SPARK PLATFORM · REFERENCE ARCHITECTURE · NATIONAL PROGRAMME · OCTOBER 2026

National implementation: how the Spark Platform connects across the entities

One sovereign core at the agency, read-only connections to every entity, signed evidence back. No agent on any endpoint, no internet path into the core.

Spark Software

No agent on any endpoint · No internet path into the core · Every decision signed, verifiable offline

WHO SIGNS IN

- **Minister and leadership** — level 1 report, signed
- **Head of Cyber Security** — level 2 report, directives, incidents
- **Analysts and entity liaisons** — level 3; up to 200 per node; entity-scoped views

Entities keep their SIEM, EDR, SOAR, firewalls and identity: nothing is replaced

HTTPS over the government private network · sign-in clearance: eight checks · roles admin, analyst, viewer

OUTSIDE — Internet and public surfaces

Entity public surfaces (1)
Domains and DNS (DNSSEC), mail records (MX, SPF, DKIM, DMARC, MTA-STS, TLS-RPT), TLS and HTTP, public web applications and AI gateways. Checked from outside, as an attacker would see them.

Entity cloud accounts
AWS: the read-only audit script runs inside the entity's own account; only its output is exported to Spark. Azure and GCP: roadmap.

External sources
Exploited-vulnerability catalogue as a data file with each release; threat intelligence as STIX 2.1 bundles from the national CERT or a commercial feed; Spark release archives, checksum-verified (signed releases: pilot).

What is looked for from outside
Exposed versions on the exploited list, missing or weak DMARC and SPF, expiring TLS, exposed admin panels, DNS that can be forged, AI gateways that leak their system prompt or follow injected instructions.

checks ←

bundles →

DMZ · CONNECTED — Connected relay node

Live public checks (1)
The scheduler assesses every registered entity domain: DNS, TLS, HTTP, mail records, DNSSEC state. Dynamic testing (OWASP, AI Fuzzer) only against targets in the authorised-targets registry, with written authorisation on file.

Intake
Intelligence bundles, catalogue updates and release archives are received here, never inside the core. Checksums and signatures are verified before anything moves on.

Secure drop
Entities outside the government network upload read-only exports over HTTPS with mutual TLS. Nothing is kept on the relay beyond the transfer window.

Sign and hand over (1)
Results, bundles and exports are signed (Ed25519) and carried into the core one way: a data diode where one exists, a verified import where not. The target list reaches the relay as a signed file by the same route.

No inbound ports anywhere: the relay only opens outbound connections, and the core never accepts a connection from the relay side. The one-way import is the only crossing.

SOVEREIGN DATA CENTRE · NO INTERNET PATH — NCSA national core

Load balancer and TLS
The only entry from the government network; certificates from the agency's own CA.

Application node 1
Spark Enterprise Edition, stateless; 32 cores, 128 GB; UI, API, scheduler.

Application node 2
Identical peer behind the balancer; active-passive today, active-active: roadmap.

Identity: local accounts, authenticator app and passkeys today · SAML, OpenID Connect and SCIM: roadmap, one release cycle

APPLICATIONS · EVERY ACTION ON THE SIGNED AUDIT CHAIN (2) (3)

- **National Cyber Command** — National picture, threat index, directives, incidents
- **Spark SHIELD** — Mail and DNS checks, warnings with SLA, exploited-CVE match
- **Domain AI Scanner** — External attack surface per entity, graded, re-assessed
- **OWASP and Compliance** — Dynamic web testing of authorised targets, mapped to controls
- **AI Fuzzer** — Campaigns against authorised AI gateways; OWASP LLM, ATLAS
- **Cloud Health Check** — AWS configuration review from the exported audit output
- **Infrastructure Monitor** — Flows, DNS and assets from exports; posture per entity
- **Adversary Elimination** — Guided hunt and control changes, run with the entity

SRE-1, the Spark AI model
A reasoning layer computes every answer; a grounding check rejects anything not in the facts; a sovereign language model on the GPU node only phrases it. No tools, no network.

Spark Learn
Classifiers and baselines trained on the platform's own data (domain risk, unsafe replies, behaviour). Models never leave the deployment.

