

● MAI-MF-LLM-26B // READY
COMMAND NET • NO INTERNET
RECORDS • ROLES • AUDIT • VOICE

MILITARY AI • HEADQUARTERS • FMS

M.AI MAINFRAME

**Your intelligence. On
your infrastructure.**

The command-scale M.AI: one fine-tuned, on-premise assistant on your own air-gapped servers, for every desk from the motor pool to the general's office.

ON-PREMISE
ON YOUR SERVER

AIR-GAPPED
BY DESIGN

FINE-TUNED
FOR
SUSTAINMENT

HUMAN
IN COMMAND

People. Data. Readiness.

NexCommand AI LLC develops artificial intelligence and enterprise software for defense, aerospace, manufacturing, maintenance and supply chain operations. We connect information, people and workflows so organizations operate with greater clarity, accountability and control, with SaaS and on-premise editions to suit each operating requirement.

THE PORTFOLIO

M.AI

Military Artificial Intelligence

Readiness, sustainment and FMS.
The Field Kit for the operator; the
Mainframe for the command.

NORA

NexCommand Organizational Reasoning Assistant

The civilian sibling for enterprises,
on the same Mainframe platform.

NEO AI

NexCommand Enterprise Orchestration Artificial Intelligence

Sales, purchasing, accounting,
inventory, manufacturing, quality
and HR in one environment.

MISSION

To help organizations turn authorized data into actionable intelligence and coordinated operations — improving readiness, strengthening accountability and supporting informed human decisions.

VISION

To become a trusted provider of operational AI and connected enterprise systems for organizations where reliability, information control and readiness matter.

VALUES

Human authority

People decide; M.AI recommends.

Integrity

Claims supported by evidence.

Operational relevance

Built around real workflows.

Security and control

Your data stays under your control.

Quality and traceability

Every answer traced to a record.

Continuous improvement

Learn from users, validate, improve.

ISO 9001 certified · Quality management

ISO 27001 certified · Information security

Founded by a retired U.S. Special Forces veteran with 30 years in maintenance, logistics and supply chain, and a licensed Professional Engineer and AS9100 lead auditor. A U.S. company in Cheyenne, Wyoming.

THE OPERATIONAL PROBLEM

Readiness is decided every morning, from data nobody can see whole.

The records that keep a fleet mission capable already exist. They sit in different systems, in a coded vocabulary, and reach the commander late.

FRICITION POINT	WHAT HAPPENS TODAY	EFFECT ON READINESS
Fragmented data	Asset status in maintenance records; requisitions, due-ins and case data elsewhere; item data in FED LOG.	No single view of why an asset is NMCS, or when it returns.
Aged requisitions	Backorders past RDD are found late, in periodic reviews.	Longer customer wait time; assets stay AWP.
Missed discrepancy windows	Short shipments and transit damage are recorded, but SF 364 / SF 361 go in late.	Lost credit; repeat requisitions.
Hidden recovery options	Serviceable stock at another depot is not visible to the decision maker.	Lateral transfers and controlled exchanges missed.
Weak component history	Serial history split between supply, transport and maintenance.	Poor rotatable management and reliability analysis.
Vocabulary barrier	MILSTRIP, DLMS, TAMMS and FMS terms are coded and in English.	Long training; knowledge leaves with rotating staff.
Headquarters sees it last	The picture reaches the general after a morning of staff reconciling slides.	Decisions wait for the data, not the other way round.

WHY A GENERAL-PURPOSE AI IS NOT THE ANSWER

Cloud dependency

Generic assistants process data in someone else's data centre.

Invented facts

A plausible NSN, status or date that does not exist.

No permissions

Cannot enforce who may see which case, unit or asset.

No doctrine

Does not know BB means backordered, or that MRO means two things.

Wrong provenance

FY2026 NDAA §1532 bars AI domiciled in covered nations from defense work.

M.AI MAINFRAME AT A GLANCE

The command's own assistant, built for sustainment.

M.AI is the conversational layer of the NexCommand AI platform. On the Mainframe she serves a whole headquarters from its own servers: she reads live supply, maintenance and FMS case records with each user's permissions, explains them in doctrinal language, and proposes the next action for a person to approve.

01 · READINESS IS SUPPLY

Every NMCS asset tied to its NSN

M.AI joins asset status to the requisition, status code, case line and due-in that is holding it down.

02 · ANSWERS, NOT SCREENS

A two-minute spoken brief

Replaces a morning of dashboard review, every figure traceable to its record.

03 · AIR-GAPPED

Your servers, your network

No cloud, no telemetry, no outbound connection. Updates cross on signed media.

04 · DOCTRINE-LITERATE

Speaks the language of sustainment

MILSTRIP, DLMS, FED LOG, SDR/TDR, TAMMS and FMS case language, in context.

05 · ERP-NEUTRAL

One assistant over every system

SAP S/4HANA, Oracle EBS, Odoo, GCSS-Army, NAVSUP and custom systems.

06 · PEOPLE DECIDE

Recommend, never command

M.AI prepares options and drafts. Authorised people approve every transaction.

