ENCRYPTED BY DESIGN 2026

AI & ENCRYPTION PLAYBOOK

FOR AFRICAN NONPROFIT TECHNOLOGY ORGANIZATIONS



A practical playbook on AI risk, encryption, privacy by design and organizational readiness

ACKNOWLEDGEMENTS

This playbook was developed through the work and contributions of individuals and organisations committed to advancing responsible AI, strong encryption, privacy, and digital security for public-interest technology.

We extend our sincere appreciation to the experts, practitioners, nonprofit organisations, civil society organisations, technology professionals, and digital rights advocates whose knowledge, experiences, and perspectives have helped shape this resource.

We are particularly grateful to the contributors who shared their expertise on AI safety, cybersecurity, privacy, encryption, data protection, digital rights, and nonprofit technology. Their practical insights have helped ensure that this playbook speaks to the realities faced by organisations working with sensitive information and communities.

The 2026 edition builds on the conversations and learning from the previous edition, with a stronger focus on the intersection of AI, encryption, and the protection of beneficiary data.

With Appreciation
TrustedTech Africa
Project Lead & Publisher

Internet Society Foundation
For supporting the Global Encryption Day 2026

Our Contributors and Reviewers

Toluwani Akinniyi
Muhammad Arshi Wasique
Omolara Esther Hamzat
Ephraim Akuetteh Mensah
Joy Uchechi Eziashi

For sharing their expertise, experiences, and recommendations

TABLE OF CONTENTS

- Section 1: Understanding AI and Digital Security Risks……3
- Section 2: Why Strong Encryption Matters……4
- Section 3: AI and Beneficiary Data Protection……5
- Section 4: Encryption by Design……7
- Section 5: Privacy by Design and the Interactive Checklist……8
- Section 6: Threat Modelling for Nonprofits……10
- Section 7: Secure AI Tools and Open Source Resources……11
- Section 8: African Policy and Regulatory Considerations……13
- Section 9: Organizational Encryption Action Plan……14
- Section 10: AI Safety and Encryption Self Assessment……15
- Case Study: AI in Practice……16
- Appendix……17

  Appendix A: Quick Reference Checklist……17

  Appendix B: Key Terms……17

  Appendix C: Useful Resources……18
- References……18

Section 1

Understanding AI and Digital Security Risks

Artificial intelligence is now part of everyday operations for nonprofit and public interest organizations. Teams use AI tools to draft reports, translate documents, summarize meetings, analyze research and prepare donor communications. For organizations with limited staff, these tools can be genuinely valuable. They also introduce governance, privacy and cybersecurity considerations that did not exist a few years ago.

Every time a staff member enters information into an AI system or connects an AI tool to email, cloud storage or another platform, organizational data may move beyond the organization's direct control. For nonprofits working with vulnerable communities, survivors, refugees, health information or politically sensitive programs, that movement of information can create real risk. The starting point for responsible AI use is therefore a simple question that should guide every decision in this playbook:

What happens to our information when we give it to an AI system?

The AI Data Lifecycle

When someone uses an AI system, information is collected, processed and sometimes stored. Retention depends on the provider and its configuration, and may include prompts, uploaded files, interaction histories and technical logs. Data may also pass through several technology providers, since an AI application often depends on a separate model provider or cloud host, creating a wider technology supply chain. Using an AI tool can therefore mean transferring organizational information to several external parties at once, and the organization remains responsible for deciding whether that risk is acceptable.

Why This Matters for Nonprofits

Consider an organization supporting survivors of gender based violence. Its records may include names, shelter locations and case histories. If a staff member copies case notes into a public AI tool to save time, the key questions are whether personal
information was transferred to an external provider, whether it was retained, and whether the organization understood the provider's safeguards before that information was shared. The same issue arises whenever witness statements, patient information or politically sensitive material are processed by AI without appropriate governance.

Key Risks to Watch

Data leakage and retention

- Staff may copy confidential information into AI tools without realizing it can be transferred to, and retained by, an external provider. Closing a chat window does not necessarily delete the information.

Third party and supply chain risk.

- The visible AI provider often depends on other cloud, model or technology providers, so the real question is not only who you signed a contract with, but who else can reach the data.

Access and connected systems.

- AI tools connected to email, cloud storage or document repositories can access far more information than a user realizes. The principle of least privilege matters: AI applications should receive only the access their task requires, protected by unique accounts and multi factor authentication.

Inaccurate outputs.

- AI systems can produce confident but incorrect or incomplete information, particularly in health, legal, safeguarding or policy contexts. Important AI generated outputs should be reviewed by a qualified person before they are relied upon.

Strong encryption remains an essential safeguard, covered in Section 2 and Section 4, but it is not a complete AI security strategy on its own. An AI system generally needs access to readable information to process it, so organizations must also ask who can access that information while it is being processed, where it is stored, how long it is retained and which third parties may receive it.

Data Minimization: The Simplest Safeguard

One of the most effective safeguards available to any nonprofit is data minimization. Imagine an organization analyzing interview transcripts related to political violence to identify common themes. The analysis may not require the names or contact details of the people interviewed. Removing that detail before submitting the material to an AI system significantly reduces the impact of any disclosure. The useful shift in thinking is to stop asking only, can we use AI for this task, and also ask, what is the minimum information this AI system actually needs.

Practical AI Security Checklist

- Classify the information as public, internal, confidential or highly sensitive before deciding whether AI is appropriate.
- Define which AI tools staff and volunteers are approved to use for organizational work.
- Secure AI accounts with unique logins and multi factor authentication rather than shared credentials.
- Review what an AI tool can access when it is connected to email, cloud storage or a database.
- Require human review of AI generated outputs that affect people, health, finance or public communication.
- Include AI related incidents, such as accidental disclosure, in existing incident reporting procedures.
- Assign one person clear responsibility for AI related risk, even without a dedicated security team.

Governance should be proportionate to risk. Many African nonprofits operate with small teams and limited budgets, so a reasonable starting point may be a one page acceptable use policy naming approved tools, prohibited information and who to contact when something goes wrong. Additional controls can be added as AI use expands. Nonprofits are trusted with sensitive information because communities believe it will be handled responsibly. AI does not transfer that responsibility to a technology provider. The organization remains accountable for what it collects, what it shares with an AI system, and what safeguards it applies throughout.

Section 2

Why Strong Encryption Matters

For a nonprofit, information is rarely just data in a computer. It can represent a person's identity, health, location or livelihood, and it can include confidential details about staff, donors, partners and operations. When such information is exposed, the consequences can extend well beyond a technical incident: people can be put at risk, programs interrupted and public trust damaged.

Strong encryption is one of the most important technical safeguards available to reduce this risk. Encryption transforms readable data into a form that cannot be understood without the correct key, so that information remains protected even if it is intercepted, copied or accessed without authorization. It should be treated as part of responsible information management, not simply an IT feature.

End to End Encryption

End to end encryption protects a communication from the moment it leaves the sender until it reaches the intended recipient, encrypted on the sender's device and decrypted only on the recipient's device. This matters for a field worker discussing a beneficiary case with a program manager, since someone intercepting that communication in transit should not be able to read it.

It is important to understand the limits. End to end encryption protects the content of a communication while it is being transmitted and stored within that system. It does not protect information copied into an unsecured document, captured in a screenshot, forwarded through another channel or stored on an insecure device. Encryption therefore needs to sit inside a wider security approach, not stand in for one.

In Transit and At Rest

Information in transit is data moving between systems, devices or people, for example an email or a file transfer. Information at rest is data stored on a device, server, database or backup. Both need protection, since securing one does not compensate for leaving the other exposed. A nonprofit's beneficiary or donor database, staff files and monitoring data are typically stored across laptops, servers, cloud platforms and backups, and a lost or stolen device with strong encryption is far harder for an unauthorized person to read.

