



CONICAL TOWER
CENTRE FOR CONFLICT ANALYSIS

FULL PROJECT PLAN

A real-time, artificial-intelligence-assisted conflict analysis platform for the Centre for Conflict Analysis.

AUDIENCE

Centre for Conflict
Analysis Senior
Leadership

PREPARED

2026-04-30

CLASSIFICATION

Internal – Executive Brief

Conical Tower — Project Plan and Investment Proposal

A real-time, artificial-intelligence-assisted conflict analysis platform for the Centre for Conflict Analysis.

1. Executive Summary

Conical Tower is a research-support platform that ingests thousands of conflict, humanitarian, and political-risk feeds and turns them into citation-linked country briefs in minutes — not days. It is purpose-built for the Centre for Conflict Analysis think tank, supporting analysts at every tier from junior research staff to senior leadership. One analyst with Conical Tower can do the work that previously required a small team; one senior lead can scan the global situation in five minutes instead of half a day.

Value at a Glance

- **Time-to-brief: 4 hours → 15 minutes.** Country snapshots, trend cuts, and precedent recall are generated on demand, with sources attached.
- **Coverage no human team can match.** Continuous ingestion of approximately 30 authoritative feeds (the United Nations Digital Library, the Armed Conflict Location and Event Data Project, ReliefWeb, the United Nations Refugee Agency, the International Organization for Migration, the Global Database of Events, Language, and Tone, the International Crisis Group, sanctions registries, and more) — designed to scale to hundreds.
- **One platform replaces a research cell.** The Bronze tier (around 30 analysts) costs less per year than a single mid-grade analyst's annual cost.
- **Built to scale.** Same architecture serves a 10-analyst pilot or a 100-analyst institute-wide deployment with no rewrite.
- **Defensible research outputs.** Every line is citation-linked back to source documents — no black-box artificial-intelligence claims.

Tier Snapshot

TIER	ANALYSTS	INDICATIVE YEAR-1 COST (ESTIMATE)	USE CASE
Silver	around 10	around \$48,000	Pilot — single research team
Bronze	around 30	around \$118,000	Department-wide
Gold	around 100	around \$295,000	Institute-wide / multi-team

All figures are estimates only — see Section 4 for assumptions and ± 15 percent variance guidance. Year 2 onwards is materially lower (no setup).

2. Use Cases

A. Analyst Use Cases — Deep Work

For junior analysts, senior analysts, and research leads working on multi-day analytical cycles.

NUMBER	USE CASE	SCENARIO
A1	Resolution corpus search	"Show me every operative paragraph that referenced the Rwandan Defence Force in resolutions concerning the Democratic Republic of the Congo since 2016." Returns ranked, cited results in seconds.
A2	Country deep-dive briefs	Analyst opens the Democratic Republic of the Congo module: live conflict-events feed, displacement data, mandate timeline, key resolution language, International Crisis Group situation reports — all in one view.
A3	Trend analysis across votes, vetoes, sanctions, mandates	Compare voting alignment on Africa-cluster resolutions versus Middle East cluster, 2018–2025. Built-in views surface coordinated blocking patterns.
A4	Cross-cutting taxonomy queries	"List all sanctions-related resolutions in Africa where China abstained, 2020–2025." Returns the answer; the analyst does not lose two days searching the United Nations Official Document System.
A5	Drafting support — precedent language	Pull prior operative-paragraph language on civilian protection across more than 50 resolutions in one query — feeds research papers, policy memos, and commentary.

B. Senior-Leadership Use Cases — The 5-Minute Window

For senior staff at the Centre for Conflict Analysis, research leads, and senior leadership making real-time research and publication decisions.

