



Business case

Assessing the corporate mail archive

Abstract

Europe's essential services still take their critical decisions by email, and that is also how they are compromised. Phishing, invoice fraud, quiet forwarding and supplier impersonation land on the same entities that NIS 2 now treats as essential or important. Live filtering only sees mail from the day it is switched on. The archive holds the years before: who spoke to whom, which file moved, where it was sent, and which rule copied a thread outside the perimeter. Leaving that corpus unread is an unquantified liability under a law that is already in force.

This case compares two ways to assess a bounded set of in-scope knowledge-worker mailboxes, then to re-scan the flow. Mistral Large 3 on La Plateforme can read unstructured French, including emoji, but it sends the corpus out as tokens, prices by volume, and cannot show the same span, rule and frequency six months later on another machine. Percipion, the on-premise symbolic engine from the French company Aeteos, maps the same text onto inspectable knowledge schemas and reasons over them on the customer's own Debian host, offline, with no fine-tune on the case file and no inference sub-processor.

On cost, custody and auditability the on-premise path is the one that can be put in front of an incident-response file, an ANSSI reviewer or a court. The recommendation is to approve the bounded assessment, build the corruption schema first, and keep any large language model, if at all, for work that is allowed to leave the premises.

Introduction

The organization has already written most of the risk down. It is sitting in the mail archive.

Europe's essential services, power, payments, networks, hospitals, plants, rails, are run by large organizations whose critical decisions still travel by email. That is also how they are attacked. Phishing, invoice fraud, stolen files, a supplier who is no longer who they claim to be, a privileged user quietly forwarding the wrong attachment. These are not hypotheticals. ENISA's 2026 landscape shows ransomware, credential theft and geopolitically motivated activity landing on the same entities that NIS 2 now calls essential or important.

In a large group, that exposure sits with a few thousand knowledge workers : the people with a mailbox and access to information that can move markets, shut a plant, or walk out the door. NIS 2 (Directive (EU) 2022/2555) has turned that fact into a personal problem for the board. In 2026 the directive is no longer a transposition debate. In most Member States it is national law : register, manage risk, report incidents on a short clock, supervise the supply chain, and prove the controls work. Fines reach 2 percent of global turnover for essential entities. Management can be held to account. After an incident, nobody will ask whether the gateway was licensed. They will ask whether the organization can show what already left the building.

Live filtering cannot answer that. It only sees mail from the day it was switched on. The archive holds the years before : who spoke to whom, which file moved, where it was forwarded, what was asked, what was disclosed, which inbox rule copied a thread to a private address. That corpus is both the attack path and the evidence locker. Leaving it unexamined is not a neutral IT choice. It is an unquantified liability under a law that is already in force, in a threat climate that is already geopolitical.

This case proposes a bounded, lawful assessment of the corporate mail archive for in-scope knowledge workers, then a daily and monthly scan of what those mailboxes send and receive. The outcome is not a reading exercise. It is a defensible picture of compromise, leakage, fraud, conflict and third-party abuse, handed to incident response and into the NIS 2 evidence file before a regulator or an attacker forces the same search at 2 a.m.

Two engines are compared on the same scope. One is a large language model, shortlisted as Mistral Large 3 on Mistral La Plateforme in the EU, with a Batch API and a cached system prompt. The other is Percipion™, the on-premise cognitive computing system from Aeteos™. Both can read unstructured French, including

emoji. They do not cost the same, they do not keep the text in the same place, and they do not give an auditor the same kind of answer. The figures that follow size the corpus, the professional services and the 36-month cost. The close says which engine the case recommends, and why.

What is an email?

An email, for this case, is any message item in the organization's mail service or archive, body, headers, attachments, links, embedded objects and system-generated metadata, that was sent, received, drafted, forwarded or filed in connection with work. The questions the archive can answer are concrete. Who talked to whom? What file moved? Where was it forwarded? What information was asked? What information was disclosed?

Assessing the mailbox is the right control, provided it is lawful, proportionate and governed. The threats in scope are compromise of accounts, leakage of confidential data, bribery and corruption, conflicts of interest, fraud, competition and sanctions breaches, supply-chain abuse, insider recruitment, and geopolitically motivated targeting. The mailbox is the right place to look because that is where those events are negotiated, executed or covered up, and because NIS 2 requires the organization to detect, report and evidence them, not merely to filter tomorrow's mail.

Different threats

An email archive is a business record. For a large organization the in-scope issues fall into three families : cyber compromise, integrity and fraud, and confidentiality and regulatory duty. Email is where those families leave a trail.