Encryption and Public Trust

People share information with nonprofits because they expect it to be handled responsibly. A beneficiary provides personal details to receive assistance, a donor provides financial information, a partner shares confidential project data. When that trust is broken, consequences can include harm to the individuals affected, reputational damage and loss of confidence in the organization. Strong encryption cannot prevent every incident, but it can meaningfully reduce the impact of one by making stolen or intercepted information far harder to use, and it demonstrates that the organization takes its responsibilities seriously.

Encryption Is Not the Whole Strategy

An encrypted database can still be accessed by an employee with excessive privileges. An encrypted laptop can still be compromised through a stolen password. Encryption cannot stop a user from deliberately forwarding sensitive information, and it does not answer whether the organization should have collected that information in the first place. Effective protection combines encryption with access controls, strong authentication, secure devices, timely updates, backups, staff awareness and clear information handling practice. For a small nonprofit, this does not mean buying every available security product. It means identifying the information that would cause the greatest harm if exposed, and prioritizing protection around it.

Getting Started

- Identify sensitive information: know what you collect, where it lives, and who has access.
- Protect communications: use channels with appropriate encryption for the sensitivity of what is shared.
- Protect devices: enable full device encryption on laptops, phones and tablets used for organizational work.
- Protect backups: a backup holds the same sensitive information as the original system and deserves the same protection.
- Review third party platforms before uploading sensitive data, including cloud and AI services.
- Manage keys and passwords carefully, since encryption only protects information when its keys are properly secured.

The goal is not to encrypt everything simply because it is possible. It is to understand what needs protection, where it is most exposed, and to build encryption into the systems and processes that handle it, a design approach covered in full in Section 4.

Section 3

AI and Beneficiary Data Protection

Beneficiary information is often the most sensitive category of data a nonprofit holds, and increasingly it is the category most likely to end up inside an AI tool, whether through a case summary, a translation task or a reporting shortcut. This section sets out practical guidance for protecting it.

What Beneficiary Data May Include

Beneficiary data can include names and contact details, household or location information, health and welfare details, financial or eligibility information, identity documents, biometric data, photographs and recordings, case narratives, safeguarding concerns, and family or relationship information. Much of this is sensitive by nature and some of it can directly endanger a person if exposed.

Why It Needs Careful Protection

Beneficiaries are usually in a position of dependency relative to the organization serving them. They rarely have the power to negotiate how their information is used, and exposure can carry consequences ranging from embarrassment and discrimination to physical danger, particularly for survivors of violence, refugees, children or politically exposed individuals. Protecting this information is inseparable from protecting the people it describes.

Collection and Minimization

Collect only the fields genuinely needed to deliver a specific service, and avoid asking for the same information repeatedly across different forms and systems. Before adding a new field to an intake form, ask whether the program can actually function without it.

Access, Storage and Sharing

Access controls

- Give staff and volunteers access to beneficiary records on a need to know basis tied to their role, use individual logins rather than shared accounts, and review access whenever someone changes role or leaves.

Secure storage

- Keep beneficiary data on approved, protected systems such as a case management platform or encrypted device, not personal email or messaging apps.

Secure sharing

- Share only what a partner, donor or government counterpart genuinely needs, through a secure channel, and prefer a summary over a full case file wherever that is sufficient.

Retention and deletion

- Agree how long beneficiary records are kept after a case closes or a program ends, consistent with donor and legal requirements, and build in an actual deletion or anonymization step rather than letting old records accumulate indefinitely.



Photographs, Recordings and Identifying Information

- Images and recordings of beneficiaries require appropriate, informed consent that accounts for the power imbalance between the organization and the person photographed. Consider whether identifying images are actually necessary for communications or fundraising, and avoid showing identifying features of survivors, children or other at risk individuals in public material. Raw field photographs and recordings deserve the same protection as any other sensitive record.

Using AI with Beneficiary Information

- Before entering beneficiary details into an AI tool, remove names and identifiers where the task allows it, confirm the tool is one the organization has approved, and treat any AI generated summary as a draft that a qualified person checks against the original notes before acting on it. Section 6 includes a worked example of exactly this scenario and what can go wrong when that review step is skipped.

Third Party Platforms

- Case management systems, mobile data collection apps, cloud storage and AI services are all third party platforms handling beneficiary data on the organization's behalf. Before adopting one, understand how the provider stores, protects, shares and retains that information, and where it is hosted.

Protecting Vulnerable Individuals and Communities

- Extra caution is warranted for children, survivors of violence, refugees, people with stigmatized health conditions, and human rights defenders, given the elevated consequences of exposure. This can mean using pseudonymous case identifiers instead of full names in shared reports, or restricting a dataset's circulation more tightly than standard practice would otherwise require.

Staff and Volunteer Responsibilities

- Staff and volunteers should receive brief, practical guidance on what counts as sensitive beneficiary information, which tools are approved to handle it, and what must never be shared informally, for example through a personal messaging app. This belongs in onboarding and volunteer agreements, not left to be learned informally.

Beneficiary Data Protection Checklist

- 1. We know what beneficiary data we collect and why each field is needed.
- 2. Access to beneficiary records is limited to staff and volunteers who need it for their role.
- 3. Beneficiary data is stored on approved, protected systems, not personal devices or messaging apps.
- 4. We have agreed how long beneficiary data is kept and how it is deleted or anonymized afterward.
- 5. Sharing beneficiary information with partners or donors is limited to what is genuinely necessary.
- 6. Photographs, recordings and identifying details are only used with appropriate consent and care.
- 7. Before beneficiary information goes into an AI tool, identifying details are minimized where possible.
- 8. Staff and volunteers have received basic guidance on handling beneficiary information responsibly.

Section 4

Encryption by Design

Encryption by design means that whenever an organization creates or adopts a new way of collecting, storing or sharing information, someone asks early how that data will be protected, both moving and at rest. This belongs in the same conversation as budget and timeline, not in a separate review after a system is already live. It matters most for the categories nonprofits already hold in abundance: beneficiary, donor, staff, financial, identity, case management and program monitoring data. Encryption is not a complete security solution by itself. It protects information from being read if intercepted or accessed without authorization, but it does not stop phishing, fix weak passwords, or prevent someone with legitimate access from misusing data. It is one strong layer within a broader approach.

Encryption Across the Data Lifecycle

Stage	What to check
Collection	Forms and apps should submit over a secure, encrypted connection (HTTPS), including mobile data collection tools.
Transmission	Confirm, rather than assume, that data moving between offices, partners or app servers is protected in transit.
Storage	Databases and file systems holding sensitive records should be encrypted at rest, especially on older or locally hosted systems.
Processing	Data is often briefly decrypted for use, so access controls and staff practice matter as much as encryption here.
Sharing	Avoid plain email attachments for sensitive files; use an access controlled platform or a password protected file with the password sent separately.
Backup	A backup holds the same sensitive information as the live system and needs the same protection, not a lesser one.
Archiving	Older records kept for legal or historical reasons still need protection matched to their sensitivity.
Deletion	Define retention periods and build in an actual deletion step; encrypted data that should no longer exist is still a risk.

Encryption and AI Systems

Sensitive information can end up in AI prompts, uploaded documents, training or fine tuning data, embeddings, vector databases, conversation histories, application logs and their backups, often within what feels like a single simple task. Before submitting nonprofit data to an AI tool, staff should ask:

Where does this information go, who can access it, how long is it retained, and what controls protect it?

Encryption does not automatically make it appropriate to submit sensitive data to an AI provider. A tool can encrypt data in transit and at rest and still retain it for training, share it with subprocessors, or store it in a jurisdiction with weaker protection than the communities the organization serves. Encryption addresses how data is protected; it does not answer whether it should have been shared with that provider at all.

Key Management, in Plain Language

Who controls the keys

- Most cloud platforms manage keys by default, which is usually fine for general data, but highly sensitive categories deserve a check on whether greater control is possible.

Where keys are stored

- Never alongside the data they protect. A password saved in the same spreadsheet it unlocks defeats the purpose.

Access and rotation

- Limit key and master password access to the smallest necessary group, review it periodically, and rotate credentials after staff turnover.

If a key is compromised

- Have even a basic plan covering who is notified, how quickly credentials change, and how affected systems are checked.

