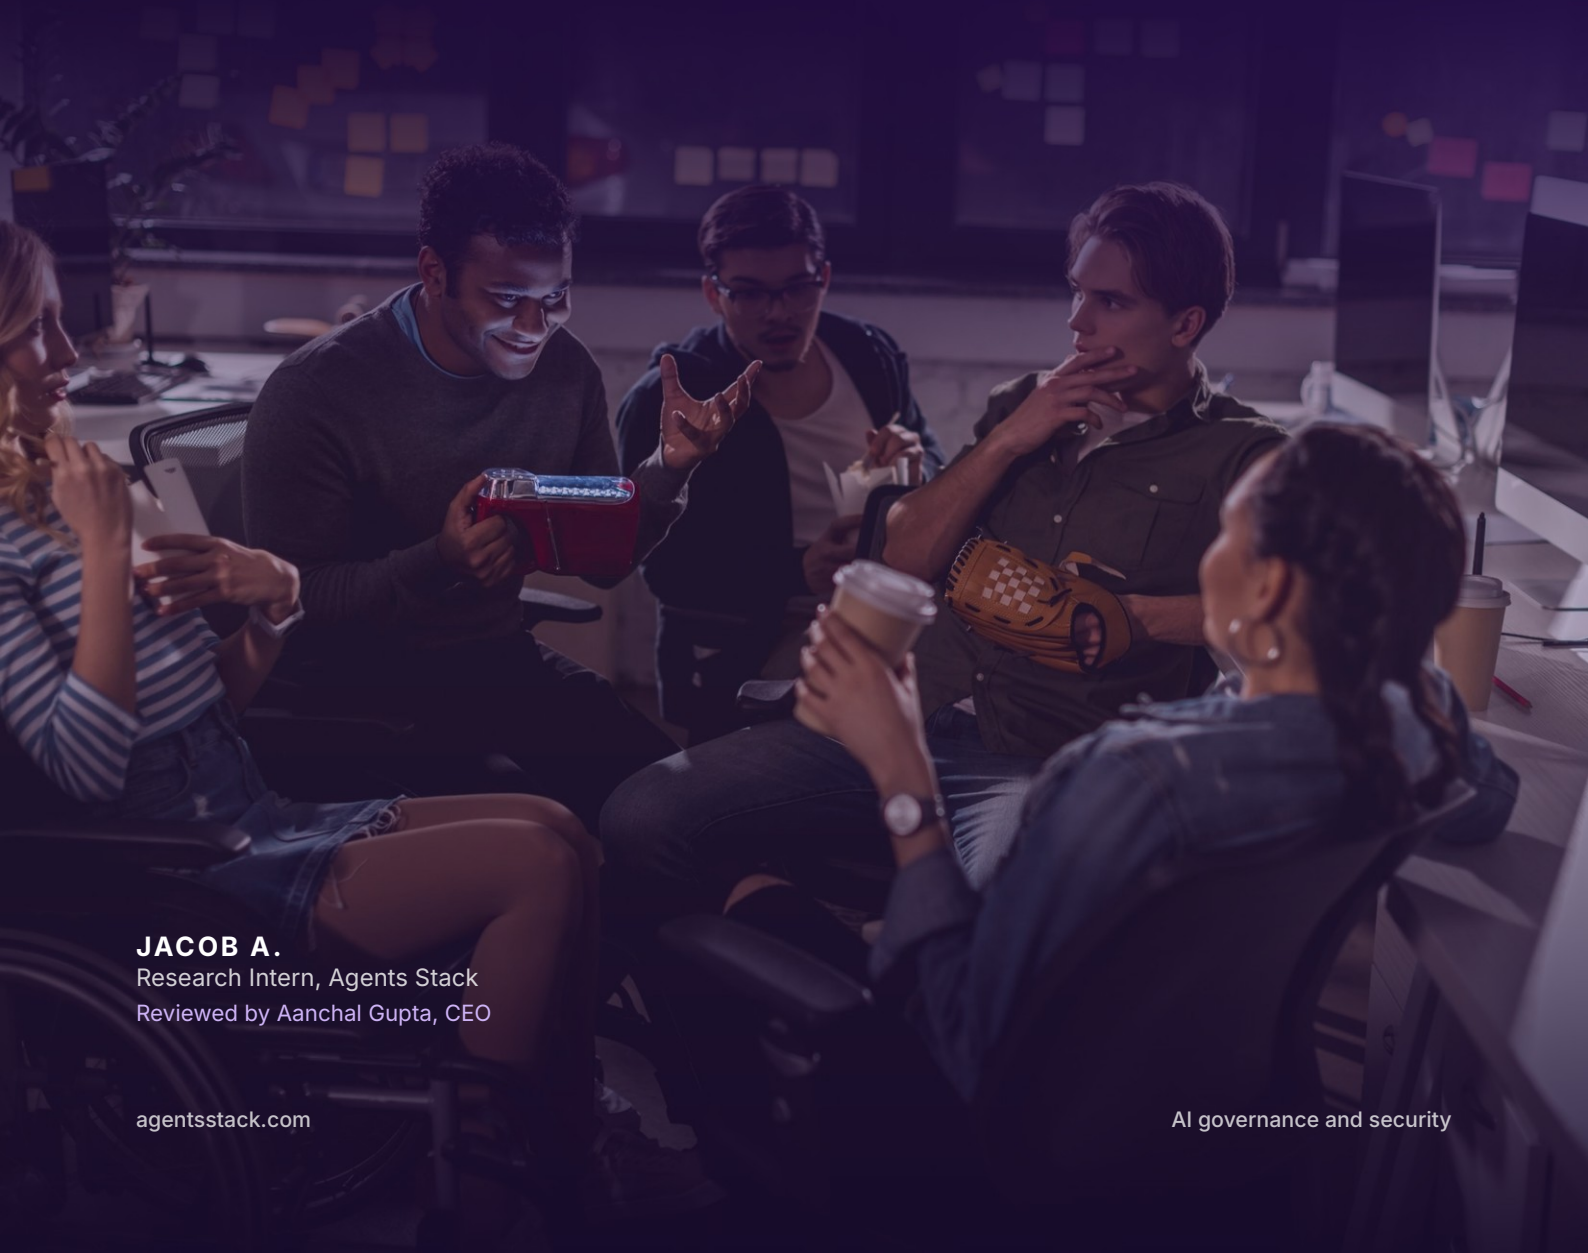
Agents Stack

WHITEPAPER 01 | OCTOBER 2026

Overcoming the Fear Factor

Solutions for AI adoption in nonprofits and think tanks

Most nonprofits, think tanks and public bodies already use AI, but few have adopted it as institutions. The gap comes down to willingness and security more than regulation.



JACOB A.
Research Intern, Agents Stack
Reviewed by Aanchal Gupta, CEO

AT A GLANCE

Key findings

AI arrived in NGOs, think tanks and government offices through individual staff, not through the institutions they work for. The gap with large companies is not one of use. It is one of institutional adoption.

92% vs 7%

Use without adoption

92% of nonprofits use AI, but only 7% report major strategic impact and just 4% have documented AI workflows.

Dec 2027

Regulation is rarely the blocker

Outside Europe, law is rarely the binding constraint. Even the EU has delayed its high-risk rules, and Singapore's agentic AI framework is voluntary.

50% vs 19%

Those who use AI most govern it least

Half of humanitarians in Sub-Saharan Africa use AI daily, yet only 19% report having an AI policy.

40% faster

The fear of replacement is not yet supported

Professionals using generative AI finished writing tasks 40% faster with 18% higher quality. Most AI use augments work rather than automating it.

1 in 5

Shadow AI widens the security gap

NGOs and think tanks are prime targets for state-sponsored attackers, and shadow AI played a part in one in five breaches studied by IBM.