Compromise of the mailbox

Phishing and credential theft, a fake executive, a fake supplier, malware, ransomware. Almost every initial-access campaign against knowledge workers still starts, or is confirmed, in email. The archive shows the lure, the payload, the first reply, and whether the account was used afterwards.

Disclosure of confidential information

Forwarding to a personal Gmail or Outlook address. Forwarding attachments that contain intellectual property, pricing, merger material, patient data or customer data. Engaging with a future employer or a journalist. Confidential files are still moved as attachments and forwards. The archive sees years of it, including after someone has left.

Corruption, bribery and influence

Offers or requests for improper payments. Discussions of “success fees”, side letters, cash, gifts, travel, jobs for relatives. Use of agents, consultants or offshore intermediaries as a cut-out. Kickbacks hidden in inflated invoices and discussed by email. Bribery is rarely booked as “bribe”. It shows up as awkward commercial language, unexplained intermediaries, and private-side coordination.

Conflict of interest

An undisclosed stake in a supplier or bidder. Side work for a competitor. Awarding work to a related party. Mixing personal and corporate negotiations on the same thread. Conflicts are relationships. Email shows who was talking to whom before a tender, a hire, or a contract award.

Fraud and asset misappropriation

Procurement fraud and bid-rigging coordination. Duplicate or fictitious vendors. Expense and reimbursement schemes set up by mail. Diversion of payments after a change-of-bank-detail request. Internal collusion between accounts payable, treasury and a vendor. The fraud event is often a payment. The agreement to commit it is in the mail.

Competition law and market abuse

Price signaling, or allocation of customers or territories. Exchange of commercially sensitive information with peers. Insider dealing or tipping around listings, results, or a transaction. These cases are built on who said what to

whom, and when. The mailbox is the contemporaneous record regulators already expect the organization to be able to search.

Sanctions, export control and third-country risk

Deals or technical data involving restricted parties or destinations. Circumvention via intermediaries. “Don’t put this in the official file.” Screening tools catch names in master data. Email catches the workaround.

Supply-chain and third-party abuse

Compromised vendor mail used against the organization. Staff coaching a supplier to bypass controls. Concentration on a single high-risk ICT provider. NIS 2 makes supply-chain security a duty. The archive shows how people and vendors actually communicate.

Insider threat and corruption of the person

Recruitment or grooming of an employee by a competitor, a criminal group, or a state actor. A gradual loyalty shift before resignation. Sabotage instructions, or disgruntled data-harvesting. The compromise of a person often looks like ordinary correspondence until volume, timing and destination are read together.

Geopolitical and state-nexus activity

Spear-phishing of executives, engineers and public-affairs staff. Collection against energy, telecoms, defence-adjacent manufacturing, transport and government. Hacktivist or proxy pressure around political events. ENISA’s current picture is that essential and important entities are the main targets, and email remains the cheap path to a person with access.

Duty of care and organizational integrity

Harassment, discrimination or retaliation recorded in mail. Destruction, or “clean up the thread”, after a concern is raised. Failure of managers to escalate a known

incident. These items are not the primary NIS 2 aim, but they appear in the same corpus and create separate legal risk if ignored.

Organization size

The addressable customer is a large organization of the type that already engages consulting firms on cybersecurity and regulatory programs : typically 10,000 or more employees, in energy, transport, banking, telecoms, health, manufacturing or the public sector. NIS 2 brings these groups into scope as essential or important entities.

Not every employee has a corporate mailbox or handles confidential data, and a first assessment is usually limited to one tenant, one country or one essential-entity perimeter, rather than the entire worldwide group. The baseline is sized on 400 in-scope mailboxes. This is a planning figure, not a claim about the full group estate. It is set deliberately at the low end of a large-organization wave so the investment can be judged on method, unit cost and outcome, not on an inflated seat count.

Four hundred mailboxes correspond to a first, bounded population of knowledge workers, one tenant, one country, or one essential-entity perimeter, rather than every licensed mailbox in a 10,000-plus employee group. The business case is approved or rejected on that structure. If the sponsor later extends the same method to 2,000 or 5,000 mailboxes, the cost scales on the published unit rate. It does not require a new invention of scope.

Knowledge workers

Knowledge workers are employees whose main job is to produce, interpret, decide on or circulate information, not to operate machines, move goods, or deliver a physical service. Peter Drucker¹, a management consultant, introduced the term for people who apply specialized knowledge rather than manual labor. In practice that means professionals, managers, specialists and office staff whose output is analysis, decisions, documents, designs, advice or coordination.

1 https://en.wikipedia.org/wiki/Peter_Drucker

Usually in scope are executives, directors and managers; finance, legal, compliance and HR, procurement, engineers, scientists and R&D, architects, product and IT, sales and key-account staff, strategy and communications, analysts and consultants, project and program staff, traders and underwriters, clinicians with systems access and control-room engineers.