For most small and medium nonprofits, good key management is mostly good password management, done consistently and documented well enough to survive staff turnover.

Encryption by Design: Practical Checklist

- 1. We know which systems and files hold sensitive beneficiary, donor, staff or financial data.
- 2. Our forms and websites that collect sensitive data use secure, encrypted connections (HTTPS).
- 3. Our core storage platforms (cloud drive, database, case management system) encrypt data at rest.
- 4. Our backups are encrypted or stored with equivalent protection, not left as plain unprotected copies.
- 5. Sensitive files are never sent as unprotected email attachments without an alternative safeguard.
- 6. We have a written or agreed retention period for sensitive data, and old data is actually deleted.
- 7. No password or encryption key is stored in the same location as the data it protects.
- 8. Access to keys, admin accounts and master passwords is limited to specific named staff.
- 9. Staff have a simple, known process to follow if a device is lost, stolen or a password is exposed.
- 10. Before using an AI tool with sensitive data, staff know to check where the data goes and how it is retained.

Implementation Roadmap

UNDERSTAND	CONNECT
Immediate	Identify the 3 to 4 systems holding your most sensitive data; confirm whether they encrypt at rest; stop any obviously risky practice such as emailing sensitive spreadsheets unprotected.
First 30 days	Tighten access to sensitive systems and shared passwords; confirm backups are encrypted; draft a retention guideline; brief staff on the AI data question above.
First 90 days	Formalize key and password rotation; build data protection questions into procurement; review AI tools already in use; set a recurring annual review of this checklist.



Section 5

Privacy by Design and the Interactive Checklist

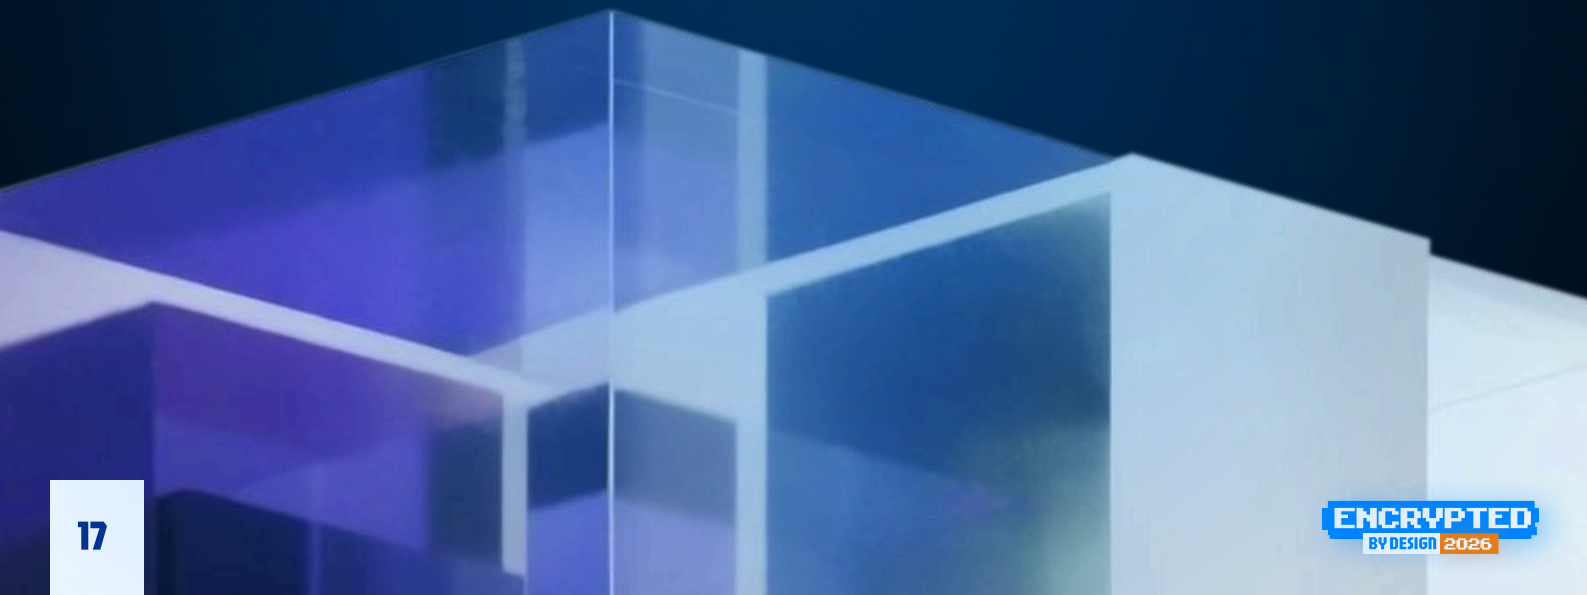
This section is built around the seven Privacy by Design principles, originally articulated by former Ontario Information and Privacy Commissioner Ann Cavoukian and widely referenced in international privacy frameworks. The table below translates each principle into what it means, what an organization should actually do, and a short nonprofit example.

Principle	What it means and why it matters	What to do	Example
1. Proactive, not reactive	Identify privacy risks before a project starts or a tool is adopted, not after something goes wrong.	Ask a basic privacy question before launching any new system or AI tool; give one person clear responsibility for raising concerns early.	Before adopting a mobile data collection app, the team checks where it stores data rather than discovering an issue after rollout.
2. Privacy as the default	People should get reasonable protection automatically, without having to request it.	Collect only what is genuinely needed; set retention periods; use role based access so staff see only what their role requires.	A case system is configured so field staff see only their assigned cases, by default.
3. Privacy embedded into design	Privacy is part of planning a system, like budget and timeline, not a review bolted on afterward.	Include a privacy line item in project planning and procurement; document key decisions so reasoning survives staff turnover.	When designing a donor database, the team decides upfront which fields are truly necessary.
4. Full functionality, positive sum	Privacy does not have to come at the expense of usability or program delivery.	When a privacy concern arises, look for an adjustment that keeps the program goal intact rather than abandoning either one.	Instead of exact addresses, a monitoring team collects location at
5. End to end security	Protection should cover the full lifecycle, collection through deletion, not just one point such as storage.	Map where sensitive data actually goes; apply access controls and pseudonymization consistently; treat sharing and deletion as security steps.	An organization finds staff export beneficiary lists to unprotected spreadsheets and closes that gap as a priority.
6. Visibility and transparency	People should understand, in plain language, how their information is used, including AI involvement.	Provide a short, plain language privacy notice at collection; disclose AI use; give people a contact point for questions.	A cash transfer registration form explains simply what is collected and why, and who to contact.
7. Respect for user privacy	People should have meaningful control over their own information, appropriate to context.	Let people find out what is held about them, correct inaccuracies, and present choices in plain language.	A beneficiary who has moved asks to update her contact details, and staff have a simple process to do it.

Not every element of these principles, particularly rights such as correction or deletion, applies in exactly the same way in every African jurisdiction. National data protection laws vary in scope and enforcement, so organizations operating across multiple countries should confirm what applies where they work, a topic covered further in Section 8.

Data Lifecycle Quick Guide

Stage	Key privacy question
Plan	Do we actually need this data?
Collect	Are we collecting only what is necessary?
Store	Is it appropriately protected?
Use	Are we using it for the stated purpose?
Share	Does the recipient actually need it?
AI use	Is it appropriate to provide this information to the AI system?
Retain	Do we still need it?
Delete	Can we securely remove it when no longer needed?
Review	Have circumstances, risks or technology changed?



Before You Put Nonprofit Data Into an AI Tool

AI Privacy Quick Check

1. Does the information contain personal or sensitive data?
2. Do we actually need to provide the data to the AI system?
3. Can we remove names, identifiers or unnecessary details first?
4. Who operates the AI service?
5. How is submitted information handled and retained?
6. Who can access the information?
7. Have we considered the privacy, contractual and legal implications?

