



THINK DELUS

INTELLIGENCE AND MONITORING SYSTEMS · NAIROBI

THE NARRATIVE INTELLIGENCE DESK

A proposal for a monitoring and evidence service supporting the follow-through of the Trust Summit and the Nairobi Declaration on Trust and Call to Action.

PREPARED FOR	Awel Uwhanganye, Chief Executive Officer, The Trust Summit
SUBMITTED TO	The Trust Summit Secretariat, Nairobi
PREPARED BY	Think Delus
DATE	23 October 2026
SUBJECT	Ten-week pilot of the Narrative Intelligence Desk
VALIDITY	This offer remains open for 30 days from the date above. All prices are subject to confirmation.

47

items of Summit coverage collected and classified, 18 to 23 October 2026

23%

of keyword matches were false positives or other Nairobi Declarations

10

weeks for the pilot, with the exit stated at week 10



ILLUSTRATIVE IMAGE · A SUMMIT IN SESSION

This proposal is submitted on the initiative of Think Delus. **Think Delus monitored the public record of the Summit; we did not attend it, and nothing in this proposal should be read as a claim of association with the Summit or its convenors.** The evidence in sections 2 and 3 comes from a demonstration we ran on published coverage of the Summit, and any figure in it can be traced to a source file we hold and will share.

CONTENTS

01	Cover and submission details	1
02	Understanding of the requirement	3
03	Proposed service	4
04	Method and quality assurance	6
05	Scope, exclusions and dependencies	7
06	Pilot plan and timeline	7
07	Commercial terms	8
08	Governance	10
09	About Think Delus	10
10	Acceptance	12

THE DEMONSTRATION, AT A GLANCE

47

items collected across news, broadcast, official and social sources

31

signal items that concerned the Summit itself

42%

of signal came from the convenors' own machinery: 13 of 31 items


3

items pressed the delivery question, among them The Eastleigh Voice and KICTANet

2 UNDERSTANDING OF THE REQUIREMENT

The Nairobi Declaration on Trust and Call to Action, adopted on 23 October 2026, did more than record agreement. Its delivery steps create standing instruments. Paragraph 22 provides for a Coalition on Global Trust. Paragraph 23 provides for an inclusive Nairobi Trust Council, hosted by Kenya and drawn from the public, private and philanthropic sectors, to follow up and review progress on the commitments the Summit made.

THE REPAIR JOB THE SUMMIT SET ITSELF




Three relationships run through the Declaration and the Summit’s sessions. Each one is discussed in public, by name, and each public account needs a maintained record.

Paragraph 24 provides for a Global Trust Facility, championed for establishment by the United Nations, held and administered by UNDP and resourced from voluntary funds from state and non-state actors over four years. The Facility is to finance trust infrastructure, including national accreditation bodies, regulatory trust indices, data protection, cyber security authorities and judicial cooperation mechanisms. Commitments of that kind will be discussed in public, by name, for years.

Paragraph 22 A Coalition on Global Trust A standing coalition to carry the Declaration’s commitments beyond the room in which they were made.	Paragraph 23 A Nairobi Trust Council Hosted by Kenya and drawn from the public, private and philanthropic sectors, to follow up and review progress on the commitments themselves.
Paragraph 24 A Global Trust Facility Championed for establishment by the United Nations, held and administered by UNDP, financing trust infrastructure over four years.	Paragraph 25 A Framework for Action To be developed by the Secretariat, supported by a robust mechanism for monitoring, tracking and reporting on implementation and progress.

Paragraph 25 places the operational duty with your office: the Secretariat is to develop a Framework for Action, supported, in the Declaration’s words, by “a robust mechanism for monitoring, tracking and reporting on implementation and progress on a biennial basis”. Paragraph 28 confirms that the Summit convenes biennially, so the first reporting cycle runs to 2028. The Summit’s concept note names the delivery instruments that Framework will need: a Commitment Tracker holding validated commitments, named owners, milestones, status and evidence of progress; a Results Framework; a

trust measurement instrument requiring technical validation, clear ownership and a defensible methodology; political-economy and risk analysis; and a RACI ownership map.