In short, a knowledge worker here is anyone with a corporate mailbox who routinely sees information that is confidential, regulated or commercially sensitive, and who could therefore leak it, be socially engineered, or be corrupted through email.

Types of emails

Some messages are short and some are long. To estimate volume, the study uses four types : the short email, the normal email, the proper explanation, and the long brief.

A short, or operational, email moves a known task one step : approve, attach, confirm, ping, acknowledge. Little or no new argument.

A normal email is a complete work message with a little context : one request, one answer, a short status. Still scannable on a phone. Roughly a short paragraph or two.

A proper explanation, a marketing note or a client update, is used when the sender must explain a decision, a change, or a client situation. There is a beginning, a point, and a next step. Still one screen or two, not a document.

A long brief, a legal note or a project dump, does the work of a memo : a legal position, a project dump, an all-staff narrative, a long forward with commentary. People often should have used an attachment or a document.

A short email averages 43 words². A normal email averages 69 words³.

2 Farshad Kooti, Luca Maria Aiello, Mihajlo Grbovic, Kristina Lerman and Amin Mantrach, "Evolution of Conversations in the Age of Email Overload", 2015, USC Information Sciences Institute and Yahoo Labs.

3 Weijiang Li, Yinmeng Lai, Sandeep Soni and Koustuv Saha, "A Comparison of Language in AI-Generated and Human-Written Emails", 2025, University of Illinois Urbana-Champaign and Emory University.

A proper explanation averages 200 words⁴. A long brief averages 622 words⁵. These counts are the visible body the receiver has to read. Footers, signatures and disclaimers are not counted.

What is a word

Here, a word is a sequence of six alphabetical characters⁶, the average used in the study, followed by a space, so seven characters in total. For English or French the average is nearer six characters; for German or Dutch it is nearer seven. The separating space counts, and punctuation has to be allowed for.

Applied to the four lengths, that convention gives 301 bytes for a short email (43×7), 483 bytes for a normal email (69×7), 1,400 bytes for a proper explanation (200×7), and 4,354 bytes for a long brief (622×7).

Standard distribution of daily emails

An average worker receives 117 emails per workday (Microsoft Work Trend Index, 2025). Newsletters, spam and alerts, in short, junk that will not be kept, fall out. Over the long term only 54 percent of those 117 emails are kept and archived (Mail Manager, “The email and document usage management benchmark report”, 2021), so about 63 emails per workday remains.

From the reply-length study by Kooti, Grbovic, Lerman and Mantrach, about 60 percent of emails are short, 22 percent normal, 12 percent a proper explanation and 6 percent long briefs. On the 63 archived emails a day, that is 38 short, 14 normal, 8 proper explanations and 4 long briefs. Those are the estimates for mail received.

4 Prospeo, 2026, on email length.

5 Arthur Samoylovich and Adam J. Frankel, “A Cost Analysis of All-Staff Emails in an Australian Hospital Network”, 2026, The University of Queensland.

6 Linjieqiong Huang, Erik D. Reichle and Xingshan Li, “Comparative Analyses of the Information Content of Letters, Characters, and Inter-Word Spaces Across Writing Systems”, 2022, CAS Key Laboratory of Behavioral Science, Institute of Psychology, Chinese Academy of Sciences.

Of the 117 emails received, only 24 percent stay in the inbox, most probably because they need an answer (Sanebox, 2025). The answer may itself be short, normal, a proper explanation or a long brief, and the reply is kept. That is about 28 emails produced per workday : 17 short, 6 normal, 3 proper explanations and 2 long briefs.

Attachments

About 20 percent of emails received or sent have at least one attachment (DigitalWarRoom, 2024). Attachments may be PDF, Word, Excel, PowerPoint, images, zip archives, CSV, video or HTML. This study counts only attachments a worker can read as text. Those represent between 55 and 75 percent of files attached in a corporate environment. The study keeps the average, 65 percent.

Most documents are containers. For a 100 kb file, after the text is extracted and headers, footers and signatures removed, roughly 70 to 95 kb remains for plain text, 30 to 70 kb for rich text, 6 to 20 kb for a Word document, and 3 to 12 kb for a PDF. As a working estimate, the assessable text per 100 kb sits between 8 and 15 percent, so 8 to 15 kb. The study keeps 12 percent. Most attached documents fall between 50 and 350 kb. The study keeps 200 kb, which on the 12 percent rule adds 24 kb of text per attachment.

Estimates for a knowledge worker

The tables below size what one knowledge worker receives, sends and archives. Sizes are estimated words in the assessable body, not the raw message on the wire. Headers, footers, signatures, disclaimers and format bytes are excluded, because the assessment runs on the extracted text.