Use the Interactive Privacy by Design Checklist

This playbook is accompanied by a companion Excel workbook, the Privacy by Design Checklist, an interactive self-assessment tool that walks through these same seven principles in about 20 to 25 short questions, scores your organization automatically, and generates a simple action plan. It takes 15 to 30 minutes to complete and is a practical self-assessment tool, not a legal compliance certification.

Companion file: **Privacy by Design Checklist**



Section 6

Threat Modelling for Nonprofits

The goal of threat modelling is not to eliminate every possible risk. It is to help an organization decide where its limited time, people and resources should be focused first. A nonprofit using AI does not need to identify every theoretical threat. It needs to understand what a system touches, who could be affected, where responsibility sits, how a problem would be detected, and what deserves attention first.

Look Beyond the AI Model

A useful threat model examines the whole activity, not only the underlying AI model, by following a simple chain:

Data → AI System → People → Decision → Impact

Stage	Key question
Data	What information enters the system, and how sensitive is it?
AI system	Where is it hosted, who operates it, and what happens to submitted information?
People	Who uses the system, and who could be affected if something goes wrong?
Decision	What can the AI influence, and who can review, challenge or override its output?
Impact	What happens to people, the organization or its mission if the system fails or is misused?

Five Questions for a Practical Threat Model

1. What data does this AI system touch, and how sensitive is it? Classify it as public, internal, sensitive or highly sensitive, and identify which provider receives it.

2. Who could be harmed if this data or decision were exposed, wrong or misused? Consider beneficiaries, staff, donors and partners, not only the organization itself.

3. Who reviews the AI's output before it affects a person or a decision? A human somewhere in the workflow is not automatically effective oversight; someone needs clear authority to challenge or override the output.

4. How would we find out if something went wrong, and how quickly? Consider both a data or security failure and an AI output failure, such as an incorrect summary.

5. Given our time and resources, what is the one thing we should fix first? The answer should be specific: what to address, and why.

A numerical risk score can create a false sense of precision that a small organization may not have the evidence to support. The final answer to question five is the most important output of the exercise.

Scale the Response to the Organization

If the organization cannot do this	Consider this instead
Establish a dedicated AI governance committee	Assign a named decision owner within an existing team
Run a formal external threat modelling exercise	Use the five questions above during project planning
Deploy an enterprise AI monitoring platform	Maintain a simple inventory of AI systems and review it periodically
Hire dedicated AI security specialists	Establish a clear internal owner and seek targeted external advice when needed

These alternatives are not substitutes for stronger controls where those are genuinely necessary. They are ways to establish basic governance and risk awareness when resources are constrained, consistent with the NIST AI Risk Management Framework's emphasis on governance and human oversight that is proportionate to an organization's size and maturity.

Section 7

Secure AI Tools and Open Source Resources

African nonprofit technology teams operate under real constraints that shape which tools are actually usable, not just theoretically best. Connectivity is not guaranteed. Sub Saharan Africa lost an estimated 1.6 billion dollars to government ordered internet shutdowns in 2024 alone. Budgets are thin and staff are often generalists rather than dedicated security engineers. The recommendations below are chosen for that reality: free or low cost, open source or independently audited, and deployable without specialist staff.

An important caution before the list: an open source tool is not automatically secure, and a tool labelled secure is not secure for every use. Evaluate the whole system, meaning the software, its configuration, who administers it and how, not the license alone. Strong protection comes from the interaction between technology, policy, staff practice and institutional routine, not from any single product.

Recommended Tools by Category

Encryption for files and devices

VeraCrypt is free, open source disk and file encryption software, independently audited in 2016. Cryptomator offers lighter file encryption before syncing to cloud storage such as Google Workspace or Microsoft 365 nonprofit plans. Together they give field staff carrying beneficiary data a practical, low cost layer of protection.

Website and transport encryption

Let's Encrypt provides free, automated TLS certificates, and OpenSSL supports certificate handling and diagnostics. For most small teams, enabling automatic certificate renewal through a host or Certbot addresses the most common transport layer risk, making HTTPS a baseline expectation rather than a premium feature.

Secure communications

Signal is open source, end to end encrypted messaging widely used by civil society organizations because it does not depend on advertising or selling user data; Proton Mail offers encrypted email with a free tier for correspondence with partners who do not use encrypted messaging. Neither solves an internet shutdown, so organizations should also keep encrypted offline backups and agree in advance how staff will communicate without internet access.

Password and secrets management

Bitwarden is an open source password manager independently audited by multiple security firms, with a free tier and a self hosted option. KeePassXC is a fully offline option recommended by the Electronic Frontier Foundation, useful where connectivity is unreliable.

Privacy enhancing technologies

Tools such as Tor Browser can help maintain access during shutdowns or platform restrictions, and privacy respecting search engines such as DuckDuckGo reduce tracking during routine research. Independent resources such as Privacy Guides track how these products and their security records change over time.

For network security testing or formal auditing, engage a qualified technical specialist or trusted partner rather than relying on general purpose scanning tools without proper authorization and documented scope; this is an area where misuse carries real legal and ethical risk.

Secure Use of AI Tools

The African Union's Continental AI Strategy, endorsed in 2024, names cybersecurity risk, privacy and weak national AI governance capacity as key challenges to AI adoption across the continent. The practical lesson for nonprofits is to set internal rules before problems occur. Before using an AI tool with sensitive information, check what happens to information entered into it, whether it is used to train the model, how long it is kept, where it is stored, whether it could be transferred to another country, and whether the organization has an agreement with the provider that protects its data. Specific products will change faster than this playbook can be updated, but these underlying questions will remain useful. Organizations without the internal expertise to answer them should seek advice from regional digital rights organizations before adopting a tool across the organization.

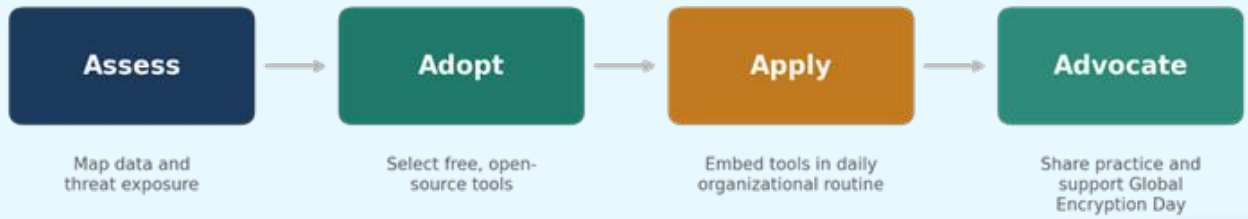
Figure 7.1 A Layered Encryption Model for Nonprofit Technology Organizations



Tool Summary

Function	Recommended tool(s)	Note for a resource constrained nonprofit
Website encryption	Let's Encrypt with Certbot; OpenSSL	Automated renewal removes recurring cost and the risk of an expired certificate.
File and disk encryption	VeraCrypt	Useful for laptops and removable media carried by field staff.
Email and document protection	GnuPG with Kleopatra or Mailvelope	Supports recipient specific encryption, but needs sound key management practice.
Password management	Bitwarden or KeePassXC	Replaces shared plaintext password lists with unique credentials per system.
Secure messaging	Signal; Session where metadata exposure is a priority	Needs organizational rules on what may be shared, not just the tool itself.

Figure 7.3 A Simple Adoption Pathway for Nonprofit Teams



Getting Started, in Order

1. Confirm every organizational website has a valid, auto renewing certificate.
2. Retire shared password spreadsheets and move the team to Bitwarden or KeePassXC within a defined period.
3. Identify which communications genuinely require end to end encryption and standardize one approved channel.
4. Make VeraCrypt containers, or an equivalent, the default for sensitive fieldwork data, not an optional extra.
5. Review the toolset and procedures at least annually, since platforms and regulation continue to change.

None of these tools is a complete solution on its own. Their value lies in combination: transport encryption protects information in motion, device encryption protects it at rest, secure messaging protects sensitive conversations, and password managers strengthen access control. Security becomes meaningful when it is embedded in routine practice rather than treated as a one time purchase. For further African civil society specific programs and support, see Appendix C.