3 million

The gaps can be closed

Organizations with sanctioned tools and human oversight already reach millions of mothers with health advice and target cash aid more accurately.

The recommendation

Own the decisions, outsource the machinery.

Keep policy, risk appetite and accountability in-house. Buy the tools, monitoring and threat intelligence. Small NGOs lean on external templates and volunteer support; large INGOs, UN agencies and governments can carry more in-house.

01 EXECUTIVE SUMMARY

AI is already inside the sector. The institutions are not.

AI has already reached NGOs, think tanks and government offices. Much like early AI adoption in the for-profit world, it arrived through individual staff rather than through the institutions they work for.

Three-quarters of humanitarian workers now use AI daily or weekly, yet fewer than a quarter of their organizations have a formal AI policy (Humanitarian Leadership Academy). The gap with large companies is not one of use. It is one of institutional adoption: budgets, policies, workflows and scale.

This report asks what is holding that official adoption back. Is the hurdle regulatory, legal, cultural or psychological? And can the market, or the organizations themselves, solve it?

The evidence across six regions points in several directions. Outside Europe, regulation is rarely the binding constraint. Even inside it, the EU has delayed its own high-risk rules to December 2027, and Singapore's agentic AI framework is voluntary.

The everyday barriers are capacity and culture. There is often no money set aside for AI, thin skills at leadership level, policies that are missing or that actively discourage staff from using AI, and a gap in conviction between executives and frontline staff. In European NGOs and think tanks especially, that lack of funding reflects the wider psychological and cultural caution that has slowed AI adoption across Europe.

Two deeper issues

A willingness gap, psychological and political

Many professionals, especially in policy institutes and other roles further from fieldwork, fear AI will replace the thinking and writing their careers rest on. Donors, communities and the public are uneasy about its use.

A security gap

NGOs and think tanks are among the sectors most targeted by state-sponsored hackers. Unofficial AI use widens that exposure: shadow AI played a part in one in five breaches studied by IBM (IBM via Help Net Security).

Both gaps can be closed. Organizations that have done so are reaching millions of mothers with health advice, targeting cash aid more accurately and analyzing tens of thousands of public responses in hours. This report sets out the products that could help the rest follow, and weighs external partners against internal approaches to governance and security.

Large companies have moved quickly: 88% now use AI in at least one business function (McKinsey). The question is why NGOs, think tanks, foundations and government bodies have not institutionalized it in the same way.

02 THE GLOBAL PICTURE

Use without adoption

Globally, NGOs are not refusing AI. Some 92% of nonprofits use it (NonProfit PRO), as do 71% of think tanks (On Think Tanks) and over 70% of public servants (ITIF).

92%
of nonprofits use AI
NonProfit PRO

71%
of think tanks use AI
On Think Tanks

70%+
of public servants use AI
ITIF

What has not spread is organizational adoption. Only 7% of nonprofits report major strategic impact, and just 4% have documented AI workflows (Webiano).

The difference from companies is commitment. Corporate high performers put over a fifth of their digital budgets into AI (Neodata Group), while more than half of nonprofit executives have no AI budget at all (Chronicle of Philanthropy).

THE AI PARADOX

The humanitarian sector has a name for this: individual uptake far outpaces organizational readiness.

03 REGION BY REGION

Use runs ahead of governance almost everywhere

AI is not spreading from the Global North to the South. The heaviest humanitarian use is in Asia-Pacific and in crisis settings such as Kenya (65% daily), Sudan and Bangladesh, while Western Europe reports the lowest (Humanitarian Leadership Academy).

Governance runs the other way: the regions that use AI most govern it least. Part of the problem may be regulatory, but the pattern also points to psychological and cultural resistance to AI in high-governance jurisdictions like the EU.