Table 1. Daily average of emails received.

Email types	Daily amount	Size in words	Size per unit (bytes)
Short	38	43	301
Normal	14	69	483
Proper explanation	8	200	1 400

Long Brief	4	622	4 354
Attachments	8	3 511	24 576

Table 2. Daily average of emails sent.

Email types	Daily amount	Size in words	Size per unit (bytes)
Short	17	43	301
Normal	6	69	483
Proper explanation	3	200	1 400
Long Brief	2	622	4 354
Attachments	4	3 511	24 576

On those estimates, each workday a knowledge worker's mailbox grows by 362,651 bytes, about 354 kb, where 1 kb is 1,024 bytes. In France people work on average 218 days a year, so each year the mailbox grows by 77,205 kb of assessable characters in bodies sent or received. Each month that is about 6,596 kb.

Table 3. Monthly average of emails archived per worker (sent and received).

Email types	Monthly amount	Size in words	Size per unit (bytes)
Short	999	43	301
Normal	363	69	483
Proper explanation	200	200	1 400
Long Brief	109	622	4 354
Attachments	217	3 511	24 576

Estimates for the workforce

The baseline is 400 in-scope mailboxes. Each grows by 6,434 kb a month. That figure is the one used for monthly computing cost.

Employee-tenure surveys group workers by time with the current employer. OECD rough figures for a European firm are 12 to 16 percent with less than one year of service, 25 to 30 percent with one to four years, 20 to 25 percent with five to nine years, and 35 to 45 percent with more than ten years.

This case uses the mid-points of those bands : 14 percent under one year, 26 percent at one to four years, 20 percent at five to nine years, and 40 percent over ten years. On 400 mailboxes that is 56, 104, 80 and 160.

Archive depth is set to 10 months for under one year, 36 months for one to four years, 84 months for five to nine years, and 144 months for over ten years. Those depths, times 6,434 kb a month, produce the per-worker sizes in Table 4.

Table 4. Amount and average size of mail archives, as published.

Time within the company	FTE	Archive size per worker (Mb)
Less than 1 year	56	63
1 to 4 years	104	226
5 to 9 years	80	528
Over 10 years	160	905

Table 5. Emails to be assessed per archive, by tenure.

Type of emails	Under 1 year	1 to 4 years	5 to 9 years	Over 10 years
Short	9 992	35 970	83 930	143 880
Normal	3 633	13 080	30 520	52 320
Proper explanation	1 998	7 194	16 786	28 776
Long Brief	1 090	3 924	9 156	15 696
Attachments	2 173	7 822	18 251	31 287

The use case

The use case is a lawful assessment of the corporate mail archive for in-scope knowledge workers, followed by a daily or monthly scan of incoming and outgoing messages. The expected outcome is not a reading exercise. It is a defensible picture of compromise, leakage, fraud, conflict and third-party abuse, handed to incident response and into the NIS 2 evidence file.

Two technologies are compared. First, a large language model. Then Percipion™.

Which large language model is relevant?

French corporate email may contain personal data, and the processing has to stand up to audit or legal use. The gating criteria are EU processing, no training on the customer's data, structured JSON or XML output, the ability to process emoji, and a posture that can be defended under the GDPR and the AI Act.

To avoid exposure of business data to the US Cloud Act, Azure OpenAI, Anthropic Claude and similar US-hosted services are out of this shortlist. Aleph Alpha is a credible European alternative, but it is weaker than Mistral on French. The study therefore uses Mistral Large 3 on Mistral La Plateforme (EU), with the Batch API and a cached system prompt.

Adapting the detection system

Mistral and Percipion™ can both process unstructured natural language, here French, and both can transcribe emoji into words and infer a meaning. Neither yet knows the organization's own terminology : product names, internal codes, subsidiaries, VIP vendors, friendly nicknames, local acronyms, typical invoice patterns. Those terms have to be added first.

If the first assessment is looking for corruption, the detection system has to learn what this organization means by gifts, hospitality, facilitation, conflict of interest, kickback language, opacity, urgency and pressure, concealment, third-party intermediaries, unusual payment routing, and affective cues such as fear,

complicity and euphoria, together with the emoji lexicon. The entity glossary has to cover products, brands, legal entities, aliases and high-risk counterparts.

A gold set of relevant emails, labeled by compliance, is then required for calibration and for the test that the system is ready to deploy. That job has to be done whether the engine is Mistral or Percipion™. There is a starter project before any production run, then a recurring maintenance project. What is adapted is the detection system : ontology, entities, labels, prompts or schemas, and the evaluation.