Section 8

African Policy and Regulatory Considerations

Understanding the legal environment an organization operates in is part of protecting the people it serves. This section gives a practical overview, not a legal encyclopedia, and it is not a substitute for jurisdiction specific legal advice.

Look Beyond the AI Model

A useful threat model examines the whole activity, not only the underlying AI model, by following a simple chain:

Three Different Things

Law and regulatory requirements

Binding obligations set by a country's legislation and enforced by a designated authority. These vary by country and must be confirmed locally.

International good practice

Frameworks such as the OECD AI Principles, UNESCO's Recommendation on the Ethics of AI, or the NIST Privacy and AI Risk Management Frameworks. These are widely respected references, but they are not automatically law in any given African country.

Recommended organizational practice

The practical guidance in this playbook itself: sensible, defensible steps that go beyond the legal minimum in many places, useful regardless of what the law strictly requires.

By recent counts, most African Union member states have enacted some form of data protection law, though far fewer have a fully functioning enforcement authority, and the pace and rigor of enforcement varies considerably from country to country. A handful of illustrative examples:

The Regulatory Landscape

By recent counts, most African Union member states have enacted some form of data protection law, though far fewer have a fully functioning enforcement authority, and the pace and rigor of enforcement varies considerably from country to country. A handful of illustrative examples:

- **Nigeria:** the Data Protection Act 2023, overseen by the Nigeria Data Protection Commission, with breach notification duties for organizations that meet applicable thresholds.
- **Kenya:** the Data Protection Act 2019, overseen by the Office of the Data Protection Commissioner.
- **South Africa:** the Protection of Personal Information Act (POPIA), overseen by the Information Regulator.
- **Ghana:** the Data Protection Act 2012, overseen by the Data Protection Commission.
- **Senegal, Côte d'Ivoire and Cameroon:** representative Francophone examples, each with its own supervisory authority and its own pace of development; some Lusophone countries' frameworks remain in draft form.

At the continental level, the African Union's Malabo Convention on Cyber Security and Personal Data Protection, adopted in 2014, entered into force in June 2023 after reaching its 15th ratification, and by recent counts has been ratified by around 16 of the African Union's 55 member states. It is therefore an important signal of continental direction, but it is not yet binding on every country, and organizations should check whether their country of operation has ratified and domesticated it. The African Union's Continental AI Strategy, endorsed in 2024, separately identifies cybersecurity, privacy and weak national AI governance capacity as continental priorities, and calls for stronger governance capacity in member states.

Cross Border Data and AI

AI tools and cloud services frequently move data across borders as a normal part of how they operate. An organization working with an AI provider, or with partners in another country, should understand where its data is actually hosted and processed, and what obligations that triggers under the law that applies to it. This is especially relevant for organizations operating across multiple African jurisdictions with different, and sometimes overlapping, legal regimes, and for Francophone and Lusophone countries where established digital rights and legal resources have historically been less available in local languages than in Anglophone Africa.

AI introduces additional considerations beyond conventional data protection, including transparency about when AI is used, meaningful human oversight of AI influenced decisions, and clear accountability when something goes wrong, themes explored further in Sections 6 and 10.

Who could be harmed if this data or decision were exposed, wrong or misused? Consider beneficiaries, staff, donors and partners, not only the organization itself.

Who reviews the AI's output before it affects a person or a decision? A human somewhere in the workflow is not automatically effective oversight; someone needs clear authority to challenge or override the output.

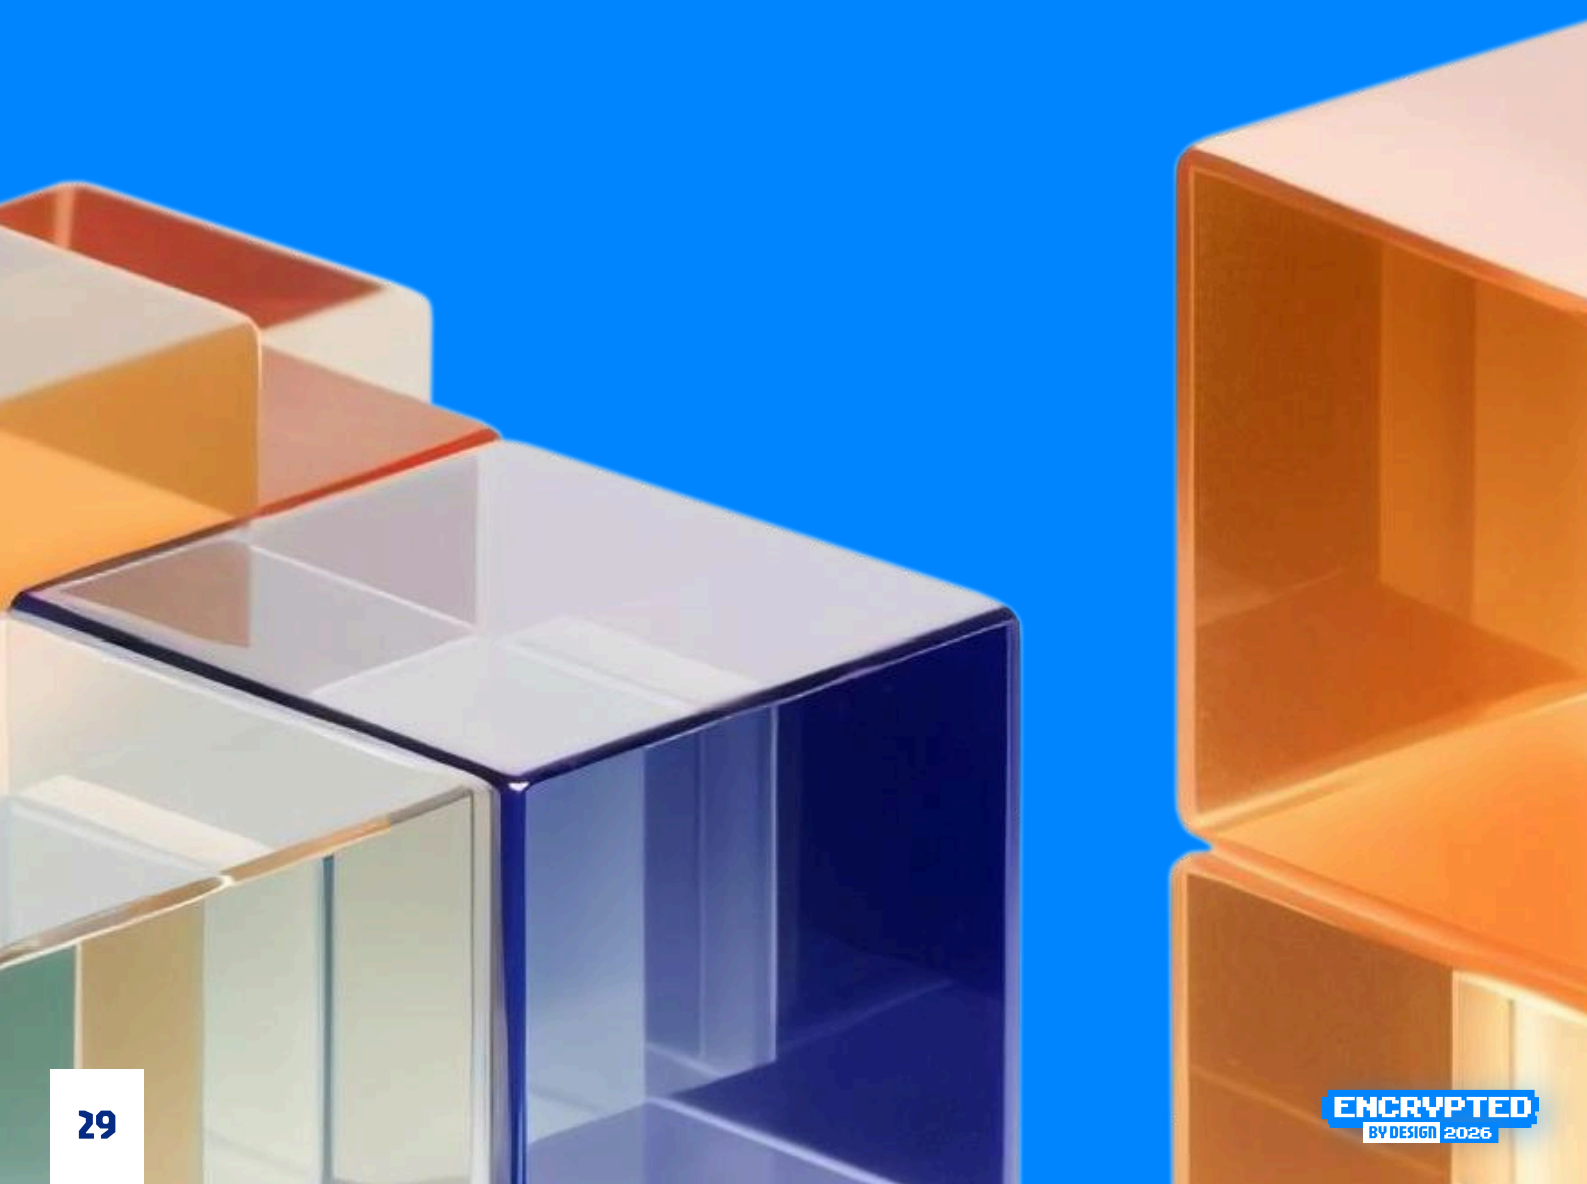
How would we find out if something went wrong, and how quickly? Consider both a data or security failure and an AI output failure, such as an incorrect summary.