ONE PLATFORM, TWO FORMS

	M.AI FIELD KIT	M.AI MAINFRAME
Who uses it	One operator, carried to the work	Many named users at a link on your network
Where it runs	Rugged single board computer (SBC) with an AI accelerator	Your own servers with GPU workers
The model	mai-fk-llm-4b*	mai-mf-llm-26b*
Its face	A 2D holographic persona	A live 3D holographic avatar
Licence	Per device	Per named user (seats)

SEVEN DESIGN PRINCIPLES

Offline by design	Figures from records, not the model	The user's permissions, always
Recommend, never command	Doctrine-literate	Auditable end to end
		Say what is true

WHERE THE TIME GOES

M.AI takes the hours out of observe and orient. The decision stays with you.

Most of a headquarters' morning is spent assembling the picture, not deciding. M.AI assembles and checks it in seconds, so the time goes back to the commander.

THE DECISION CYCLE, BEFORE AND AFTER (ILLUSTRATIVE)

Without M.AI

Staff assemble the picture by hand

Observe: gather reports

Orient: reconcile sources

Decide

Act

With M.AI

Picture assembled and verified in seconds

Obs.

Orient

Decide

Act

time returned to the commander

M.AI compresses the observe and orient phases. Actual gains are measured in each pilot against the customer's own baseline.

WHERE M.AI TAKES THE TIME OUT

TASK	TODAY	WITH M.AI	MEASURE
Daily readiness picture	Reports pulled from several systems, reconciled into slides	Spoken brief at login; every figure traced	Hours → minutes
Component history	Supply, transport and maintenance records searched by hand	One serial-number timeline on request	Days → minutes
Aged requisitions	Found in a periodic review	Listed by PD and days past RDD, follow-ups drafted	Weekly → every morning
Discrepancy reports	Noticed near or after the deadline	Deadlines tracked; SF 364 drafts ready for review	Missed → on time
Regulation look-up	Publications searched page by page	The paragraph that applies, cited	Search → cite
New staff officer	Months to learn the codes	Asks in plain words from day one	Months → days

PILOT TARGETS WE SIGN UP TO

≥ 98%

of figures in answers verified against source records

≥ 90%

correct on the national evaluation set at pilot; ≥ 95% at rollout

0

permission or data-leakage incidents

100%

of PD 01–03 requisitions past RDD reviewed every week

Measures of effectiveness from the M.AI white paper. Before/after rows describe the work, not guaranteed results; baselines and targets are agreed per site.

OPERATIONS OVERVIEW

Your whole operation, on one screen.

Unlock the Mainframe and M.AI is already waiting: the live 3D assistant on the left, your status and next steps on the right, and one question box for anything else. Nothing on this screen leaves your building.



1 Local inference

The M.AI LLM runs on your server. Ready means the model is loaded and answering.

3 Internet

Checked, never needed. AI processing stays local.

5 System health

GPU, memory and storage on your own hardware.

7 Recent activity

Every unlock, question and answer, from the audit log.

2 ERP connection

Your system of record, read under your permissions.

4 Operational brief

What to do next, from your own workspace.

6 Connected systems

ERP, knowledge vault and Customer Console at a glance.

8 M.AI in 3D


Talk, mute, brief me and lock, beside the live assistant.

M.AI Mainframe internal preview v0.1.4 with synthetic demonstration data. The ERP connector is configured at installation.


MEET M.AI

A live 3D assistant your people will talk to.


On the Mainframe, M.AI is a true 3D holographic character in your service's uniform. She listens, looks at the person speaking and answers in their language, rendered live in the browser from your own server.



M.AI · United States



M.AI · Saudi Arabia



M.AI · Tanzania

Looks at you

Head and eyes follow the person speaking, the pointer or the page. Natural blinks and glances.

Camera stays off

Camera and microphone stay off until a user turns them on, and stop the moment the screen locks.

Speaks your language

Local English and Arabic voices with lip-sync; speed, pitch and depth are adjustable.

Fits the room

Calibrate for a laptop or a large operations-room screen, with eye-contact presets.

Stands to attention

When M.AI answers she faces you, upright and steady. Professional, happy, concerned or listening.

All on your server

Rendered in the browser from your own server. No video or audio leaves the building.

Avatar settings

Character, lighting, hologram intensity, expressions, attention mode, camera and microphone selection, room calibration and voice are all set from Settings, and remembered on that screen.

Rendered live by M.AI Mainframe internal preview v0.1.4 (avatar module 0.4.5). Further national characters follow the 14 editions.

THE MORNING BRIEF

Two minutes. Spoken. Every figure from a record.

The brief is M.AI's signature. It is prepared before the duty day, refreshed at login if records changed, and follows the shape a staff officer would use.

EXAMPLE BRIEF · ENGLISH · DEMO DATA

"Good morning, General."

"Fleet FMC rate is eighty-one percent. Four vehicles are NMCS and three are NMCM."

"The four NMCS vehicles are held by two NSNs. The NSN ending one-one-eight-seven is backordered, status code B-B, on case line three, and is twelve days past its required delivery date."

"Two serviceable units, condition code A, are on hand at the Northern depot. A lateral transfer would return two vehicles to FMC today. A controlled exchange from a long-term NMCM vehicle could return a third, subject to your approval."

