Policy Hackathon – RFP template

Intelligent Government: Reimagining Civic Infrastructure

Instructions: Please present your proposal - Either design your own 2-pager to showcase your idea or fill this template out with your final solution. Once completed, **please email it to Whitney Westbrook (wwestbrook@chathamhouse.org) by Monday, March 18 at 9:00 am GMT.**

TARGET MINISTRY	ESTIMATED VALUE	VENDOR / TEAM NAME
Ministry of Foreign Affairs and International Development (MFA)	*$6-10 million*	*Aegis*

EXECUTIVE SUMMARY

The Ministry of Foreign Affairs faces growing cyber threats, including espionage, malicious attacks, and data breaches that risk exposing sensitive diplomatic information and agent identities. This proposal introduces Guardian Angel, an AI-assisted system that analyses employee access patterns to detect behavioural irregularities and potential security risks. Built on existing secure analytics platforms already used within government systems, it enables cost-effective and rapid deployment. By identifying both internal and external threats early and alerting human security teams, Guardian Angel helps prevent breaches while strengthening Valdoria's cyber resilience and safeguarding its democratic values and data privacy through its adherence to the Paneuro AI Act.

PROBLEM & OPPORTUNITY

The Ministry of Foreign Affairs manages highly sensitive data, positioning it as a prime target for cyber threats. Existing monitoring systems are largely reactive, limiting the ability to detect latent/emerging threats.

Manual monitoring of metadata consumes significant staff time and stretches limited cybersecurity resources dependent on data scale. This results in delayed detections, increasing the risk of high-cost breaches.

AI behavioural analytic tool that allows for proactive, rather than reactive, threat detection by identifying anomalies in staff metadata. This greatly enhances cyber resilience in a cost-effective and scalable manner.

EVALUATION CRITERIA CHECKLIST – Ensure your proposal addresses:

1. System design & differentiation between pilot and scaling phases

PROPOSED AI SYSTEM

- **System name:** *Guardian Angel*
- **AI capabilities used:** *Behavioural analytics using existing, pre-trained machine learning models; entity behaviour analysis; pattern recognition; and risk scoring.*
- **How it works:**
  Guardian Angel will spend the first 90 days learning from employee metadata only such as login times, device usage, and file access patterns. It will not analyse content – only behavioural patterns. If significant deviations from these patterns occur, a risk score will be generated and an alert sent to Valdoria's national cyber security team (A separate, intelligence department) for human review, ensuring accuracy and minimising false positives.

2. Cost and value for money
3. Data security & privacy (Valdorian Data Protection Act compliance)
4. Transparency and digital inclusion
5. Implementation timeline & citizen support
6. Geopolitical implications: security, resilience, sovereignty

IMPLEMENTATION TIMELINE

Outline your rollout plan across the full 5-year contract.

Phase	Description
Pilot (Yr 1–2)	Guardian Angel will launch at MFA Headquarters, covering system access logs, secure file systems, and authentication activity. All HQ staff (around 400 people) will have individual baselines established over the first 90 days. Alerts will begin in a controlled phase by Month 4, with full operational capability achieved by Month 6. Throughout the pilot, the system will be continuously refined to reduce false positives and improve accuracy. By the end of Year 2, an independent audit will assess performance and inform the decision to expand, with the aim of significantly reducing detection time and improving early identification of security risks by about 50%.
Scale (Yr 3–5)	Guardian Angel will expand to a select number of high-risk Valdorian embassies in Year 3 such as China and Russia. Based on performance and audit results, gradual expansion to additional locations will occur where risk levels justify deployment. This could include neutral-based embassies, international development offices, and international organisations.
Key Milestones	**Year 1:** Alerts operational at MFA Headquarters (initial deployment phase) **Year 2:** Independent audit completed and decision made on phased expansion **Year 3:** Rollout begins at a select number of high-risk embassies **Year 5:** Deployment extended to priority international locations, with further expansion planned based on risk and performance

PRICING & VALUE FOR MONEY

Break down your total contract cost and justify the investment.