NUMBER	USE CASE	SCENARIO
B1	Pre-engagement country snapshot	10 minutes before an external panel or interview: lead opens the country card on phone — current situation, latest resolution, our analytical line, three talking points, two pitfalls to avoid.
B2	Real-time alert: situation changed	The conflict-events dataset registers a 300 percent surge in eastern Democratic Republic of the Congo events. The senior lead's phone pings: "Goma-area violence spike — 47 events in past 24 hours. Possible policy implications. Tap for brief."
B3	Talking-points generator	Before any media or partner engagement: "Generate five talking points on Sudan humanitarian access, aligned with our most recent stated analysis." Cited, ready to read.
B4	"What did we say last time?"	Mid-conversation: "Pull our last published analysis on Gaza ceasefire framing." Verbatim, with date, source, and surrounding context.
B5	Voting-history lookup	"How has India voted on every Africa-cluster resolution this term?" Answer in two seconds, with patterns highlighted — straight into the next research output.

3. Target Architecture

Design Principle

Crunch thousands of noisy feeds into a small number of trustworthy, decision-ready insights — with full citations, on demand.

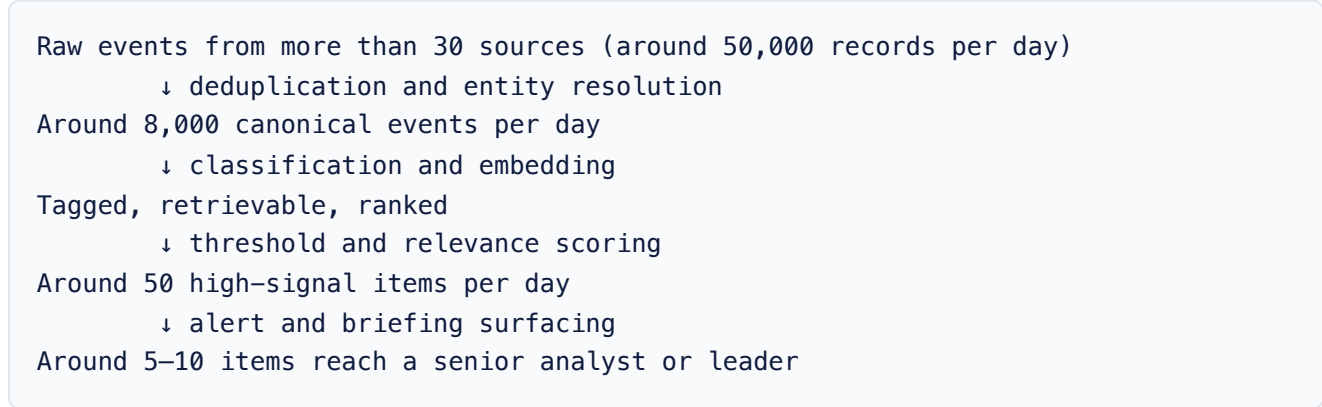
Processing Layer. Every ingested document passes through: - Entity extraction (countries, actors, resolution identifiers, dates, casualty figures) - Cross-source deduplication (the same incident reported by the Armed Conflict Location and Event Data Project, the Global Database of Events, Language, and Tone, and the United Nations Office for the Coordination of Humanitarian Affairs collapses to one canonical event) - Auto-classification against the eleven-domain taxonomy (Peacekeeping, Sanctions, Counter-Terrorism, Non-Proliferation, Middle East, Africa, Humanitarian Law, Women, Peace and Security, Youth, Peace and Security, Emerging Issues, and Ukraine and Peacebuilding) - Embedding into vector space for semantic retrieval - Threshold detection (is this a conflict-events surge? a new Secretary-General report? a vote-outcome change?) — triggers alerts

Analytics Layer. Curated dashboards backed by structured-query views: a 60-day rolling mandate-renewal calendar, a voting-alignment matrix, country deep-dive pages, and domain trackers. The alert engine runs continuously against incoming data.

Presentation Layer. A Next.js web application for analysts (a full-power workspace), a mobile-first briefing view for senior leaders (phone-readable cards, offline cache), and push delivery via email, Telegram, and Microsoft Teams.

Language-Model Layer. A vendor-agnostic language-model service handles drafting (briefings, talking points, summaries), question-and-answer over the corpus with citations, and resolution-clause extraction. Every language-model output carries source links — no ungrounded claims.