FIGURE 1

AI use, adoption, law and the main hurdle by region

REGION	HUMANITARIAN DAILY AI USE	GOVERNMENT ADOPTION	BINDING AI LAW?	MAIN HURDLE
North America	52%*	US: uneven	No; federal rules bind agencies only	Funding, staff buy-in
Europe and UK	39% (Western Europe)*	Germany, France: cautious; UK: uneven	EU AI Act; high-risk rules delayed to Dec 2027	Caution, leadership skills, funding
Asia-Pacific	55%	Singapore, India: advanced	Mostly voluntary	Frameworks too complex for small organizations
Sub-Saharan Africa	50% (19% have a policy)	South Africa: uneven	Draft policies and bills	Governance capacity, data sovereignty
Middle East and North Africa	51%	Saudi Arabia: advanced	Varies by country	Crisis-context data risk, thin evidence
Latin America	40%*	Brazil: uneven	No; Brazil's bill stalled	Governance and inclusion

Sources: Humanitarian Leadership Academy and Data Friendly Space, January 2026 pulse survey; Public First for the Center for Data Innovation, Public Sector AI Adoption Index 2026. *Smaller samples, treat as indicative.

North America

Only 4% of nonprofits have fully integrated AI. 44% of executives see major untapped opportunity, against 23% of staff (Chronicle of Philanthropy).

Europe and the UK

The EU has delayed its high-risk rules to December 2027 (Hogan Lovells), yet adoption stays cautious. 76% of UK charities use AI, but only 44% have a digital strategy, and a third rate their CEO's AI skills as poor (Third Sector).

Asia-Pacific

Singapore launched the first agentic AI governance framework (Baker McKenzie), and about 80% of its public officers have used its secure AI assistant (Singapore MDDI). Community organizations still need simpler versions.

Sub-Saharan Africa

The paradox is sharpest here. Half of humanitarians in the region use AI daily, but only 19% report having an AI policy, and among local organizations worldwide the figure is just 13% (Humanitarian Leadership Academy). Kenya's draft policy would keep Kenyan data under Kenyan law (The Star).

> **The people with the least protection are using AI the most. Somebody should put that at the top of the to-do list.**

Middle East and North Africa

Gulf governments push adoption from the top, while crisis-driven use, such as 60% daily use in Sudan, runs ahead of data safeguards (Humanitarian Leadership Academy). No regional NGO survey exists.

> **In a war zone, a leaked spreadsheet doubles as a target list.**

Latin America

Brazil's AI bill has stalled since December 2024 (CASRAI), and the IDB rates the region's responsible-AI maturity at 2.9 out of 5 (IDB Lab).

04 THE WILLINGNESS GAP

Psychological: the fear of replacement

In NGOs, think tanks and ministries, careers are built on thinking and writing, which is exactly what generative AI does visibly well. Many professionals see AI as a substitute for their work.

In the US, 52% of workers are worried about AI's impact on the workplace (Pew Research Center). Inside nonprofits, the staff closest to communities are the most likely to refuse AI.

So far, the evidence does not support the fear:

40%
faster on writing tasks, with 18% higher quality
Noy and Zhang, Science

57%
of AI tasks involve augmentation rather than automation
Anthropic Economic Index

None
discernible economy-wide labor disruption since ChatGPT's release
Yale Budget Lab via Fortune

> **The robots are coming for the policy analyst's formatting. The analyst should be delighted: more time to meet policymakers on the ground and think harder about how policy should work.**

The honest caveat. Employment of 22- to 25-year-olds in the most AI-exposed jobs fell 13% relative to others, concentrated where AI automates rather than augments (Stanford). The risk falls mainly on routine entry-level tasks. The tacit judgment of experienced professionals stays valuable, and organizations can choose to protect junior learning roles.

AI is a very fast intern. The person who knows which minister to call, which donor to thank and which village to visit first still runs the show.

The hesitation runs deepest in the West. Western European humanitarians report the lowest daily AI use of any region, Germany and France remain cautious public-sector adopters, and over half of US workers are worried about AI at work. In these countries, professional identity is often built on authorship.