Each of those instruments assumes a current, evidenced picture of the public account: what is being said about the Summit and its commitments, by whom, in what terms, and whether that account is holding. That picture does not currently exist as a maintained record. The requirement, as we understand it, is for an instrument that produces it continuously and can be cited without embarrassment. That is what this proposal offers.

ND-03 · SERVICE

3 PROPOSED SERVICE

The Narrative Intelligence Desk is the record of what is being said about you, classified so you can act on it, with the sources to prove it.

The Desk monitors coverage of the Trust Summit and its follow-through, classifies each item by frame, and delivers a short cycle of products your team can use without further processing.

SOURCES, MONITORED DAILY

- Web news**
Kenya and East Africa titles, online editions
- Broadcast, published on the web**
Station pages, streams and clips
- Official channels**
Summit, government and partner properties
- Social sources**
Where platform access is licensed



THE NARRATIVE INTELLIGENCE DESK

Each item read, de-duplicated and classified against a written rubric: relevance, frame and tone, with its source attached.



PRODUCTS, ON A FIXED CYCLE

- Daily alert digest**
07:30 EAT, risks flagged with reasons
- Weekly brief**
Mondays 09:00 EAT, three to five pages
- Corpus file**
All items, sources and classifications
- Monthly frame review**
Classification tested with your team

Daily alert digest, 07:30 EAT

Anything published in the previous cycle that scores as a risk to the Summit's commitments or reputation, with the source and the reason it was flagged.

Weekly brief, Mondays 09:00 EAT

Three to five pages. Coverage classified by frame: who is making the accountability argument, who is making the cost argument, who is amplifying the Summit's own account, and who is scrutinising it.

The corpus file

It holds all items collected, with sources and classifications, so that any number in a brief can be traced to the item it came from.

Monthly frame review

A working session with your team in which our classification is tested against your own reading and adjusted where yours is better informed.

SPECIMEN · WEEKLY BRIEF, PAGE ONE

DESK FORMAT, SHOWN AT REDUCED SCALE

WEEKLY BRIEF: THE SUMMIT AND ITS FOLLOW-THROUGH

Coverage classified by frame, with sources. Any figure traces to the corpus file.

RISK · FLAGGED

ACCOUNTABILITY

COST

AMPLIFICATION

SCRUTINY

Item counts by frame · Source and classification logged per item · Prepared for the Secretariat's Monday reading

We have already run the method on your Summit

Between 18 and 23 October 2026 we collected and classified the public coverage of the Summit itself: 47 items across news, broadcast, official and social sources, each classified against a written rubric. Three results from that demonstration explain why the Desk is built the way it is.

11 NON-SIGNAL

5 ADJ


Signal, 31 items

Non-signal, 11 items

Adjacent, 5 items


Total collected: 47 items

Keyword noise



Of keyword-matched items were not about this Summit at all: 11 of the 47 items collected. Seven were false positives on ordinary news. Four were collisions with other documents called a Nairobi Declaration.

The convenors' own voice



Of signal coverage came from the convenors' own machinery: 13 of 31 signal items. Independent coverage accounts for the remaining 18.

First-party, partner or amplification channels

13 of 31

Independent coverage

18 of 31

First, 23 per cent of keyword-matched content was not about the Summit at all. Seven items were false positives on ordinary news, and four were collisions with other documents called a Nairobi Declaration. A keyword monitor would have reported that quarter as coverage of your event. Second, 42 per cent of the genuine signal came from the Summit's own properties, government channels, partner organisations and broadcasters' amplification pages, so volume flatters the event and

independent scrutiny is thinner than the count suggests. Third, only three items pressed the delivery question, among them The Eastleigh Voice and KICTANet. Those are the voices that will be quoted if delivery slips, and they are the ones the Desk watches most closely.