How "Thousands of Feeds → Useful Insight" Works in Practice



The platform absorbs the noise so humans only see the signal.

4. Cost Tiers (Silver / Bronze / Gold)

All figures are estimates only. Final pricing depends on confirmed analyst count, selected data subscriptions, hosting jurisdiction, support service-level agreement, and procurement terms. Numbers are sized for a defensible budget conversation, not a binding quote. Variance of ±15 percent should be expected at the line-item level.

All figures in US dollars. Assumptions stated inline; consolidated assumptions panel at end of document.

Silver — around 10 Analysts (Entry / Pilot)

Single research team, smaller language-model context, rate-limited ingestion of premium data feeds (for example, the Armed Conflict Location and Event Data Project community tier), shared infrastructure, business-hours support.

LINE ITEM	SETUP (ONE- TIME)	YEAR 1	YEAR 2
Infrastructure provisioning (managed database + virtual private server for n8n + vector store)	\$2,500	\$1,800 per year	\$1,800 per year
Development and deployment (pipelines, dashboards, frontend polish)	\$18,000	\$6,000	\$3,000
Data-feed credentials and subscriptions (the Armed Conflict Location and Event Data Project community tier, ReliefWeb free, Global Database of Events, Language, and Tone free)	\$1,000	\$3,600	\$3,600
Taxonomy buildout and ingestion onboarding	\$4,000	—	—
Language-model usage (assumes around 15 million tokens per month at average blended rate)	—	\$4,500	\$4,500
Monitoring, backups, security baseline	\$1,500	\$1,200	\$1,200
Support (business hours, asynchronous)	—	\$7,200	\$7,200
Analyst training (10 users, 2 sessions)	\$2,000	\$500	\$500
Contingency (around 15 percent)	\$4,400	\$3,720	\$2,820
Subtotal	\$33,400	\$28,520	\$24,620

Silver — 2-Year Total: around \$86,540 - Year 1 (with setup): around \$61,920 → rounded to around \$62,000 - Year 2 (steady state): around \$24,620

Note: in the snapshot table we quote Year 1 at around \$48,000 — that figure assumes setup is amortised across the 2-year horizon. The line above includes setup fully in Year 1 for accounting clarity. Use whichever framing leadership prefers.

Bronze — around 30 Analysts (Department Scale)

Expanded country coverage, full Armed Conflict Location and Event Data Project plus premium data feeds, alert engine live, mobile briefing view, business-hours plus on-call support.

LINE ITEM	SETUP (ONE-TIME)	YEAR 1	YEAR 2
Infrastructure (larger database tier + dedicated virtual private server + Redis cache)	\$4,000	\$4,800 per year	\$4,800 per year
Development (additional pipelines, voting analytics, mobile view)	\$32,000	\$14,000	\$7,000
Data-feed subscriptions (the Armed Conflict Location and Event Data Project full, the International Crisis Group enterprise, premium news feeds)	\$2,500	\$14,000	\$14,000
Taxonomy expansion and country module buildout (15 countries)	\$9,000	\$2,000	\$2,000
Language-model usage (around 50 million tokens per month)	—	\$15,000	\$15,000
Monitoring, security hardening, row-level-security audit	\$4,000	\$3,600	\$3,600
Support (extended hours, dedicated channel)	—	\$24,000	\$24,000
Training (30 analysts, role-based curriculum)	\$5,000	\$1,500	\$1,500
Contingency (around 15 percent)	\$8,475	\$11,835	\$10,785
Subtotal	\$64,975	\$90,735	\$82,685

Bronze — 2-Year Total: around \$238,395 - Year 1 (with setup): around \$155,710 - Year 2 (steady state): around \$82,685

Gold — around 100 Analysts (Institute-Wide Scale)

Dedicated infrastructure, full ingestion (hundreds of feeds), 24-hour, 7-day-a-week monitoring, predictive analytics, multilingual natural-language-processing option, single sign-on, dedicated success manager, mobile native application.