"Two supply discrepancy reports for short shipments are approaching their submission limit. SF 364 drafts are ready for review."

"Would you like to approve the lateral transfer, see the NMCS list, or review the SDR drafts?"

THE SHAPE OF EVERY BRIEF

- 1 **Situation**
FMC / NMCS / NMCM, by fleet
- 2 **Drivers**
The NSNs and faults costing availability
- 3 **Supply actions**
Aged requisitions, due-ins, PD
- 4 **Discrepancies**
SDR / TDR deadlines
- 5 **Carried over**
Yesterday's actions and decisions
- 6 **Recommendation**
One, with its basis
- 7 **Decision options**
For the commander to choose

Net-style pronunciation

Codes spoken the way logisticians say them: "B-B", "zero-two", "DOD-ack". Captions show the standard codes.

Role-based

A general hears drivers and risk; a maintenance officer hears tails, work orders and parts.

Knows what was decided

Yesterday's recorded decisions carry into today's brief, so nothing is briefed twice.

Example adapted from the M.AI white paper; NSNs, case lines and figures are fictitious.

ASK · VERIFY · DECIDE

Every figure traced to a record. Every decision yours.

Ask in plain words. M.AI finds the records first, writes the answer second, and checks every figure against the source before you see it. When there is a choice to make, she lays out the options and you decide.

Show all NMCS vehicles with an open high-priority requisition 1 of 50 questions

Side pane + New chat

Hello, Colonel. I'm MAI. Ask me for a brief, an NSN, or anything about your data.

Show all NMCS vehicles with an open high-priori

BLUF: FMC rate 86%. NMCS 3, NMCM 2 of 35 assets.

asset	status	nsn	description	dataset	row
21-044	NMCS	5330-99-000-1187	Final drive leak	equipment_status	2
21-051	NMCS	5330-99-000-1187	Final drive leak	equipment_status	3
19-013	NMCS	4320-99-000-4521	Hydraulic pump	equipment_status	4
19-020	NMCM		Scheduled service	equipment_status	5
17-002	NMCM		Long-term repair	equipment_status	6

VERIFIED

Confidence: high Regulation: DA PAM 750-8 3-2 · AR 700-138 (not loaded) · DoDI 3110.05 (not loaded)

Why this answer

Save as

Word

Excel

PDF

Text

Bottom line first

The answer, then the records that support it, row by row.

Verified

Every NSN, status and figure matched to a record; the regulation cited.

No data, no guess

If the records are not loaded, M.AI says so.

Word PDF Excel

Refresh

2 asset(s) NMCS on NSN 5330-99-000-1187 (Final drive leak)

Con

- Assets: 21-044, 21-051
- Serviceable stock recorded: 4
- Open requisitions for this NSN: 1

A Lateral transfer from serviceable stock

Risk LOW

Recommended

Returns 2 of 2 asset(s) to FMC once the part arrives and is installed.

Pros

- 4 serviceable on hand at Northern depot
- No new procurement or cost

Cons

- Reduces the holding activity's stock
- Transport and a turn-in/issue document (DD

B Controlled exchange from NMCM asset 19-020

Risk MEDIUM

Can return 1 asset to FMC without waiting for supply, if the donor has a serviceable part installed.

Pros

- Does not depend on the supply system
- Donor is already non-mission-capable

Cons

- Degrades the donor further and adds maintenance man-hours
- Donor must be inspected first

C Follow up the open requisition and review its priority

Risk LOW

Keeps the requisition on track to replace the stock used; no immediate change to readiness.

Pros

- Document TST001-6150-0001 is status BB at PD 02, 16 days past RDD
- Keeps the supply record accurate

Cons

- Slow: depends on the source of supply
- A PD change must be justified by the unit's FAD and urgency of

D Do nothing (wait for no supply)

Risk HIGH

2 asset(s) stay NMCS until the delivered; requisition already past RDD.

Pros

- No extra workload or transport
- No stock moved from other activities

Cons

- 2 asset unavail
- Readin grows 1 day pat

Decision cards: each option with pros, cons, risk and approving authority, built only from your records. Record the decision; it goes to the audit log. Synthetic demonstration data.

DECISION SUPPORT · PEACETIME AND WARTIME

Courses of action, with the evidence behind each.

M.AI is a decision-support system for readiness and sustainment. From live records, the command's own history and its doctrine, she builds options, projects their effect and risk, and says which she would take and why. The commander decides.

DECISION AREA	PEACETIME	CRISIS AND WARTIME
Readiness generation	Balance phase inspections, training and deployments	Maximise available assets for a surge and sustain them
Maintenance priorities	Order work by readiness gain and backlog age	Prioritise repairs and BDAR by time to return to mission
Parts and stock	Set PLL / ASL and safety stock from demand history	Allocate scarce LRUs; lateral transfer; controlled exchange
Requisitions and FMS	Chase aged requisitions; plan case milestones	Recommend PD escalation where justified
Distribution	Optimise routing and consolidation	Re-plan when routes or nodes are disrupted
Force-level risk	Forecast availability for planned activity	Readiness risk by course of action

THREE COURSES OF ACTION, ONE RECOMMENDATION (EXERCISE DATA)

OPTION 1 Wait for the backordered NSN + No stock moved – 3 aircraft NMCS ≈ 12 more days EFFECT +0 FMC this week RISK High	OPTION 2 RECOMMENDED Lateral transfer from Northern depot + Condition code A stock + Within commander authority – Depot drops below safety stock EFFECT +1 FMC today RISK Low	OPTION 3 Controlled exchange from long-term NMCM + +1 FMC in 2 days – Extra man-hours; needs approval EFFECT +1 FMC in 2 days RISK Medium
--	---	--

M.AI recommends: Option 2 now; prepare Option 3 as a fallback. Confidence high (4 records, 0 gaps). Basis: stock record 05:30, requisition status 06:00. **Your decision.**

Commanders decide

M.AI never issues orders or executes on its own.