- **PostgreSQL** — Primary and synchronous standby (NVMe); every row scoped by organisation
- **Object storage** (7) — Evidence bundles, signed reports, backups; restore tested
- **Signed audit chain** — Ed25519 hash chain over every result and decision; verified at sign-in and offline
- **GPU, scheduler** (6) — Language models on the host; continuous re-assessment; Prometheus metrics

INTERFACES · THE ONLY WAYS IN AND OUT

- **REST API and OpenAPI** — Scoped service tokens; every UI read is an API call; /api/docs
- **Webhooks and e-mail** — HMAC-SHA256 on every warning, incident and directive event
- **TAXII 2.1 and STIX** — Indicators, policy decisions, exchange with other nodes
- **Imports** — JSON, NDJSON, CSV, STIX bundles via admin console or API
- **Entity portal** — Entity-scoped view behind the same clearance sign-in
- **Signed reports, evidence** — Three-level situation report; evidence bundle, offline-verifiable

(1) Read-only exports: assets, flows, DNS, AWS audit, scanner results →

(1) ← SIEM and EDR detections in as incidents (REST); alerts as CEF or JSON: pilot

(4) Warnings, directives with SLA, assignments: webhooks to ticketing, SOAR, SIEM →

(4) ← Entity portal: owners acknowledge, report compliance, see their posture

(2) Authorised testing of AI gateways and web apps (written authorisation on file) →

(1) TAXII 2.1: indicators, protective-DNS policy, exchange with tier-1 nodes ↔

(7) Signed situation reports and evidence bundles →

GOVERNMENT PRIVATE NETWORK — Connected entities

National SOC and CERT (existing, kept)
Alerts in as incidents; Spark incidents and warnings out by webhook; intelligence both ways over TAXII. Spark feeds it and is fed by it, never replaces it.

- **Ministry, government services** (5) — Public portals, citizen AI assistant, cloud estate — Keeps: SIEM · EDR · SOAR · firewall and WAF · identity · cloud
- **Energy and water utility** (5) — IT and OT inventories from exports; OT protocol monitoring: partner — Keeps: SIEM · EDR · SOAR · firewall and WAF · identity · cloud
- **Telecommunications operator** (5) — National resolver, DNS, large public surface — Keeps: SIEM · EDR · SOAR · firewall and WAF · identity · cloud
- **Financial institution** (5) — Regulated reporting, mail authentication, web and API estate — Keeps: SIEM · EDR · SOAR · firewall and WAF · identity · cloud
- **Healthcare provider** (5) — Patient portals, AI triage gateway, cloud accounts — Keeps: SIEM · EDR · SOAR · firewall and WAF · identity · cloud

Tier-1 entity with its own Spark node
Spark Enterprise Edition on the entity's hardware with its own database; exchanges with the national core by TAXII 2.1 and signed reports, never shared tables.

National resolver, protective DNS
Consumes SHIELD policy decisions from the TAXII collection and blocks at the resolver. Pilot.

What each entity provides
A named owner, an authorised-targets list with written authorisation, scheduled read-only exports, a webhook endpoint or mailbox for directives, a service account for its SIEM or SOAR.

LEGEND

- API call over HTTPS with mutual TLS
- File export or import (JSON, NDJSON, CSV, STIX)
- One-way transfer into the core (data diode or verified import)
- (4) Step of the end-to-end loop (right)
- Pilot: runs today, proven on representative data
- Roadmap: not built
- Trust boundary or kept system

THE END-TO-END LOOP: THE NUMBERS ON THE ARROWS

(1) Collect → (2) Detect → (3) Prioritise → (4) Assign → (5) Act → (6) Verify → (7) Report

1 exports, live checks, detections · 2 catalogue match, failed policy, injected AI · 3 severity, SLA, threat index, SRE-1 · 4 owner, directive, incident stage
5 the entity, with its own tools · 6 re-assessment and acknowledgement · 7 signed, three levels, verifiable offline with the public key

THREE STAGES, EACH WITH A GATE

- **Proof of value** — one node, SQLite, a handful of entities
- **Pilot** — one application node, PostgreSQL, up to 50 entities
- **National** — two nodes, primary and standby, GPU node, relay, hundreds of entities

Spark Software Middle East Cybersecurity Solutions · Reference architecture for the national programme · Sizes, limits and measured figures: section 3 of the team brief · Laboratory figures, no production deployment yet