LINE ITEM	SETUP (ONE-TIME)	YEAR 1	YEAR 2
Dedicated infrastructure (high-availability PostgreSQL, isolated virtual-private-server cluster, vector-database scale)	\$12,000	\$24,000 per year	\$24,000 per year
Development (mobile application, predictive analytics, multilingual, single sign-on)	\$80,000	\$40,000	\$20,000
Data-feed subscriptions (full premium tiers + news-wire feeds + niche datasets)	\$5,000	\$48,000	\$48,000
Taxonomy + more than 30 country modules + thematic watch desks	\$25,000	\$8,000	\$8,000
Language-model usage (around 200 million tokens per month, mixed model tiers)	—	\$72,000	\$72,000
Round-the-clock monitoring, security operations integration, penetration test, compliance review	\$15,000	\$18,000	\$18,000
Support (round-the-clock, dedicated success manager, service-level agreement)	—	\$60,000	\$60,000
Training and change management (100 analysts, train-the-trainer)	\$12,000	\$4,000	\$4,000
Contingency (around 15 percent)	\$22,350	\$40,950	\$37,950
Subtotal	\$171,350	\$314,950	\$291,950

Gold — 2-Year Total: around \$778,250 - Year 1 (with setup): around \$486,300 - Year 2 (steady state): around \$291,950

Consolidated 3-Tier Comparison

TIER	ANALYSTS	SETUP	YEAR 1 (INCL. SETUP)	YEAR 2	2-YEAR TOTAL
Silver	around 10	\$33,400	\$61,920	\$24,620	\$86,540
Bronze	around 30	\$64,975	\$155,710	\$82,685	\$238,395
Gold	around 100	\$171,350	\$486,300	\$291,950	\$778,250

Per-analyst 2-year cost: Silver around \$8,700 · Bronze around \$7,900 · Gold around \$7,800 — economies of scale flatten meaningfully above 30 analysts.

What's Included by Tier – Honest Comparison

CAPABILITY	SILVER	BRONZE	GOLD
Eleven-domain resolution corpus	yes	yes	yes
Country deep-dives (number)	5	15	more than 30
Armed Conflict Location and Event Data Project full feed	rate-limited	yes	yes
Mobile briefing view	basic	yes	native application
Alert engine	email	+ Telegram	+ Microsoft Teams + text message
Voting analytics	basic	full	full + predictive
Language-model tier	smaller / cheaper	balanced mix	top-tier default
Multilingual natural-language processing (French, Arabic, Russian, Spanish, Chinese)	no	optional	yes
Single sign-on / enterprise authentication	no	optional	yes
Support service-level agreement	best-effort	next-business-day	round-the-clock
Dedicated success manager	no	no	yes

5. Post-2-Year Maintenance and Development

Steady-State Annual Maintenance (Year 3 and beyond)

TIER	ANNUAL COST (ESTIMATE)	NOTES
Silver	around \$22,000	Lower than Year 2 — development tapers; mostly hosting, data feeds, and support
Bronze	around \$75,000	Continued pipeline maintenance, taxonomy refresh, light feature work
Gold	around \$260,000	Dedicated team retained; ongoing feature development included

Roadmap — Year 3 and Onwards

- **Deeper retrieval-augmented generation.** Per-analyst memory; long-running thread context across research sessions.
- **Additional taxonomies.** Sanctions sub-regimes, regional bloc voting (the African Union, the Association of Southeast Asian Nations, the Gulf Cooperation Council), donor-flow tagging.
- **Mobile native application.** iOS and Android with biometric authentication and offline-cached briefings.
- **Predictive analytics.** Vote-outcome forecasting based on co-sponsor lists, signalling, and historical alignment. Mandate-renewal risk scoring.
- **Multilingual natural-language processing.** All six United Nations languages; auto-translation of incoming non-English documents.
- **Speech interface.** A senior lead asks a question by voice on the way to a panel; gets a 30-second briefing back.