Work packages, either engine

Governance and legal

Scope and decision rights : who can see a suspect email.

- Legal basis, Sapin II mapping, and a DPIA / AIPD.
- AI Act posture : human oversight, logging, prohibited-use check.
- A data-processing agreement with the model host, a no-training clause, and EU processing — required where an API is used.
- Retention, access, and an evidence chain of custody.

Domain adaptation

Corruption ontology, version 1 : categories, weak versus strong semantics, concealment, intermediaries, gifts and hospitality, conflicts of interest, payment routing, pressure and urgency.

- An affect layer : emotions and affective states that co-occur with corruption language.
- Emoji and informal lexicon, including French slang.
- An enterprise entity book : products, services, legal entities, cost centres, VIP vendors, aliases, internal codes.
- An annotation guide : what is a hit, what is a hard negative, what is too weak to flag.
- Annotator training for compliance reviewers.

Gold set

A sampling plan across the five email lengths, attachments included.

- Labeling of 1,500 to 2,500 emails, with dual review on disagreements.
- An inter-annotator agreement report.
- A held-out test set that is never used to tune the prompt or the schema.

System build

A cached system prompt and a strict JSON schema for the language-model path : frequencies per category, spans, emoji, affect. Percipion™ uses a knowledge schema rather than a prompt.

- An ingest pipeline : EML or PST to text, emoji preserved, then batch, then JSON validation.
- An evidence store : one JSON per file, source hash, prompt or schema version.
- An optional verbatim report with highlights.
- An evaluation harness : precision and recall by category, and calibration of semantic frequency.

Transfer and go-live

A calibration workshop with compliance : threshold and false-positive budget.

- A runbook, on-call, and a prompt-change or schema-change procedure.
- User acceptance on a sealed mailbox sample, then a production pilot of one month or one archive.
- Client time is extra and is not in the consultant bill. Expect 20 to 40 days of compliance and legal time on the starter.

A starter of this kind takes 60 to 120 man-days. The estimate keeps 60. Quarterly maintenance takes 5 to 15 man-days. The estimate keeps 5. Those figures are the language-model path, before the Percipion™ reduction below.

Where Percipion shortens the project

The important difference sits in domain adaptation.

Adding a word to a large language model has to be prepared. Sentence generation is the fast part. For one word, asking a model for 500 distinct sentences, half affirmative and half negated, takes about 10 to 30 minutes of API time. A person then needs 30 to 90 minutes to delete copies, fix wrong uses of the word, and check that the negated half really denies the claim. Current project estimates are based on this.

Adding a word to a knowledge schema used by Percipion™ takes about a minute. It is faster. Part of the professional services on the starter therefore drops if Percipion™ is used. On previous projects, that starter is estimated at 45 man-days rather than 60. The same effect reaches maintenance, whose purpose is to refresh the model or the schema with new words. Percipion™ maintenance is estimated at 4 man-days a quarter rather than 5.

The day rates used are 1,400 euros for a consultant on the Mistral path, 840 euros for an Aeteos™ consultant, and 600 euros as the best offshore rate. The 1,400 euro figure should be read as an assumed integrator rate for a Mistral implementation, not as a Mistral list price, unless a quote is attached.

Offshore delivery

Some work packages can be done offshore, which reduces delivery cost. The Maghreb is the right lever for francophone technical execution. Madagascar is the right lever for volume. Maghreb is nearshore : same time zone, professional French, a three-hour flight, and the most mature market for a large account.

Morocco nearshore grids for 2025–2026, and Tunisia and Morocco delivery centres, advertise about 40 percent under the French market. Algeria has a strong mathematics and AI pool, ENSIA and the engineering schools, not yet absorbed by service centres, but the nearshore fabric is less industrialised.

Madagascar, francophone and offshore, is also an option. French is an official language, the offset is one or two hours, and the cost of living is very low. Franco-Malagasy firms, Bocasay and peers, bill from about 180 euros a day, with a claimed 40 to 50 percent saving against France at equal supervision.

Not everything can move. Technical leadership, reviews, the risk of error, and staff turnover stay close to the account. Strong profiles are not cheap even in

those locations. Morocco already sells a trained language-model engineer at 500 to 700 euros a day, not 250. The gap shrinks on the strong profiles. The data itself is the other constraint. A bank or a health account will often not send the training set to Antananarivo, and sometimes not to Casablanca either, without clauses and EU hosting. Offshore is possible. The best estimated blended day rate used here is 600 euros.

Table 6. Professional services for one use case (corruption).

