a denominator for every institution-level target. These measures would test whether early learning becomes durable practice and whether that practice protects democratic functions.

Measurement question	Suggested indicator and verification
Did guidance enter institutional practice?	Number and share of supported institutions with approved procedures, incident playbooks, AI procurement checks, or staff induction incorporating the SSSCIP guidance; verify through documents and staff interviews.
Did response improve?	Median time from first alert to verification and response; proportion of incidents with evidence retained, validated correction, and documented follow-up; assess false positives and false negatives.
Did institutions remain operational during disruption?	Availability of citizen-facing communications and key services during cyber or energy incidents; attendance recovery after outages; use of offline verification materials.

Measurement question	Suggested indicator and verification
Did capacity survive staff turnover?	Share of trained units with at least two trained staff, replacement onboarding within a defined period, and documented handover; reassess at 3, 6, and 12 months.
Was use inclusive and rights-respecting?	Training and tool uptake disaggregated by gender, region, role, and institution; review whether public responses are transparent, proportionate, and respectful of freedom of expression.
Was support sustainable?	Renewal of monitoring-tool access, active users and alert use over time, domestic or institutional financing, and continued use after project-supported licences end.

Methodological notes and evidence register

The evidence base combines the project's narrative outcome and impact report signed 16 April 2026, the MEL results framework through March 2026, the risk update supplied for this synthesis, and the evidence types described in those records. The project report cites attendance lists, pre- and post-training surveys, a 44-person online follow-up, key-informant interviews, participant messages, Let's Data usage analytics, official SSSCIP communication, and documentary approval of the recommendations.

Evidence source	Supports	Limit
Attendance and participant lists	Reach by cohort, gender, institution type, and role.	The aggregate total is inconsistent across source passages; attendance does not prove learning.
Pre- and post-training surveys	Cohort-level self-assessed knowledge and confidence shifts.	Responses are aggregate rather than individually matched; no independent skills test or comparison group.
Follow-up survey and stakeholder messages	Examples of applied learning and early practice change.	Follow-up n=44 is a small, self-selected subset; not representative of all online participants.
Key-informant interviews and written institutional feedback	Usefulness, tool use, strategy work, and operational relevance.	Stakeholder accounts are project-collected; ongoing changes require later document verification.
Let's Data product analytics	Onboarding registrations and interaction with alerts.	Open rate does not measure accuracy, response speed, or decision quality.
SSSCIP order and official communication	Formal approval of recommendations and acknowledgement of IIG input.	Approval does not demonstrate adoption by all intended institutions or enforceable compliance.
Risk register and delivery records	Context, anticipated risks, mitigation, and whether activities proceeded.	Risk ratings are not outcome measures; absence of a recorded incident is not proof of no exposure.

Data reconciliation note

The source files report 635 online participants and 94 in-person participants, a core total of 729. A separate 14-person women's cohort is reported, while some narrative passages give 743 participants overall. One aggregate gender figure lists 438 women and 306 men, which sums to 744; another lists 423 women and 306 men for the 729-person core cohort. This report uses 729 for the two core cohorts and keeps the 14-person cohort separate. The project dataset should reconcile the one-person difference and clarify whether the total of 743 includes all supplementary participants.

The MEL framework set a target of at least three research products, while the reporting record names two. It also set a training target that included 100 civil-society participants; the supplied output breakdown does not evidence that group. The report should not claim these specific targets were met without confirming the final output register.

All percentages are reproduced from the supplied project records. Percentage-point changes are calculated from reported pre- and post-training shares.