THE EASTLEIGH VOICE,
REPORTING THE SUMMIT

**“calls out gaps in
Kenya’s public
accountability”**

KICTANET

“trust is not a soft skill”

THE DELIVERY QUESTION

**3 items in 47 asked
what happens next. The
Desk watches those
voices closest.**

ND-04 · METHOD

4 METHOD AND QUALITY ASSURANCE

The Desk’s claim on your confidence rests on its method. Relevance, frame and tone are assigned against a written classification rubric, applied by analysts, and the rubric is available to your team on request. De-duplication and entity resolution happen before anything is counted, which is what removes the false positives and name collisions described above. The same item carried by six outlets is one item, and a different declaration with a similar name is a different subject.

STEP 01

Collect

Sources swept on a daily cycle, across the agreed source set.



STEP 02

Resolve

De-duplication and entity resolution run before any counting.



STEP 03

Classify

Relevance, frame and tone assigned against the written rubric.



STEP 04

Review and brief

Critical items checked by a second analyst, then issued with sources.

Any item classified as critical, or as a reputational risk, is reviewed by a second analyst before it reaches you. That two-person rule applies to the daily digest without exception, because an alert that is wrong costs more than an alert that is late. Once a quarter, 20 items are drawn at random and re-classified blind by an analyst who did not see the original classification. The agreement rate is published in the brief that follows the audit, and if it falls below 85 per cent the rubric is revised and the revision is recorded.

Quarterly blind audit: 20 items re-classified, agreement rate published in the following brief

FLOOR 85%



Below the 85 per cent floor, the rubric is revised and the revision is recorded in the corpus file.

THE CORRECTION PROMISE

If we publish a factual error, we correct it within four working hours of the error being identified, and the correction is recorded in the corpus file so the record shows what changed and when.

5 SCOPE, EXCLUSIONS AND DEPENDENCIES

✓ In scope for the pilot

- One subject: the Trust Summit, the Nairobi Declaration and their follow-through.
- Kenya and East Africa, across web news, broadcast content published on the web, official channels, and social sources where platform access is licensed.
- The four products listed in section 3, at the stated times.

✗ Out of scope

- Survey-based trust measurement and opinion polling.
- Legal or defamation assessment, and takedown requests of any kind.
- Crisis communications strategy. The Desk reports the record; the response belongs to your own team.

A BOUNDARY, STATED IN ADVANCE

We do not accept monitoring work on terms that require findings to be softened or withheld. A desk that sells the honest record cannot sell a favourable one.

The dependency that matters

Volume data from X and Meta is not included in this proposal. It depends on licensed platform API access, which the pilot price does not carry; where the Secretariat holds or obtains such access, or wishes us to, that coverage is available as a separately priced add-on described in section 7. Our demonstration was built without it, and the findings above stand on sources any reader can check.

6 PILOT PLAN AND TIMELINE

The pilot runs for ten weeks from a start date we agree with you. It is long enough to establish a baseline, run two monthly cycles of the frame review in shortened form, and give your team a record it can judge.

WEEKS 1 TO 2

Configure and baseline

WEEKS 3 TO 8


Full operation

WEEK 8

Formal review

WEEKS 9 TO 10

Operate and decide



PHASE	WHAT HAPPENS
Weeks 1 to 2	Configuration and baseline. Source set and search terms agreed with your team, the classification rubric calibrated against the Summit corpus, and a baseline brief issued covering the Summit period itself.
Weeks 3 to 8	Full operation. Daily alert digests, weekly briefs each Monday, and the corpus file maintained throughout.
Week 8	Review. A formal frame review with your team: the classification is tested against your reading, the source set is challenged, and we present the evidence for and against continuing.
Weeks 9 to 10	Operation continues. At the end of week 10 the Secretariat decides whether to end the pilot, extend it, or convert it to an annual contract at the rate in section 7.

The exit, stated plainly. The pilot ends at week 10 unless the Secretariat converts or extends it. There is no automatic rollover.

ND-07 · COMMERCIAL

7 COMMERCIAL TERMS

All amounts are in Kenya shillings and are **subject to confirmation**.

SETUP AND BASELINE, ONE-OFF

KES 220,000

Configuration, rubric calibration and the baseline brief covering the Summit period.