Cost Item	Amount (USD)
Development & Integration	*$1.5–3M*
Pilot Deployment (Yr 1–2)	*$1.5–2M*
Scaling Phase (Yr 3–5)	*$2–3M*
Maintenance & Support	*$1–2M*

Total Contract Value	**$6–10 million**

These figures are justifiable by their ability to significantly reduce the risk of costly cyber breaches, which can result in severe financial, operational, and diplomatic damage. By using existing technologies and a targeted, phased rollout, Guardian Angel delivers high-impact security improvements in a cost-effective and scalable way for Valdoria. Figures are broad because they are dependent on success and technological improvement. A financial oversight committee held by the MFA in cooperation with the Treasury will be held every year to make sure funds aren't misappropriated or wasted.

GOVERNANCE & SAFEGUARDS

Explain how your system upholds accountability, protects citizens, and maintains public trust:

- **Human oversight:** *Every alert must be reviewed by a named Valdoria's national cyber security member before any response is initiated. More serious alerts will require sign-off from two personnel. This directly reflects the Paneuro AI Act - Article 14, which affirms that high-risk AI systems should be under human control. Final decisions are made by human staff to prevent false flags.*
- **Bias & fairness:** *Every individual has their own baseline, rather than pre-existing baselines for demographic groups. This will stop profiling by nationality, religion, ethnicity or gender. The model will be audited every six months to check for patterns in false positives or issues with baselines. Any employee can challenge an alert related to them after action has occurred. Ongoing monitoring to cease discriminatory outcomes will be based on Paneuro AI Act.*
- **Data privacy:** *Guardian Angel analyses metadata, including login times, file access patterns and the volumes of data transfer. It does not read message contents. All data is stored specifically on Valdorian government systems to prevent data being potentially sold to advertisers. It will be encrypted and only accessible by the national cyber security team.*
- **Transparency**. *Every alert has an explanation in plain language over the cause of the trigger and which data points were involved, so analysts can understand the reasoning behind every flag. All system activity will be stored by the national cyber security team, distinct and neutral, to prevent any bias. For example, a member of the MFA might not take a flag seriously due to personal connections with a flagged MFA member. The MFA will release an annual anonymised public record available to members of Parliament and select committees, so representatives and leaders can see how the system is used. Only the national cyber security team has access to non-anonymised specific data.*
- **Digital inclusion:** *Guardian Angel operates exclusively within MFA's internal infrastructure. Due to the sensitive nature of the MFA, citizens won't have access to its details. All staff that are subject to monitoring will be clearly informed and will have to sign a T&C to agree to have their data monitored by Guardian Angel. Staff should also be given a group presentation, and then individual meetings for new staff explaining exactly what Guardian Angel does.*

GEOPOLITICAL & RISK CONSIDERATIONS

- **Technology sovereignty:** Guardian Angel will prioritise deployment on secure, government-approved infrastructure hosted within Valdoria, utilising trusted commercial platforms to reduce costs. Contracts will ensure that the Valdorian state retains full ownership and control of all sensitive data, with strict limitations on reuse beyond the Ministry of Foreign Affairs. Partnerships will be limited to vetted domestic providers with established records of working with government systems.

- **Supply-chain resilience:** Due to ongoing global trade wars and geopolitical instability, Guardian Angel will minimise reliance on external suppliers by using widely available, vendor-supported technologies and ensuring interoperability across systems. Where external providers are required, Valdoria will prioritise diversified and trusted partners to reduce dependency risks. The system will also be designed to function independently of any single external provider, ensuring continuity in the event of geopolitical disruptions. We intend to utilise Valdoria's own mineral wealth, using resources such as copper, lithium and cobalt, if we have any issues producing these minerals, we will import them from friendly nations.
- **Security:** Guardian Angel itself may become a target for cyber attacks or adversarial manipulation, including attempts to evade detection or generate false alerts. To mitigate this, the system will incorporate regular security testing, access controls, and human oversight in all high-risk decisions. Models will be continuously monitored and updated to reduce bias and adversarial exploitation, while strict access controls will ensure that only authorised personnel can view sensitive outputs.

WHY VALDORIA SHOULD CHOOSE US

Guardian Angel by Aegis stands out by combining advanced behavioural analytics with a cost-efficient, implementation-ready design. Unlike competitors that rely on expensive, fully bespoke AI systems or broad surveillance approaches, this proposal leverages existing platforms, focuses on metadata rather than intrusive content monitoring, and uses a phased, risk-based rollout. This allows Valdoria to achieve high-impact cybersecurity improvements quickly, affordably, and in line with data privacy standards, making it both more practical and more scalable than alternative solutions.