Yet AI cannot do the most important part of that work. It can draft a brief, but it cannot decide which question matters or set out a vision for the future of a policy area.

KEEP HIRING INTERNS

Internships are how the sector grows its future senior staff. With entry-level roles the most exposed, protecting them protects the pipeline of judgment. Interns also broaden the range of views inside an organization and keep it close to its community-driven mission. Students bring new questions, languages, lived experience and generational perspectives that a model trained on existing text cannot supply on its own. The smart move is to redesign internships around fieldwork, research design, checking AI output and community engagement, and away from formatting.

Political: legitimacy, trust and values

Mission-driven organizations depend on trust, which makes AI a political question for them. Donors are happy to meet the robot. They just want an introduction.

FIGURE 2

Five pressures that make AI a political question

PRESSURE	EVIDENCE
Donors	31% of US donors would be less likely to give to a charity using AI; 92% want clear explanations of AI use
Reputation	Amnesty deleted AI-generated protest images after a backlash and paused their use until it had a policy
Public trust	69% of UK citizens are unaware of any AI-run public service
Values	Nonprofit staff rank environmental impact as their top AI concern, above job loss
Funding	Only 12% of US nonprofit executives have received a grant for AI

Sources: Chronicle of Philanthropy (donor study and nonprofit executive survey); Kindsight, 2025 Donor Perceptions of AI Report; PetaPixel; IT Brief UK.

WHAT CLOSES THE GAP

Disclose AI use by default, involve frontline staff in setting the rules, commit to augmentation rather than job cuts, and have a policy in place before any controversy. Where guidance is clear, 91% of public servants feel confident using AI (ITIF).

05 CYBERSECURITY

The security gap is real. Avoidance only hides it.

Nonprofit executives name data, privacy and security as their biggest barrier to AI (Chronicle of Philanthropy). The caution is well founded.

Nearly 80% of nation-state cyberattacks in 2020-21 targeted government, think tanks and NGOs (Microsoft Digital Defense Report). In 2022, a breach at a contractor exposed data on 515,000 vulnerable people held by the Red Cross (Al Jazeera).

FIGURE 3

The security gap in numbers

INDICATOR	FIGURE	SOURCE
NGOs hit by a breach or critical data incident	**53%**	NetHope, 2025 State of Humanitarian and Development Cybersecurity Report
UK charities with an incident response plan	**19%**	UK Cyber Security Breaches Survey 2025
Breaches involving shadow AI	**1 in 5**	IBM, Cost of a Data Breach Report 2025 (adds US$670,000 on average)
Organizations with AI breaches lacking AI access controls	**97%**	IBM, Cost of a Data Breach Report 2025

Figures as reported by each source. IBM figures via Help Net Security.

AI widens this exposure in three ways

- **Shadow AI leaks data.** 59% of employees use unapproved AI tools, and 75% of those share sensitive information (Techstrong.ai).
- **Agents can be hijacked.** Nearly 30% of humanitarians use AI agents (Humanitarian Leadership Academy), and prompt injection is the top risk for AI applications (OWASP Top 10 for LLM Applications).
- **Attackers use AI.** Phishing and deepfake calls impersonating executives and donors are getting cheaper and more convincing (CGNET).

ON OUTRIGHT BANS

Bans simply push staff use of AI out of sight. The answer is approved tools, hard boundaries for beneficiary data, access controls on AI, and limited permissions with human approval for agents.

> **A ban on AI is a policy on paper and a confidential memo on someone's personal ChatGPT account.**

06 THE ART OF THE POSSIBLE

Organizations that already crossed both gaps

Some organizations have already crossed both gaps. Each gave people a safe, sanctioned tool and kept humans in charge of judgment.