Project type	Mistral	Offshore	Aeteos
Starter	€84 000	€36 000	€37 800
Maintenance	€7 000	€3 000	€3 360

These estimates are for one use case : assessing messages for evidence of corruption. The threats in scope are wider, and other models or knowledge domains should be added. Some activities are done once. Others repeat. For each additional use case count 60 percent of the initial starter, and states that the maintenance part remains the same. That last sentence should be defined before the multi-use-case totals are relied on.

Table 7. Professional services for each additional use case.

Project type	Mistral	Offshore	Aeteos
Starter	€50 400	€21 600	€22 680
Maintenance	€7 000	€3 000	€3 360

Assessment costs

Two analyses run on the 400 in-scope mailboxes. The first assesses the archives. The second assesses monthly incoming and outgoing mail.

Percipion™, from Aeteos™, is billed all-included. The current cost is a flat 2,000 euros per month for one server. There is no per-seat fee and no per-document fee. Once the server is in place, the monthly scan does not add a usage charge.

Mistral is priced on the bytes in each prompt, which is to say on tokens.

Mistral Large 3 on la Plateforme is \$0.50 input, \$0.05 cached input and \$1.50 output per million tokens. The live EUR toggle shows €0.44 input and €1.30 output per million. Cached input is 10% of the input rate, so €0.044 per million. The Batch API is a 50% discount on those rates (Prices from October 3rd 2026).

Our central token assumptions is the same for every file:

- Document text at 4 bytes per token (French UTF-8 on Mistral's tokenizer). A money emoji is usually 1–2 tokens and does not move this estimate.
- Cached system prompt of 2,500 tokens (categories, weak/strong signals, affect, emoji rules, JSON schema). That is 1,889 million cached tokens.
- 200 uncached wrapper tokens per file.
- 1,200 output tokens per file for a full evidence JSON. That is 907 million output tokens.
- No file needs splitting. The largest attachment is about 6,100 tokens, against a 256k context. Batch files are capped at 512 MB and 1 million requests, so this is several batch jobs, not one, but that does not change the token bill.
- Output is about 70% of the bill. We selected the base audit JSON to minimize costs.

The archive prices, on the narrative headcount of 400, are :

- Under one year : 21 euros per archive, 56 mailboxes.
- One to four years : 77 euros per archive, 104 mailboxes.
- Five to nine years : 179 euros per archive, 80 mailboxes.
- Over ten years : 306 euros per archive, 160 mailboxes.

After the archives, a monthly assessment covers mail received and sent by those 400 people. On Percipion™ that scan has no additional charge. On Mistral this run adds 850 euros a month.


The first three months set the project up. Archive assessment is then spread over nine months. In parallel, a monthly assessment runs on the mail of the 400 people.

Results

The results show cost over 36 months. Reported figures are total cost of ownership : license or server, monthly bill, project phase and maintenance phase.

Table 8. Cost for one use case.


	Mistral	Offshore	Aeteos
TCO after 36 months	€261 514	€169 514	€146 760



On one use case, Mistral-path consultants cost more than either alternative. Offshore delivery is the cheapest solution for four months. After that, Mistral offshore costs are on average 40 percent higher than Percipion™.

Table 9. Cost for two use cases.


	Mistral	Offshore	Aeteos
TCO after 36 months	€489 428	€324 628	€206 400



On two use cases, Mistral-path consultants again cost more than either alternative. Offshore is the cheapest solution for three months. After that, Mistral offshore costs are on average 84 percent higher than Percipion™.

Table 10. Cost for three use cases.

	Mistral	Offshore	Aeteos
TCO after 36 months	€717 342	€479 742	€266 040



On three use cases, the same pattern holds for three months, after which Mistral costs are on average 107 percent higher than Percipion™.

Discussion

The baseline of 400 mailboxes is low for a large organization. A sponsor should plan for 2,000 to 5,000 mailboxes in a 10,000-plus employee group. The study prices one, two and three use cases. The threats in scope are wider : compromise of accounts, leakage of confidential data, bribery and corruption, conflicts of interest, fraud, competition and sanctions breaches, supply-chain abuse, insider recruitment, and geopolitically motivated targeting. NIS 2 requires the organization to detect, report and evidence them. Further use cases should be planned, not treated as optional.

The size estimates assume mail is cleaned before it is assessed. If it is not, volume rises and irrelevant text is scored. That penalty hits every token-priced engine, because the business pays for the amount of data processed. Trial runs before production were also not estimated. They raise the opening bill on a token engine. They do not raise Percipion's™ flat monthly rate. The more data Percipion™ assesses, the lower the unit cost to the business.

Percipion™, from the French family company Aeteos™, is an on-premise symbolic system. It maps text onto knowledge schemas and reasons over those schemas. Mistral Large 3 on La Plateforme is a statistical model : the email is sent to an API, and the model writes a JSON judgment. On the questions below they diverge as soon as the text is not allowed to leave the customer's premises.