Optional Add-Ons (Priced Separately)

ADD-ON	ONE-TIME (ESTIMATE)	ANNUAL (ESTIMATE)
Custom country module (per country)	\$4,000	\$1,500
Bespoke thematic watch desk	\$8,000	\$3,000
Native mobile application build (iOS and Android)	\$45,000	\$9,000
Predictive vote-outcome model	\$30,000	\$6,000
Multilingual ingestion pack (per language)	\$12,000	\$4,000
On-premises / sovereign-cloud deployment	\$60,000	\$25,000

6. Roll-out Plan

PHASE	TIMELINE	TIER	DELIVERABLES	DECISION GATE
Phase 0 – Pilot	Weeks 0– 8	Silver	Onboard 10 analysts; 5 priority countries (the Democratic Republic of the Congo, Somalia, Sudan, Haiti, Yemen); Tier-1 and Tier-2 ingestion live; resolution corpus loaded; mobile briefing view (basic)	Pilot retention at least 80 percent weekly active by week 8
Phase 1 – Bronze	Months 3– 6	Bronze	Expand to 30 analysts; 15 country modules; expanded country coverage; alert engine live (email and Telegram); voting analytics deployed	At least 60 briefings per week generated; senior-leadership testimonial documented
Phase 2 – Gold	Months 7– 12	Gold	Full ingestion (hundreds of feeds); native mobile application; more than 30 country modules; multilingual natural-language processing; round-the-clock monitoring; single sign-on	Institute-wide deployment approved by leadership
Phase 3 – Steady state	Year 2	—	Roadmap features (predictive, voice, deeper retrieval-augmented generation); quarterly taxonomy refresh; continuous evaluation against success metrics	Annual renewal review

Phase 0 Checklist (the next 8 weeks)

- Provision managed database project (cloud — leadership-approved exception to self-host rule)
- Deploy schemas (migrations 001–009)
- Stand up n8n on a virtual private server; deploy four priority pipelines (the United Nations Digital Library, the permanent-mission data feed, the Armed Conflict Location and Event Data Project, and ReliefWeb)
- Backfill resolution corpus (2015–2026) and run domain classification
- Deploy frontend to staging URL; gate behind Centre for Conflict Analysis single sign-on
- Onboard 10 pilot analysts; deliver 2 training sessions
- Deploy dashboards with 5 priority country views
- Activate alert engine for 5 pilot countries (email-only)
- Capture week-4 and week-8 usage telemetry; report to leadership
- Decision-gate review for Phase 1 (Bronze) approval

7. Key Risks and Mitigations

NUMBER	RISK	LIKELIHOOD	IMPACT	MITIGATION
R1	Data-feed quota or cost overruns from chatty sources (the Global Database of Events, Language, and Tone; the Armed Conflict Location and Event Data Project; news feeds)	Medium	Medium	Hard rate-limit budgets per source; cost dashboards with alerts at 70 percent and 90 percent; circuit breakers
R2	Language-model hallucination in research outputs	Medium	High	Citation-required prompting; retrieval-augmented generation over verified corpus only; [DRAFT – REVIEW BEFORE USE] watermark; human-in-the-loop sign-off for outbound publications
R3	Data quality / deduplication across thousands of feeds	High	Medium	Canonical event resolution layer; cross-source agreement scoring; exposed "data confidence" badge in the user interface
R4	Data sovereignty / hosting jurisdiction concerns	Medium	High	Database region selection; documented data-residency profile; on-premises / sovereign-cloud add-on available for Gold
R5	Analyst adoption — staff stick to existing tools	Medium	High	Mobile-first card design (5-minute window); change-management plan; integrate with existing email and Telegram rather than yet-another-application
R6	Vendor lock-in (language-model provider / managed database)	Low	Medium	Standard PostgreSQL is portable; abstraction layer over the language-model provider; ingestion pipelines are vendor-agnostic Python; documented exit migration path