FIGURE 4

Four examples and their results

EXAMPLE	RESULT
Jacaranda Health, Kenya AI triages mothers' SMS questions and flags urgent cases	Reaches about 3 million mothers; 20% more prenatal check-ups
Togo and GiveDirectly Machine learning finds the poorest households	Cut wrongful exclusion from cash aid by 4% to 21%
UK government, Consult Analyzes public consultations	50,000+ responses in two hours for £240
Singapore government, Pair Secure assistant for public officers	Used by about 80% of 150,000 officers

Sources: AI4D on Jacaranda Health, J-PAL on Togo and GiveDirectly, Civil Service World on Consult and Singapore MDDI on Pair.

07 WHAT COULD CLOSE THE GAPS

Products the sector needs

FIGURE 5

Products mapped to the gap each one closes

PRODUCT	WHAT IT DOES
SECURITY GAP	
Secure AI workspace with data guard	A sanctioned tool that redacts beneficiary data and removes the reason for shadow AI
Shadow AI discovery	Shows which AI tools staff use and what data flows to them
Agent governance	Registry, permissions, human approval and logs for AI agents
Fraud and phishing defense	Deepfake detection and payment verification, priced for small NGOs
WILLINGNESS GAP	
Policy and compliance in a box	Templates mapped to the EU AI Act, Singapore's framework and SAFE AI
Literacy and transparency tools	Staff and board training, plus donor-facing disclosure of AI use
CAPACITY GAP	
Sector copilots	Grant writing, donor reporting and policy research assistants
Pooled, low-bandwidth services	Shared, locally hosted and multilingual tools for small organizations

The natural starting points are the secure workspace, policy in a box and agent governance. Pricing matters as much as the product: tiered by budget, sponsored by funders, or bought through consortia.

> **Forget the moonshot. The sector needs a locked filing cabinet with a very smart assistant inside.**

08 EXTERNAL PARTNERS OR INTERNAL APPROACHES?

Own the decisions, outsource the machinery

FIGURE 6

External frameworks and partners compared with internal capability

	EXTERNAL FRAMEWORKS AND PARTNERS	INTERNAL PROTOCOLS AND CAPABILITY
Speed and cost	Fast; predictable subscription	Slow to build; high fixed cost
Credibility	Recognized by donors and regulators	Depends on own reputation
Expertise	Specialists and threat intelligence on hand	Hard to hire; one organization's view
Fit to mission	Generic; may miss protection risks	Reflects own values and communities
Ownership	Risk of box-ticking	Staff who write rules tend to follow them
Data control	Third-party and vendor risk, as in the Red Cross breach	Full control, which matters under data sovereignty rules

Both have real merits. The evidence favors a hybrid: **own the decisions, outsource the machinery.** Keep policy, risk appetite and accountability in-house; buy the tools, monitoring and threat intelligence.

> **Outsource the plumbing, never the conscience.**

The balance shifts with size. Small NGOs should rely mostly on external templates, tools and volunteer support such as CyberPeace Builders (The Hague). Large INGOs, UN agencies and governments can carry more in-house, as Singapore shows.

09 CONCLUSION

The technology is ready. The excuses are running out.

Outside Europe, regulation is the weakest explanation for slow AI adoption, and capacity gaps are universal. The barriers that best explain the lag are willingness and security.

The willingness gap is psychological, a fear of replacement the evidence does not yet support, and political, rooted in donor and public trust. The security gap is real, but avoidance only hides AI use.

Both gaps can be closed with the right partners and governance. The organizations already doing so had the same technology as their peers, plus the willingness to try and the governance to do it safely.

It is time to start imagining what is possible.

HOW AGENTS STACK CAN HELP

Start with knowing what AI is already in use

The hybrid model above starts with knowing what AI is already in use and setting the rules for it. The Agents Stack AI Governance and Assessment is a fixed-scope engagement that maps your AI estate, including shadow AI, classifies each use by risk, defines oversight for AI agents, and ties your policies to evidence you can show a board, a donor or a regulator. You keep the decisions and the findings.