Question 1 – Process : Is the heart of those modernization opportunities people reading unstructured text, emails, messages, reports, notes, seized files, tickets, open sources, and deciding what it means?

Yes for both. The work to be modernized is still a person deciding what unstructured French mail means. Percipion™ is built for that reading : messages, reports, seized files, weak signals, emoji, valence, and a published corruption schema. Mistral does the same job only if the taxonomy is placed in a prompt and the output is forced to JSON. The analyst stays in charge either way. The difference is that Percipion™ is a schema reader, and Mistral is a paraphraser that can invent a bridge the schema would not have drawn.

Question 2 – Sovereignty : Must the solution that reads this text be designed, owned and operated inside Europe, with no parent company, model weights, training corpus or sub-processor outside the legal perimeter?

Percipion™ matches the question as written. Aeteos™ is French and family-owned. The engine is claimed to be ANSI C99 with no third-party component, and it runs on the customer's Debian host. There is no foreign parent, no weight file, and no inference sub-processor.

Mistral AI is also French, and La Plateforme processes in the EU by default, which is why it won the earlier shortlist (LLM comparison). It still fails this stricter test. The weights and the training corpus sit outside the customer's perimeter, and Mistral is a sub-processor for every token.

Question 3 – Security : Could this workload ever leave the premises, even as tokens sent to an API? If the vendor's system can learn from the text, store it, or call home, is that a leak of the case?

If the rule is that the text never leaves the premises, not even as tokens, Mistral's API is out. A commercial Mistral call is not supposed to train on the mail, but the mail still goes to Mistral's servers.

Percipion's™ design is offline and on-premise, with no capture by Aeteos™ and no fine-tune on the case file, so the corpus cannot be poisoned by feeding it. That holds only if the host is actually air-gapped. A connector that phones home would break the claim. The default is that it does not.

Question 4 – Trustworthiness : Does the same text need to produce the same conclusion, for the same reasons, every time, including six months from now, on another machine, with another analyst watching?

Percipion™ is the engine that can aim at the same text, the same schema version and the same conclusion, six months later, on another machine. Symbolic execution is reviewable and, on Aeteos's™ account, deterministic enough to replay.

Mistral will not give that. Temperature zero reduces noise. It does not freeze the model. A vendor upgrade, a cache miss, or a prompt tweak can move a borderline bribery score. For an audit sample that must be re-run, that is a defect, not a taste preference.

Question 5 – Sustainability : If the hardware changes, the organization leaves the cloud, or the process is still needed in ten years, must the software run on ordinary machines, on little power, in a language that outlives frameworks and model vendors?

Percipion™ is the longevity bet : C, Linux, ordinary CPUs, little power, no model vendor in the loop. If Aeteos™ disappeared, the customer would still have a frozen binary and a schema version, though not necessarily the source.

Mistral as an API dies with the contract or the company. The open weights of Mistral Large 3 are an exit, but the exit is a GPU stack and a runtime that will be obsolete before the ten-year mark. For a process that must still run in 2036 on whatever hardware the lab has, Percipion™ fits the question and the API does not.

Question 6 – Legal : If a judge, an ANSSI auditor or an AI Act supervisor asks why this text was flagged, on which rule, and which human remained in charge, can the organization answer without opening a black box?

A judge, an ANSSI auditor or an AI Act supervisor can be shown a Percipion™ trace of this form : this span, this schema, this rule, this frequency, and the

analyst accepted or rejected it. That is an answer that does not require opening a weight matrix.

Mistral can be prompted to cite a span and a category. The citation is a generated story about a probability, not the rule that fired.

The host is part of the evidence.

A process that has to be shown to a regulator in 2030 should not be bound to a processor generation that Europe is already trying to leave!

The dominant enterprise instruction set is still x86, a complex instruction set owned and implemented by Intel and AMD. Around that instruction set sits a closed control plane: Intel Management Engine and AMD Platform Security Processor run below the operating system, with a privilege the tenant cannot fully inspect.

Closed GPU firmware and the CUDA driver stack add a second opaque layer.

None of this is proof that a given server is exfiltrating data. It is a larger attack surface than an essential entity can defend, and a dependency on suppliers outside the Union for the machine that holds the evidence.

Europe's answer is not a generic move from CISC to RISC. Arm is already a reduced instruction set, and the first European high-end server CPU now reaching machines, SiPearl's Rhea1, is an Arm design: samples were delivered to Bull in September 2026 for the CPU cluster of JUPITER.

The open path is RISC-V.