Not a weapon

Not a weapons, fire-control or targeting system.

Transparent

Assumptions, sources and confidence shown.


Reviewable

Every recommendation kept in the audit log.

KNOWLEDGE · COMPLIANCE · AUDIT


The unit's knowledge, mapped and accountable.

Point the Mainframe at your document stores and regulation libraries. M.AI reads them in place, maps each file to the rule that governs it, and keeps an append-only record of everything she does.



Knowledge locations

Shared folders, file servers and libraries, linked and watched. New files are read as they arrive; nothing is copied off your network.



Compliance map

Every file matched to the publication and paragraph that governs it, with coverage by publication. Reviewers confirm or reject; M.AI learns from both.



Append-only

Unlocks, questions, answers, imports and decisions, hash-chained.

Ask the log

Search it in plain words: which answers used this record?


Export the evidence

Excel, CSV, PDF, W3C PROV and STIX 2.1 for your security team.

RETRIEVAL-AUGMENTED GENERATION (RAG)


Retrieval first. The model second. Verification last.

M.AI does not answer from what a language model happens to remember. On the Mainframe she searches the command's records and libraries with each user's permissions, hands only those to the M.AI LLM, and checks every figure in the draft before anyone sees or hears it.



SERVER INDEX

Linked and watched, never copied off your network: full-text and meaning-based ranking over every page, row and paragraph, each labelled with who may see it. One index serves every desk; each user sees only what their role allows.



Why it matters


At command scale one invented NSN would travel into hundreds of decisions. Retrieval with a verification gate keeps the knowledge in your records, not in the model's weights, so a new publication or a new fleet needs loading, not retraining, and the 26B model stays honest.

Same pipeline in the Field Kit (v1.0.0.108) and the Mainframe's M.AI core. Mainframe ERP connectors: ♦ in development.

COMPLIANCE MAPPING

Every file mapped to the rule that governs it.

M.AI reads each document as it is linked and suggests the regulation paragraphs that govern it. Reviewers confirm or reject each suggestion, and M.AI learns from those decisions. The map is a live graph you can pan, zoom and follow.



Compliance map in the M.AI core (shown in Field Kit v1.0.0.108; the Mainframe carries the same function). Blue: records · green: forms · violet: regulation sections · gold: publications. Gold dashed: suggested; green solid: confirmed. Synthetic test publications.

EVIDENCE, STRONGEST FIRST

01

Citations in the document

"IAW AR 710-2, para 2-10" links straight to that paragraph.

02

The document type

Form number and layout point to the publications that usually govern it.

03

Shared vocabulary

Full-text ranking between the document and every regulation section.

04

Your reviewers

What was confirmed or rejected before, for the same kind of document.

AT HEADQUARTERS

Coverage, command-wide

Which units' files map to which publications, and which files nothing governs yet.

One reviewer teaches everyone

One confirmation improves the next suggestion for every desk.

Never a regulation it was not given

Every mapping cites a section in your loaded library.

Regulation packs

Import and export a library as a pack; a CSV of every mapping for inspectors.

AUDIT LOG · PROVENANCE

Ask the audit log, in plain words.

On the Mainframe every user's sign-in, import, look-up, question, verification result, decision and export goes into one append-only, hash-chained log inside the encrypted vault. Nothing in the product can rewrite it, an administrator included. Inspectors and security teams search it the way they would ask a colleague.

Audit log

Every unlock, import, look-up, question, verification result, decision and file export, kept inside the encrypted vault. Newest first.

Export the whole audit log: [Excel](#) [CSV](#) [PDF](#) [Provenance \(PROV\)](#) [Graph \(JSON\)](#)

which answers used the hand receipt?

[Ask](#) [List](#) [Graph](#)

3 matching entries.

2026-10-18 07:23 **What do the regulations say about hand receipts?**

BLUF: Most relevant record: [AR 710-2 2-11] 2-11. Hand receipts

[Why this answer](#)

2026-10-18 07:15 **What does the regulation say about hand receipts?**

BLUF: Most relevant record: [AR 710-2 2-11] 2-11. Hand receipts


[Why this answer](#)

2026-10-18 07:15 **Is there a scanned hand receipt in my files?**

BLUF: Yes. 1 scanned file(s) in your own records match "hand receipt": hand_receipt_scanned.pdf.


[Why this answer](#)

"Which answers used the hand receipt?" Search by meaning: filler dropped, words stemmed, the lexicon widens terms, and "withheld" or "last week" become filters. M.AI core, Field Kit v1.0.0.108, synthetic data.