Given our time and resources, what is the one thing we should fix first? The answer should be specific: what to address, and why.

Encryption and Compliance

Many data protection laws expect organizations to apply appropriate technical and organizational measures to protect personal data, and encryption is commonly recognized as one such measure. Encryption alone, however, does not guarantee legal compliance. Meeting a specific country's requirements depends on the full set of obligations that law imposes, not on any single technical control.

Organizations should obtain jurisdiction specific legal advice where necessary, particularly when operating across multiple countries, handling especially sensitive data, or facing a specific regulatory question. This playbook is intended to support good practice, not to serve as a substitute for that advice.



Section 9

Organizational Encryption Action Plan

You do not need to be a cybersecurity expert to complete this plan. Its purpose is to help your organization identify what needs protection, prioritize the biggest risks, and take practical steps forward. Appendix A provides a one page checklist to use alongside this process.

Step 1: Identify What You Need to Protect

Think about the information your organization collects, creates or stores: beneficiary and community information, staff and volunteer records, financial and donor data, legal or case information, research data, project and advocacy plans, and passwords or account credentials. For each category, ask: if the wrong person got access to this, could someone be harmed, or could the organization be seriously affected? If yes, mark it High Priority.

Step 2: Identify Your Biggest Risks

You do not need to predict every possible attack. Start with common situations: a laptop or phone is lost or stolen, a password is guessed or stolen, an email or cloud account is compromised, sensitive information is sent to the wrong person, staff use an unsafe communication channel, a former staff member still has access, or information is lost because there is no backup. For each risk, weigh how serious the consequences would be against how likely it is to happen, and start with risks that score high on both.