The instruction set is not owned by one company. The DARE programme, led by the Barcelona Supercomputing Center, is funded at close to €240 million for 2025–2028, with a little over €102 million from the EU, to produce European RISC-V chiplets for general compute, vector acceleration and AI inference.

Openchip's BER10, taped out on an open instruction set, and the Fraunhofer RISC-V secure element manufactured in Dresden are the same bet at a smaller scale: a core that can be read, rather than a core that must be trusted.

The European Processor Initiative has been explicit that the concern is no longer only supply. It is back doors and kill switches in technology sourced from outside Europe.

That is where the two engines in this case diverge again.

A large language model that is efficient today is, in practice, a CUDA workload. CUDA does not run on a RISC-V host. Important parts of the surrounding stack do not either: many Python binary wheels, GPU kernels and inference libraries are still built first for x86_64 and NVIDIA, with aarch64 second and RISC-V incomplete.

An organisation that makes a CUDA appliance the system of record for its mail archive is choosing a machine that European sovereign compute cannot yet replace, and a driver surface it cannot audit.

Percipion™ is written in ANSI C99, free of third-party components in the engine, and compatible with both RISC and CISC, on ordinary Debian. If that claim holds under test, the assessment does not need a GPU, a CUDA driver or a Python runtime, and the same binary and schema can move to a European Arm or RISC-V host when one is qualified.

Reduced instruction sets are not virtuous by themselves.

An open implementation can be inspected; a closed one cannot. The advantage for this file is a smaller trusted computing base, no opaque management engine as a condition of running the assessment, lower energy on a CPU-only host, and no licence from the instruction-set owner standing between the organization and the next generation of European silicon.

Selecting a CISC and CUDA host now, because that is where today's models run, fixes the evidence process to the dependency Europe is funding its way out of. The bounded assessment should be specified so that the host can follow.

Conclusion

The archive is already the evidence locker and already the attack path. NIS 2 does not ask whether a gateway was licensed. It asks whether the organization can show what left the building, on which signal, and under whose authority. A bounded assessment of 400 in-scope knowledge-worker mailboxes is the right first decision. It is small enough to approve, and it is the method that later scales to 2,000 or 5,000 mailboxes without being reinvented.

A large language model can read the French. Mistral Large 3 on La Plateforme is the right name on a shortlist that excludes US cloud APIs, and the professional-services path around it is real : ontology, gold set, DPIA, human oversight. It is also the more expensive path, the path that sends the corpus out as tokens, the

path whose answer can move when the model moves, and the path that cannot show a rule, only a generated citation. Offshore delivery lowers the opening invoice. It does not change the engine. By month 36, on the published figures, even the offshore Mistral path sits above Percipion™, and the gap widens as use cases are added.

Percipion™ is the cognitive computing system this case is built to recommend. It is not a chatbot wrapped around the mailbox. It maps the text onto knowledge schemas and reasons over those schemas, on the customer's own Debian host, offline, with no capture by Aeteos™ and no fine-tune on the case file. The analyst remains in charge. What the analyst is shown is a span, a schema, a rule and a frequency, the same trace six months later, on another machine, in front of another analyst. That is the shape of an answer a judge, an ANSSI auditor or an AI Act supervisor can be given without opening a weight matrix.

The commercial shape matches the duty. One server, 2,000 euros a month, no seat fee and no document fee. Adding a term to the schema is a minute of specialist work, not an afternoon of synthetic sentences. The starter is shorter, the quarterly maintenance is shorter, and the monthly scan of 400 mailboxes does not open a new invoice. Trial runs and poorly cleaned mail, which punish a token meter, do not punish the flat rate. The more of the archive the organization has the courage to assess, the stronger the unit-cost case becomes.

Sovereignty is not a slogan attached at the end. Aeteos™ is a French family company. The engine is claimed in ANSI C99, with no third-party component, on ordinary CPUs and little power. There is no foreign parent, no weight file, and no inference sub-processor. If the vendor were no longer there in 2036, a frozen binary and a schema version would still run. A process that has to outlive a model release cycle needs exactly that.

The recommendation is therefore plain. Approve the bounded assessment. Do the governance once, properly : legal basis, works-council consultation where French labor law requires it, Sapin II mapping, DPIA, human oversight, chain of custody. Build the corruption schema first, then the next domains, leakage, fraud, conflict, supply-chain abuse, insider recruitment, geopolitical targeting, on the same host. Run Percipion™ on the archive, then on the monthly flow. Keep Mistral, if at all, for work that is allowed to leave the premises. Do not make it the system of record for mail that may one day have to be explained in court.

Percipion™ is how the organization reads what it has already written, keeps the reading inside the perimeter, and can still show its working when someone asks.