BOOK AN AI GOVERNANCE ASSESSMENT

agentsstack.com/lp/ai-governance-assessment

SOURCES & EDITORIAL NOTES

Evidence you can inspect

HOW TO READ THIS PAPER

This paper was researched and written by Jacob A., Research Intern at Agents Stack, and reviewed by Aanchal Gupta, CEO of Agents Stack, for accuracy, structure and sourcing. Figures are reported as published by each cited source; survey results are self-reported and some regional samples are small. The examples, product categories and recommendations are editorial proposals, not verified Agents Stack customer results.

This paper does not provide legal, financial or regulatory advice. No affiliation with or endorsement by cited organizations is implied.

1. Humanitarian Leadership Academy and Data Friendly Space, How are humanitarians using AI? January 2026 pulse survey briefing note
2. IBM, Cost of a Data Breach Report 2025, via Help Net Security
3. McKinsey & Company, AI at work but not at scale (State of AI 2025)
4. NonProfit PRO, Nonprofit AI adoption hits 92%, but only 7% see major impact
5. On Think Tanks, State of the Sector Report 2025
6. Public First for the Center for Data Innovation, Public Sector AI Adoption Index 2026, via ITIF
7. Webiano, Every charity uses AI now and almost none are ready
8. Neodata Group, McKinsey State of AI 2025
9. Chronicle of Philanthropy, The nonprofit AI gap: bosses are bullish, staffs are wary
10. Hogan Lovells, EU legislators agree to delay for high-risk AI rules
11. Third Sector, Zoe Amar on the Charity Digital Skills Report 2025
12. Baker McKenzie, Singapore governance framework for agentic AI launched
13. Singapore Ministry of Digital Development and Information, progress of adopting AI tools within public service
14. The Star (Kenya), State tightens rules on AI use of Kenyans' data
15. CASRAI, Brazil's PL 2338/2023 status
16. IDB Lab, fAIr Tech Radar
17. Pew Research Center, U.S. workers are more worried than hopeful about future AI use in the workplace
18. Noy and Zhang, Experimental evidence on the productivity effects of generative AI, Science
19. Anthropic, The Anthropic Economic Index
20. Fortune, If AI is roiling the job market, the data isn't showing it (Yale Budget Lab)
21. Brynjolfsson, Chandar and Chen, Canaries in the coal mine?, Stanford
22. Chronicle of Philanthropy, How do donors feel about AI?
23. Kindsight, 2025 Donor Perceptions of AI Report
24. PetaPixel, Amnesty deletes AI images after backlash
25. IT Brief UK, UK public sector races ahead with AI as trust lags
26. Microsoft, Digital Defense Report: new insights on nation-state attacks
27. Al Jazeera, Cyberattack on ICRC exposes data on 515,000 vulnerable people
28. NetHope, 2025 State of Humanitarian and Development Cybersecurity Report
29. Northern Ireland Cyber Security Centre, Cyber Security Breaches Survey 2025 statistics
30. Techstrong.ai, Shadow AI is the silent data risk taking over workplaces
31. OWASP Top 10 for LLM Applications 2025, explainer via DEV Community
32. CGNET, Trust under attack: AI deepfakes and nonprofit security
33. AI4D, PROMPTS (Jacaranda Health)
34. J-PAL, Machine learning and mobile phone data for social protection
35. Civil Service World, DSIT says AI tool could free up 75,000 civil servant days a year
36. The Hague, CyberPeace Institute support for NGOs

Agents Stack

AI governance and security for organizations that operate across borders.

Agents Stack helps enterprises, public bodies and mission-driven organizations adopt AI with the governance, security and evidence their boards, donors and regulators expect.

agentsstack.com

9 Raffles Place, #06-01 Republic Plaza, Singapore 048619

© 2026 Agents Stack. Whitepaper 01, October 2026. This paper is for general information only and does not constitute legal, financial or regulatory advice. Figures are attributed to their original sources. No affiliation with or endorsement by cited organizations is implied.