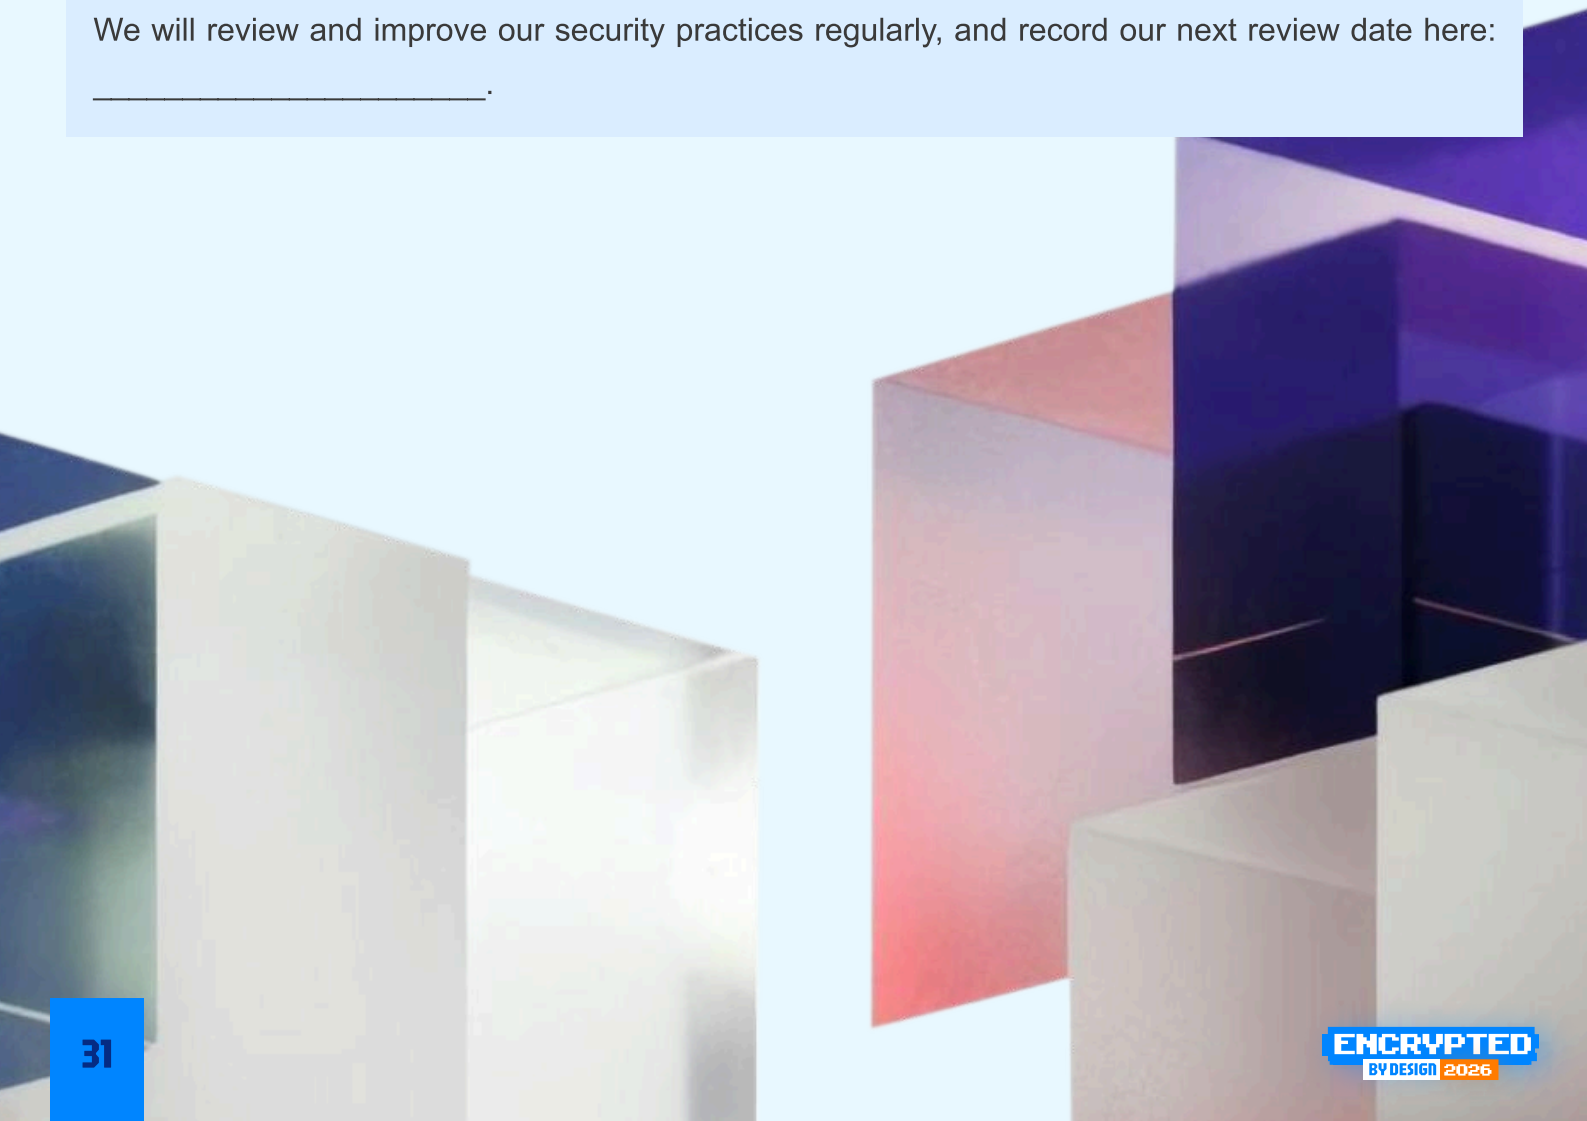
Step 3: Assign Priorities and Act

Use Appendix A to check current practice across devices, accounts, communication, files and backups, and people. For each gap identified, record a specific action, a responsible person, and a target date using the roadmap below.

Timeframe	Priority actions	Responsible	Status
Immediate	Protect critical accounts with strong passwords and MFA; identify and stop the riskiest current practice.		
First 30 days	Encrypt devices; strengthen password practices; retire shared password spreadsheets.		
First 90 days	Formalize backup, retention and key management practices; review AI tools in active use.		

Our Commitment

We will protect information based on the risks our organization and the people we work with face. We will review and improve our security practices regularly, and record our next review date here: _____.



Section 10

AI Safety and Encryption Self Assessment

This self assessment is distinct from the Privacy by Design Checklist in Section 5. The Privacy by Design Checklist focuses on embedding privacy into organizational processes and systems generally. This tool focuses specifically on organizational readiness to deploy and use AI securely and responsibly. The two are complementary and intentionally avoid repeating the same questions.

It draws on established frameworks including the NIST AI Risk Management Framework, the CIS Critical Security Controls, NIST's digital identity guidelines on multi factor authentication, the African Union's Continental AI Strategy, and applicable national data protection law. Complete it as a team, ideally with someone responsible for IT or data, a program lead and a senior decision maker, before introducing an AI tool organization wide, and review it at least annually or after any major change.

Scoring

Score	Meaning
0	Not addressed, or the organization is not aware of the issue
1	Some informal or inconsistent practice exists
2	A written policy or process exists
3	The policy is implemented, monitored and reviewed regularly

If the team genuinely does not know an answer, record "not sure" rather than guessing. This itself identifies where the organization needs more information or training.

The Ten Assessment Areas

Score	Meaning	Score
1. Sensitive data inventory	Is there a current, written inventory of the personal, financial, health or beneficiary data the organization collects?	
2. AI tools in use	Is there a current list of every AI tool in active use, including those adopted informally by individual staff?	
3. Encryption at rest	Is sensitive data on servers, shared drives, laptops and removable media encrypted?	
4. Encryption in transit	Is sensitive data encrypted whenever it moves between systems, including to and from any AI tool?	
5. Secure communication	Are end to end encrypted channels used for sensitive communication, with a documented fallback for connectivity disruptions?	
6. Access and identity	Is multi factor authentication enforced on systems exposing personal data, on a least privilege basis?	
7. Third party and AI vendor risk	For each AI tool, has the organization confirmed the vendor's retention, training use, hosting location and cross border transfer implications?	
8. Incident response	Is there a written incident response plan consistent with applicable breach notification requirements?	
9. AI governance	Is there a leadership approved policy setting rules for what data may be entered into AI tools?	
10. Staff training	Have staff been trained in the past twelve months on encryption basics, phishing recognition and safe AI use?	

Interpreting the Total

- 0 to 10, foundational gaps: core protections are largely absent. Prioritize encryption and multi factor authentication (areas 3, 4 and 6) immediately, and check registration requirements with the applicable data protection authority.
- 11 to 20, developing: practices exist but are inconsistent or undocumented. Convert ad hoc practice into written, leadership approved policy, prioritizing vendor risk and incident response (areas 7 and 8).
- 21 to 30, maturing: policies exist and are largely enforced. Focus on review cycles, staff refreshers, and monitoring regulatory developments that could affect cross border data flows.

Use Section 9 to turn any gaps identified here into a concrete action plan with owners and deadlines.

Case Study: AI in Practice

The two scenarios below are illustrative, composite examples built to demonstrate the threat modelling approach from Section 6 in practice. They are not drawn from any specific organization and do not describe actual events.

Case Study A: Sensitive Beneficiary Data

Situation

A nonprofit supporting vulnerable beneficiaries introduced a cloud based AI tool to summarize case intake notes, expecting to free up caseworker time for direct support.

Risk

Intake information sent to the tool included names, health and welfare details, location, financial circumstances and safeguarding concerns, all processed by a third party service the organization had not fully assessed.

What the organization did.

Applying the five threat modelling questions, the team identified that highly sensitive data was entering a third party system without a clear understanding of retention or access, and that inaccurate AI summaries could be acted on without detection.

What changed

The organization minimized identifying details before submission where the task allowed it, restricted access to the tool, and required caseworkers to check every AI generated summary against the original notes before acting on it.

Key lesson

A human reviewer in the workflow does not by itself make an AI process safe. The organization must understand what data enters the system, where it goes, and what controls are necessary before it is trusted with sensitive information.

Case Study: AI in Practice

Situation

A nonprofit with more applicants than it could support introduced an AI system to recommend which cases should receive priority, intending staff to review every recommendation.

Risk.

No individual or role was formally responsible for approving, challenging or overriding the AI's recommendations, so a plausible but incorrect recommendation could pass through unreviewed under time pressure.

What the organization did

Recognizing that reviewing an output is not the same as owning a decision, the organization named a specific decision owner with explicit authority to question, challenge and override AI recommendations.

What changed

. Program staff were told clearly that AI output was a recommendation, not a determination, and a process was set up to record and escalate patterns of error.