MONTHLY SUBSCRIPTION, DESK TIER

KES 295,000

Daily digests, weekly briefs, the corpus file and the monthly frame review.

TEN-WEEK PILOT, TOTAL

KES 900,000 approximately

220,000 setup plus ten weeks at the monthly rate. The pilot this proposal asks you to approve.

ANNUAL CONTRACT, IF CONVERTED

KES 3,540,000

Twelve months at the Desk tier rate, with a quarterly service review added.

Terms

- Invoiced monthly in advance. The first invoice is issued on signature and includes the setup amount.
- Minimum term three months. After the minimum term, either party may end the service on 30 days' notice.
- Platform API costs, where the Secretariat asks us to incur them, are passed through at cost plus 15 per cent and shown as a separate line on the invoice.

- Changes to scope are handled as change orders, priced and agreed before the work is done.

Add-ons, priced separately

ADD-ON	AMOUNT (KES)
Broadcast airtime monitoring and share of voice	240,000 per month, plus platform fees
Synthetic content and coordination check	180,000 per cycle
Crisis sprint, five working days	450,000
Additional subject	95,000 per month
Additional country	130,000 per month

Add-ons are never bundled into the pilot price. Each is taken only on the Secretariat's written instruction.

Market context


For orientation, published prices for media monitoring and social listening tools vary by scope. **Brand24** publishes plans from \$199 per month billed annually for its Individual plan to \$599 per month billed annually for its Business plan, with Enterprise from \$1,499 per month billed annually. **Mentionlytics** publishes plans at \$69, \$169, \$299 and \$499 per month in its plan comparison table. **BuzzSumo** publishes its PR and Comms plan at \$299 per month billed yearly. **Meltwater** publishes indicative starting points, from \$10,000 per year for Starter and from \$25,000 per year for Pro, and states that final pricing is confirmed through a scoping conversation. **Cision, Brandwatch and Talkwalker** publish no list price and route buyers to a pricing request, a demo or a custom quote.

These figures are for platform access on the terms stated by each provider. They do not include analyst time, briefing or reporting work, and the scope covered differs between providers.

Sources, accessed 23 October 2026: brand24.com/prices, mentionlytics.com/pricing, buzzsumo.com/pricing, meltwater.com/en/blog/meltwater-pricing, cision.com/pricing, brandwatch.com/pricing, talkwalker.com/pricing.

8 GOVERNANCE

Your named contact at Think Delus is Sammy Lusiola, Commercial Lead, who owns this engagement and attends the monthly frame review. Operational queries are acknowledged within one working day. A factual error in anything we have issued is corrected within four working hours of identification, as set out in section 4, and the correction is logged in the corpus file.



If an issue is not resolved at desk level, it escalates to the Commercial Lead, and from there to a director of Think Delus, with a written account of the issue at each step. During the pilot, governance is the monthly frame review in section 3 and the formal review at week 8. On conversion to an annual contract, a quarterly service review is added, at which the audit results described in section 4 are tabled.

9 ABOUT THINK DELUS


Think Delus builds intelligence and monitoring systems for public-sector and regulated markets in Africa. We are based in Nairobi and deliver across Kenya, Tanzania, Ethiopia, Uganda and the wider region. Our clients include water utilities and a national water agency in Kenya, and broadcasters and donor programmes across five markets. The Narrative Intelligence Desk applies the same discipline as the rest of that work: a written method, a traceable record, and findings we are prepared to defend line by line.

BASED IN NAIROBI **KENYA** **TANZANIA** **ETHIOPIA** **UGANDA** **FIVE MARKETS, BROADCASTERS AND DONOR PROGRAMMES**



10 ACCEPTANCE

To accept this proposal and start the pilot, sign below and return one copy to Think Delus. On signature we issue the first invoice, including the setup amount, and agree the pilot start date with your team.



For Think Delus

SIGNATURE

NAME AND TITLE

DATE

For The Trust Summit Secretariat

SIGNATURE

NAME AND TITLE

DATE

AGREED PILOT START DATE