THE LOG AS A GRAPH

Answers, the records they used and the links between them, in W3C PROV terms.



SEMANTICA

Provenance your own
tools can read

The log exports as Excel, CSV and PDF, and as **W3C PROV** (Turtle) and graph JSON, written on the server. Semantica, the open-source knowledge and reasoning project, loads the PROV export at the bench to query reasoning chains across the whole command's history. A published standard, so you are not locked to any one reader.

EXPORTS [Excel](#) [CSV](#) [PDF](#) [W3C PROV \(Turtle\)](#) [Graph JSON](#) [STIX 2.1](#)

SPEAKING THE LANGUAGE OF SUSTAINMENT

Fluent in the codes your people use every day.

M.AI is loaded with a sustainment lexicon built around the FMS environment. Each term carries its meaning, how it is spoken, and the context it applies in. When context is unclear, she asks.

TERM	MAINTENANCE / OPERATIONS	SUPPLY / TRANSPORT / CONTRACTING
MRO	Maintenance, Repair and Overhaul	Materiel Release Order
POD	—	Port of Debarkation · Proof of Delivery
PM	Preventive Maintenance	Program Manager
CM	Corrective Maintenance	Configuration Management
TO	Technical Order	Task Order

STATUS CODES, IN PLAIN LANGUAGE	PRIORITY AND CONDITION LOGIC
BA Being processed for release and shipment	PD 01-03 Issue Priority Group 1
BB Backordered	PD 04-08 Group 2
BV Procured, on contract for direct delivery	PD 09-15 Group 3
CA Rejected, with a reason to correct	PD is checked against FAD and need; priority inflation is flagged. Only condition code A is proposed for lateral transfer; code F drives carcass return and repair.

THIRTEEN LEXICON DOMAINS

FMS and security cooperation	Cataloguing and FED LOG	MILSTRIP / DLMS	Supply status and stock	Readiness
Maintenance (MRO)	Components and reliability	Aviation maintenance	Transportation	Warehousing and inventory
Discrepancy and quality	Repairable exchange	Procurement and sources		

FORMS SHE READS, CROSS-REFERENCES AND DRAFTS

DD 1348-1A	DD 1348-6	DD 1348-5	DD 1149	DD 1150	DD 1750	DD 1907	DD 2890	SF 364
SF 361	SF 368	DA 2404	DA 5988-E	DA 2407	DA 2406	DA 2408	DA 1687	

COMMAND METRICS IN EVERY BRIEF

FMC / PMC / NMC	NMCS / NMCM	Ao	CWT	Fill rate	Requisition aging	TAT	Backlog	RTS / RTM
NFF	MTBF / MTBUR / MTTR							

USE CASES · AVIATION AND GROUND MAINTENANCE

From the flight line to the motor pool.

Eight ways a maintenance organisation uses the Mainframe every day.

PHASE MAINTENANCE PLANNER

“Which aircraft enter phase in the next 30 days, and what parts are we short?”

Phase flow forecast against inspections due, PLL / ASL shortfalls and kits to order now.

PRODUCTION CONTROL

“What deferred discrepancies are open on the Apaches?”

Open and deferred discrepancies by tail, time-change items and inspections due, with the TM reference.

BATTALION MAINTENANCE OFFICER

“Which NMCS vehicles could I return this week with a lateral transfer?”

Matches condition code A stock to AWP vehicles, ranked by estimated return to service.

SHOP FOREMAN

“Which work orders have been AWP for more than 14 days?”

Work-order aging by document number and status code, with the next action for each.

RELIABILITY ENGINEER

“Show repeat faults on the hydraulic pump over 90 days.”

Repeat and no-fault-found removals grouped by serial, with MTBUR and the drivers of downtime.

COMPONENT TRACKER

“Trace serial XXXX-0412 from requisition to turn-in.”

One timeline across supply, transport and maintenance: receipt, install, removal, repair, return.

TECHNICAL INSPECTOR

“Which inspections fall due before the exercise, and which publications apply?”

Due list by tail and serial, each tied to the governing paragraph in the regulation library.

BATTLE-DAMAGE PLANNER (EXERCISE)

“Which damaged vehicles return fastest with the parts on hand?”

Ranks repairs by time to return to mission and flags controlled exchanges that need approval.

Illustrative scenarios built on M.AI's features and the white paper's use cases; each is scoped and proven in discovery and pilot. Synthetic data.

USE CASES · SUPPLY, WAREHOUSE AND DISTRIBUTION

Every requisition, bin and shipment, answerable.

From the receiving dock to the requisition desk.

SUPPLY SUPPORT ACTIVITY

“Which receipts need an SDR, and how many days are left?”

Short, wrong and damaged receipts with the time remaining; SF 364 drafts ready for review.

REQUISITION MANAGER

“List PD 01–03 requisitions more than 30 days past RDD.”

Aging bands with status codes and follow-up or cancellation drafts, ready for MOV.

WAREHOUSE SUPERVISOR

“Which locations failed the last cycle count?”

Variances by bin and NSN, the adjustments pending and a location-survey list for the shift.