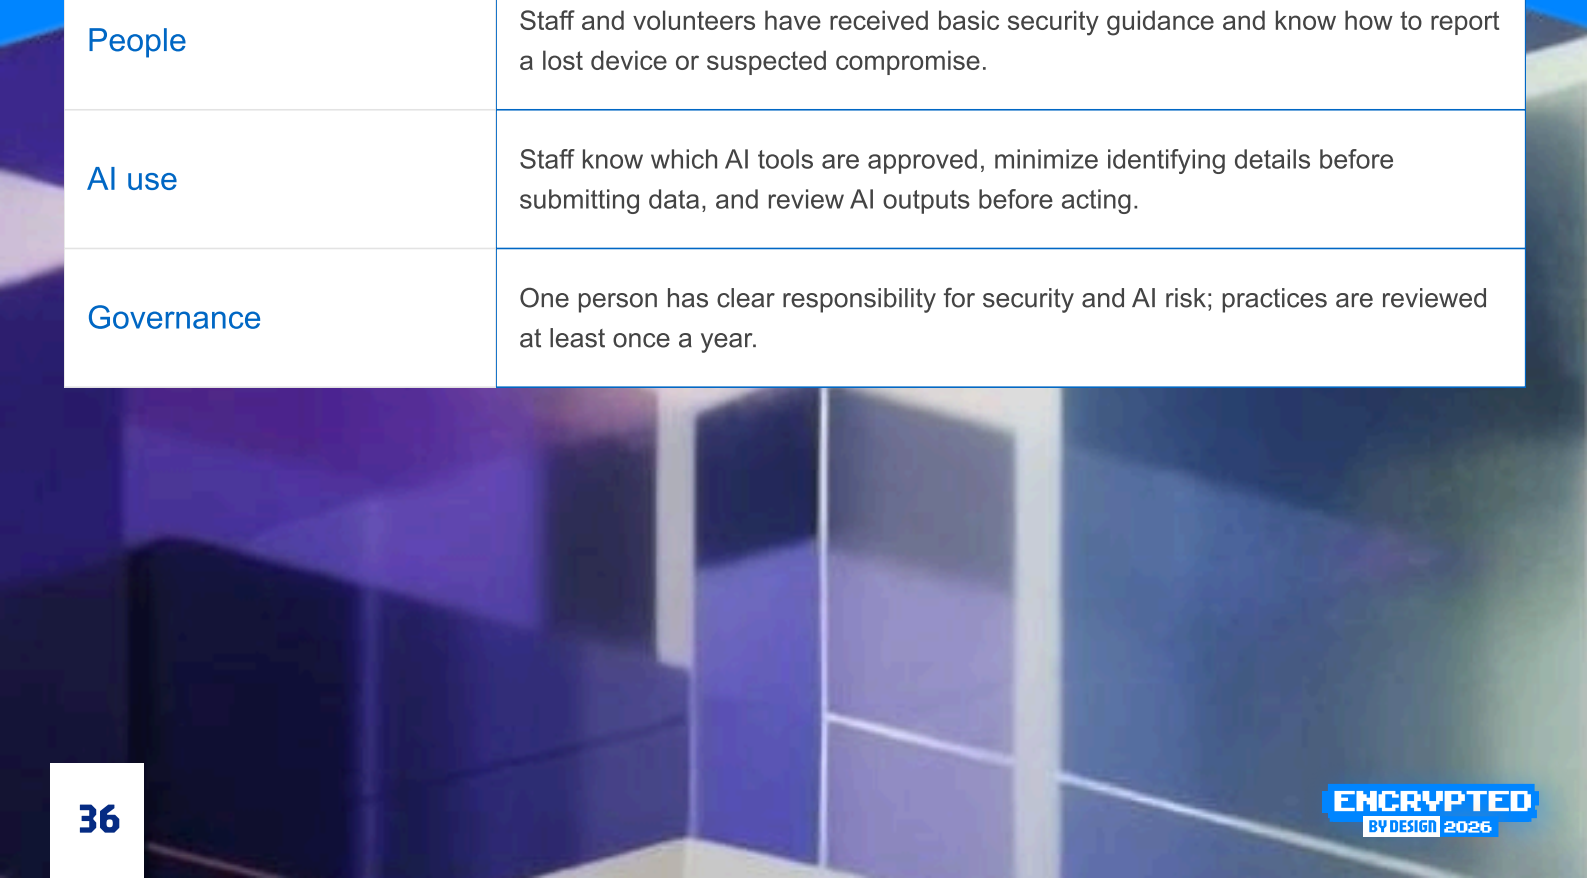
Key lesson

Clicking approve is not automatically governance. Authority to challenge and override is. Establishing who owns a decision can matter as much as assessing the underlying technology.

Appendix A: Quick Reference Checklist

A one page summary of the most important actions across this playbook. Mark each item Yes, No or In Progress.

Area	Action
Devices	Organizational laptops and phones are encrypted, use a strong password or PIN, lock automatically, and stay updated.
Accounts	Important accounts use strong, unique passwords and multi factor authentication; access is removed when someone leaves.
Communication	Sensitive conversations use an approved, encrypted channel; recipients are checked before sending sensitive information.
Files and backups	Sensitive files are stored securely with limited access; backups exist and are protected, not left as unprotected copies.
People	Staff and volunteers have received basic security guidance and know how to report a lost device or suspected compromise.
AI use	Staff know which AI tools are approved, minimize identifying details before submitting data, and review AI outputs before acting.
Governance	One person has clear responsibility for security and AI risk; practices are reviewed at least once a year.



Appendix B: Key Terms

Encryption
transforming readable information into a form that cannot be understood without the correct key.

End to end encryption
encryption that protects a communication from the sender's device to the recipient's device, so intermediaries cannot read it.

Encryption key
the piece of information that locks or unlocks encrypted data; protecting it is as important as encrypting the data itself.

Personal data
information that identifies or could identify a specific individual.

Sensitive data
personal data whose exposure could cause significant harm, such as health, financial or beneficiary case information.

Pseudonymization
replacing identifying details with a code or reference, so the data cannot be attributed to a person without additional information kept separately. Pseudonymized data can usually be re-linked to a person and is not the same as anonymized data.

Encryption key
the piece of information that locks or unlocks encrypted data; protecting it is as important as encrypting the data itself.

Anonymization
Processing data so thoroughly that the individual it describes can no longer be identified by any reasonably available means; unlike pseudonymization, this is generally not reversible.

AI system
Software that uses models to generate, classify, summarize, translate or recommend based on data it is given.

Threat modelling
A structured process for identifying what needs protection, who could cause harm, and what to prioritize.

Privacy by Design
An approach that builds privacy protection into systems and processes from the start, rather than adding it afterward



Appendix C: Useful Resources

Global Encryption Coalition:
Policy resources and practical guides on why encryption matters and how it works, globalencryption.org.

Internet Society
Policy briefs on encryption as a default safeguard for data in transit and at rest, internetsociety.org.

CIPESA (Collaboration on International ICT Policy for East and Southern Africa)
Digital rights research and security training across the region, cipesa.org.

Paradigm Initiative
Pan African digital rights nonprofit offering digital security training and policy advocacy.

Africa Digital Rights Fund
Funding and networks connecting African civil society organizations to digital security expertise.

Cloudflare Project Galileo
Free protection against distributed denial of service and web application attacks for eligible nonprofits.

Security First, Umbrella app:
A free, open source guide to digital and physical security for civil society groups in higher risk environments.

Privacy Guides:
Independently maintained comparisons of privacy focused browsers, search engines and other tools, as products change over time.

References

This playbook draws on established privacy, security and AI governance frameworks. References are provided for further reading and to support the claims made in this document, not to suggest that any single framework is automatically binding law in a given country. Where legal obligations are discussed, confirm the specific requirements of the jurisdiction in which your organization operates.

- Tabassi, E. (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0). NIST AI 100-1. National Institute of Standards and Technology.
- National Institute of Standards and Technology. NIST Generative AI Profile (NIST AI 600-1).
- National Institute of Standards and Technology. NIST Privacy Framework: A Tool for Improving Privacy Through Enterprise Risk Management.
- National Institute of Standards and Technology. SP 800-63B, Digital Identity Guidelines.
- Center for Internet Security. CIS Critical Security Controls v8, Control 3: Data Protection.
- Information Commissioner's Office (United Kingdom). Guidance on Data Protection by Design and Default.
- Organisation for Economic Co-operation and Development. OECD AI Principles, adopted 2019, updated 2024.
- UNESCO. Recommendation on the Ethics of Artificial Intelligence, adopted 2021.
- OWASP Foundation. (2025). OWASP Top 10 for Large Language Model Applications.
- African Union. Continental Artificial Intelligence Strategy, endorsed 2024.
- African Union. Malabo Convention on Cyber Security and Personal Data Protection, adopted 2014, entered into force June 2023.
- Nigeria Data Protection Commission. Nigeria Data Protection Act 2023 and General Application and Implementation Directive 2025.
- Kenya. Data Protection Act, 2019.
- South Africa. Protection of Personal Information Act (POPIA), 2013.
- Ghana. Data Protection Act, 2012 (Act 843).
- Zimbabwe. Cyber and Data Protection Act [Chapter 12:07], 2021.
- Internet Society. Policy brief: Encryption; and Policy brief: Solving Crime Without Breaking Encryption.
- Global Encryption Coalition. Global Encryption Day statements and resource library.
- Internet Security Research Group / Let's Encrypt. Ten years of Let's Encrypt certificates (2025).
- VeraCrypt Project, GnuPG Project, Bitwarden, KeePassXC and Signal Foundation. Official product documentation and security disclosures.
- Cavoukian, A. Privacy by Design: The 7 Foundational Principles.
- CIPESA and Paradigm Initiative. Digital security resources for African civil society organizations.