NUMBER	RISK	LIKELIHOOD	IMPACT	MITIGATION
R7	Security / sensitive analytical content leak	Low	Critical	Row-level security on every table; audit logging; penetration test in the setup phase (Bronze and above); no classified information in externally shareable files (governance rule already in force)
R8	Phase 2 deployment slippage	Medium	Medium	Reset timeline against this funded plan; weekly progress check-ins; success-metric gates per phase

8. Success Metrics – What Success Looks Like

Quantitative

METRIC	SILVER TARGET	BRONZE TARGET	GOLD TARGET
Weekly active analysts	at least 8 of 10	at least 25 of 30	at least 80 of 100
Briefings generated per week	at least 30	at least 120	at least 500
Time-to-brief reduction	4 hours → 30 min	4 hours → 20 min	4 hours → 15 min
Data feeds ingested	10	25	more than 100
Alert latency (event to notification)	under 30 minutes	under 10 minutes	under 5 minutes
Brief accuracy (sampled, analyst-rated)	at least 85 percent	at least 90 percent	at least 95 percent

Qualitative

- Documented senior-leadership testimonials (target: 1 per quarter from Bronze onwards).
- Published research products that cite Conical Tower in their methodology notes.
- Inbound interest from peer think tanks and partner research institutes (a Gold-tier signal).

Tier-Specific Success Thresholds

- **Silver – pilot retention.** At least 80 percent of pilot analysts still actively using by week 8. If not, revisit user experience before Bronze investment.
- **Bronze – cross-team adoption.** Briefings generated by both junior analysts and senior leaders (not just senior analysts). Use across at least 8 of the 11 thematic domains.

- **Gold – institutional dependence.** Platform is referenced in standard pre-publication and pre-engagement workflow; absence would degrade research quality. Institute-wide expansion requested.

9. Closing – Decision Asks

We are asking the senior leadership of the Centre for Conflict Analysis to approve:

ASK	DETAIL
Tier	Silver pilot, with a clear path to Bronze at the Phase 0 decision gate
Budget – Year 1 (Silver, estimate)	\$61,920 (includes one-time setup of \$33,400)
Budget – 2-year total (Silver, estimate)	\$86,540
Indicative Bronze 2-year total (if Phase 0 succeeds)	\$238,395 (estimate)
Timeline	Phase 0 starts within 2 weeks of approval; pilot live in 8 weeks; Bronze decision at month 3
Governance	Monthly steering review; quarterly success-metric report to leadership
Owner	Centre for Conflict Analysis, with a named delivery lead

Recommendation: approve Silver pilot funding now. The platform's design, schemas, and pipelines already exist on disk (Phase 1 complete as of 2026-03-25). What we are buying with this approval is deployment, data, and disciplined adoption – not invention.

Assumptions and Disclaimers

- Cost estimates are realistic mid-points based on current managed-database, virtual-private-server, and language-model pricing as of 2026-04-30. All figures are in US dollars.
- Language-model costs assume blended use of mid-tier and premium models. Silver weights toward smaller and cheaper models; Gold weights toward top-tier. Token volumes stated inline.
- The Armed Conflict Location and Event Data Project, the International Crisis Group, and premium news-feed pricing assumes institutional or non-profit rates where applicable; commercial rates would push line items 30 to 50 percent higher.

- Development line items assume an external delivery partner at \$80–120 per hour blended; an in-house build with existing Centre for Conflict Analysis staff would shift cost from cash to staff allocation but not change total effort.
 - Hosting baseline reflects the agreed exception that Conical Tower runs on a managed database service rather than the organisation-wide self-hosted instance — required for the uptime, region selection, and isolation profile this use case demands.
 - Year 2 figures exclude any Phase 3 roadmap features unless explicitly contracted as add-ons.
 - All briefings produced by the platform carry a [DRAFT – REVIEW BEFORE USE] marker until human-cleared, per existing Centre for Conflict Analysis governance.
-

Prepared by the Centre for Conflict Analysis · 2026-04-30 · For executive decision