TRANSPORTATION OFFICER

“Where is TCN ... and is it damaged?”

In-transit status from the loaded records, proof of delivery, and an SF 361 draft if needed.

STOCKAGE PLANNER

“Which ASL lines should change after 12 months of demand?”

Adds and deletes with demand history, lead times and the reorder point behind each.

HAZMAT AND SHELF LIFE

“Which items expire in the next 60 days, and where are they?”

Shelf-life list by location with the handling and disposal rules that apply.

REPAIRABLE EXCHANGE

“Which carcasses are overdue for return?”

Unserviceable due-ins past their window, with the exchange credit at risk on each.

PROPERTY BOOK OFFICER

“Which hand receipts haven't been reconciled this quarter?”

Open hand receipts by holder, sensitive items first, with the signature record that is missing.

Illustrative scenarios built on M.AI's features and the white paper's use cases; each is scoped and proven in discovery and pilot. Synthetic data.

USE CASES • AIR FORCE AND NAVY

Same assistant, every service.

Air force ranks and forms, naval terms, aircrew questions: configured per edition.

AIRCRAFT MAINTENANCE GROUP

“Which F-15s are down for parts, and on which NSNs?”

Aircraft held by supply with the requisition, priority and due-in for each part.

AVIONICS BACK SHOP

“Which LRUs came back no-fault-found this quarter?”

NFF rate by part and serial, so test procedures and stockage can be corrected.

AIRCREW

“Is my aircraft released for flight?”

Status, open and deferred discrepancies, inspections and time-change items due.

AIR MOBILITY / AERIAL PORT

“Which exercise pallets are still at the port of embarkation?”

Shipment units by TCN and priority from the loaded transport records, with the oldest first.

SHIP'S SUPPLY OFFICER

“Which casualty reports are waiting on parts?”

Mission-degrading equipment linked to requisitions, due-ins and repair actions.

NAVAL AVIATION

“Which aircraft are down for the same component?”

Common drivers across the squadron, cannibalization candidates flagged for approval.

ENGINEERING DEPARTMENT

“Open the maintenance history for pump serial ...”

Fault and removal history with FED LOG item data and the interchangeable NSNs.

MUNITIONS ACCOUNTABILITY

“Reconcile this month's issues against the signature records.”

Issue documents matched to DD 1907 signature records, with every mismatch listed.

Illustrative scenarios built on M.AI's features and the white paper's use cases; each is scoped and proven in discovery and pilot. Synthetic data.

USE CASES · HEADQUARTERS, FMS AND TRAINING

Where the Mainframe earns its seat.

The questions a headquarters asks across units, cases and languages.

COMMANDING GENERAL

“Brief me.”

A two-minute spoken picture across every unit that reports to the Mainframe, with decisions to take.

G4 / J4 STAFF

“Compare COAs for the surge to 85% by D+14.”

Options projected from past surges and current stock, with assumptions and risk for each.

FMS CASE MANAGER

“Which case lines are ready for closure review?”

Line status from the case records, open requisitions and SDRs that still block closure.

FINANCE OFFICER

“What is billed versus expended on case line 3?”

Financial position by line and a reconciliation list for the program review.

COALITION STAFF MEETING

“Summarise the readiness picture in Arabic and English.”

One answer, both languages, codes intact; spoken in each edition's own voice.

INSPECTOR GENERAL / QUALITY

“Show every record mapped to AR 710-2 for the inspection.”

The compliance map filtered by publication, confirmed and pending, exported to PDF.

SCHOOLHOUSE INSTRUCTOR

“Answer the new lieutenants from the regulations, with citations.”

Answers cite the paragraph; the instructor sees what was asked and where students struggle.

INFORMATION SECURITY OFFICER

“Who asked for data outside their unit this week?”

Refused requests from the audit log, with user, time and the rule that stopped them.

Illustrative scenarios built on M.AI's features and the white paper's use cases; each is scoped and proven in discovery and pilot. Synthetic data.

WHO M.AI SERVES

From flag officer to technician, from the same records.

M.AI adapts to role, echelon and authority: a strategic picture for a general, a precise work list for a sergeant.

ROLE	TYPICAL QUESTION	WHAT M.AI DELIVERS
Admirals, generals, air marshals	What is our readiness today and what is driving it?	Two-minute brief; top drivers; risk to mission; COAs
Commanders	Which assets can I return to service this week?	FMC / NMCS / NMCM picture; recovery options; priorities
Staff officers (J3/J4, G4, A4, N4, FMS)	Compare the COAs for the exercise surge.	COA comparison; case-line status; LOA explanations
Ship and naval engineering officers	Which casualty reports are waiting on parts?	Mission-degrading equipment linked to parts and repairs
Junior officers	What is the status of my platoon's vehicles?	Asset status, deferred faults, inspections due, draft reports
NCOs and warrant officers	Which work orders are AWP?	Work-order aging, parts status, PLL / ASL shortfalls
Pilots and aircrew	Is my aircraft released for flight?	Status, discrepancies, inspections and time-change items
Technicians and mechanics	Which NSN replaces this part?	Fault history; FED LOG data including I&S; TM references
Supply and warehouse staff	Which receipts need an SDR?	Discrepancies, SF 364 drafts, condition codes, locations
FMS case managers and finance	Which case lines are ready for closure?	Line status, financial position, closure readiness

THE MAINFRAME

Upper half: headquarters and staff

Many named users at desks and in the operations room, on the command network.

THE FIELD KIT


Lower half: away from a desk

One operator in the hangar, the warehouse or the forward site. Neither is limited to its half.

CONNECTED TO YOUR SYSTEMS

One assistant, one security model, over every ERP you run.

M.AI does not replace your ERP or maintenance system. She reads from it, and writes to it only through approved interfaces after a named person approves. The host stays the system of record.



Read-only first Writes only through approved interfaces.	No write without a name A named person approves each one.	System of record The host stays the truth.	Your dictionary Every field mapped in discovery.
--	---	--	--

ONE RECORD MODEL · THE FIELDS M.AI REASONS WITH

NSN	NIIN	P/N	CAGE	FSC	U/I	Document No.	DoDAAC	FMS case · line	Project code	PD
RDD	Status code	Due-in	TCN	Condition code	Serial / lot	Work order	Discrepancy	Fault code		
Removed / installed NSN	Man-hours	AWP	Inspection due	Flight hours · cycles	RTS estimate	QA sign-off				



M.AI COMPANION

M.AI beside the screen your people already use.

The M.AI Companion brings M.AI into the side panel of Chrome or Edge, next to your ERP, tracker or portal. Ask about your records without leaving the page, or let M.AI work on a tab you share, one approved step at a time.

Demo ERP Maintenance · 2nd Sustainment Brigade (synthetic)

Open work orders
Fleet: tactical vehicles · 35 assets

86% FMC rate **3** NMCS **2** NMCM **4** Open high-priority requisitions

WO	Asset	Status	Fault	NSN awaited	Opened
WO-26-0412	21-044	NMCS	Final drive leak	5330-99-000-1187	02 Oct
WO-26-0415	21-051	NMCS	Final drive leak	5330-99-000-1187	03 Oct
WO-26-0419	19-013	NMCS	Hydraulic pump	4320-99-000-4521	05 Oct
WO-26-0421	19-020	NMCM	Scheduled service	—	07 Oct
WO-26-0425	17-002	NMCM	Long-term repair	—	09 Oct
WO-26-0428	18-031	FMC	Closed: tyre change	—	10 Oct

Synthetic demonstration data. Not a real system.

M.AI

M.AI connected · 0 shared tab(s)

Brief me **Refresh data**

BLUF: Ask M.AI about your records, or enable task mode to work on the current web tab.

- FMC rate 86%, NMCS 3, NMCM 2, PMC 0.
- 4 open requisitions: 2 at PD 01-03, 2 past RDD, 1 backordered (BB).
- Oldest high-priority item past RDD: document TST001-6150-0001, NSN 5330-99-000-1187, 15 days late.
- No file in your folders is new or changed today.
- Chats today: 4 (5 questions). You asked about: "What do the regulations say about hand receipts?"; "Is there a scanned hand receipt in my files?"; "What does the regulation say about hand receipts?"; "Show all NMCS vehicles with an open high-priority requisition".

M.AI Companion beside a demonstration ERP page (Field Kit v1.0.0.108). Synthetic data; not a real system.

HOW IT WORKS

Pair in seconds

A short-lived code from M.AI Settings. No passphrase is shared with the browser.

Your edition, in the panel

The national persona, Brief me, Refresh data and voice, in Chrome or Edge.

Same answers, same checks

Records, regulations and verification exactly as in M.AI.

TASK MODE

Reads the tab you share

Only tabs you add to the purple M.AI group. Nothing is shared automatically.

One bounded step at a time

Visible cursor, typing, scrolling and evidence highlighted on the page.

You approve the clicks

Clicks and dropdown changes wait for your review. Stop takes over at once.

GUARDRAILS

Nothing leaves M.AI

All model processing stays on your own hardware; websites keep their own traffic.

Access ends cleanly

Removing a tab, Stop or locking ends the session.

Recorded

Every page action goes to the encrypted audit log. Passwords and file pickers stay with you.

Every desk, one assistant

Staff keep their ERP, tracker or portal open and ask M.AI in the side panel.

Answers from the Mainframe

The panel answers from the command's Mainframe under each user's own permissions. ♦

Works with the systems you have

Web ERPs, portals and trackers: nothing to install on the host system.

THE M.AI LLM

One family, two sizes, one discipline.

The Mainframe runs mai-mf-llm-26b: a sparse model that answers at command scale while activating only a fraction of itself per answer. It is built in layers, from your own sources upward, and released only through the gate.

mai-fk-llm-4b* M.AI FIELD KIT SIZE Compact class, quantised for the device RUNS ON The kit's own AI accelerator CONTEXT A question, a record set, a regulation	mai-mf-llm-26b* M.AI MAINFRAME SIZE 26B sparse, about 4B active per answer RUNS ON The site GPUs, serving many users CONTEXT LOAs, publications, case files, at command scale	nora-mf-llm-26b* NORA MAINFRAME SIZE The same 26B foundation RUNS ON The site GPUs CONTEXT Enterprise questions, in the customer's vocabulary
--	---	---

HOW THE M.AI LLM IS BUILT AND TRAINED, LAYER BY LAYER



WHAT IS NEXCOMMAND'S, AND WHAT IS NOT

The sustainment adapter

Lexicon, doctrine and tool discipline. NexCommand's own.

The tool library

The model decides which to call; the tools decide what is true.

Verification

Every NSN, quantity, date and status checked before it is shown or spoken.

The evaluation gate


Exact-match on codes and figures, in English and Arabic, or no release.

* Each M.AI LLM is adapted by NexCommand from an open-weight foundation model of U.S. origin; licences and notices are provided in the product. ♦ mai-mf-llm-26b is in development.

REFERENCE ARCHITECTURE

The same core as the Field Kit, at command scale.

A core that orchestrates each request, enforces permissions and writes the audit trail; open-weight models for language, speech and search; and one strict rule: figures come from records, not from the model.



THE LIFE OF ONE REQUEST

01	02	03	04	05	06
Ask	Identify	Plan	Retrieve	Verify	Answer
Typed or spoken; speech transcribed on the server.	The Core loads role, unit, language and permissions.	The model plans and calls approved tools only.	Each call runs with the user's own permissions.	Every NSN, quantity, date and status checked.	Text and voice, sources, reasoning, options.


COMPONENT	TECHNOLOGY	ROLE
M.AI Core	Service on your server	Orchestration, permissions, tool library, audit and reasoning trail
Model runtime	On the site GPU workers	Serves many users at once; stateless workers
Language model	mai-mf-llm-26b*	Understands requests, selects tools, composes answers
Knowledge index	PostgreSQL with vector search, multilingual	LOAs, technical publications, SOPs, minutes
Connectors	One per host system	Map host data to the NexCommand record model
Speech	Local speech-to-text and voices	English and Arabic; codes pronounced net-style
3D persona	Rendered in the browser	Idle, listening, thinking, speaking; lip-sync
Customer Console	Separate install	Licences, users and seats, updates to Field Kits, sync

* Open-weight foundation model, adapted by NexCommand; notices in the product. ♦ In development: shown as designed.

THE AIR-GAPPED AI CENTER

Engineered as a room, not just a server.

Your Authorized Technical Partner delivers a complete offline site: room, power, cooling, firewalls, service hosts, storage, backup, the Customer Console and its media workflow, then the AI workers. No variant connects the AI center to the Internet.



SIZED FOR YOUR USERS · PLANNING REFERENCE

NAMED USERS	PEAK REQUESTS (20%)	DUAL-GPU WORKERS	SERVICE HOSTS	RACKS	IT LOAD kW
100	20	2	2	2	5.1
500	100	7	2	3	10.4
1,000	200	13	3	5	17.5
3,000	600	38	4	11	44.8
5,000	1,000	63	6	18	72.9
10,000	2,000	125	10	35	142.1

≤ 3 s p95 first token, warm	20 tokens / s per stream	8K token context	24 h initial recovery point
---------------------------------------	------------------------------------	----------------------------	---------------------------------------

Continuity. Nothing depends on the internet, so losing outside links changes nothing. Field Kits keep working with what they hold and reconcile when they next reach the server. Backups carry the vault's encryption; restoring on new hardware is a licensed operation.

Planning reference from the M.AI Mainframe Partner Site Engineering Guide, Rev 3 (NC-MF-ENG-001), at 20% simultaneous generation. Quantities are confirmed after a measured prototype; site design by your partner.

SECURITY AND RESPONSIBLE AI

Security is the product.

Adding M.AI must not widen the attack surface of your ERP or weaken its access model. These are the controls, and the principles behind them.

■ Air-gapped

All processing on your hardware; no internet, no cloud, no telemetry. Updates on approved, signed media.

■ Transactions

Read-only by default; drafts are submitted by people through host approvals.

■ Prompt-injection defence

Text in documents is data, never instructions; only approved tools can be called.

■ Signed everything

Licences, updates and model packs signed by NexCommand; a server refuses anything else.

■ Access control

Every lookup runs with the user's identity and permissions; documents carry access labels.

■ Answer integrity

Figures only from returned records, verified; missing data reported as missing.

■ Audit

Every request, tool call, answer and approval logged, hash-chained, exportable.

■ Assurance testing

National evaluation sets, red-team and regression tests on every release; STIX 2.1 export.

U.S. DOD AI ETHICAL PRINCIPLES · NATO PRINCIPLES OF RESPONSIBLE USE

Responsible

Humans remain accountable for every decision.

Equitable

Training data curated; recommendations tested across units.

Traceable

Sources, assumptions and confidence with every answer.

Reliable

Evaluation, regression and adversarial testing before release.

Governable

Administrators can restrict tools or take M.AI offline at any time.



ISO 27001

NexCommand is certified for information security management.

ISO 9001

Certified quality management behind every release.

ITAR · EAR

LOA technical data stays inside the authorised environment.

NDAA §1532

U.S.-origin open-weight foundations only.

ASSURANCE · PRIVACY

She attacks herself. She remembers nothing.

Two promises a security officer can check: a built-in attack test that reports in a standard format, and a model that holds no memory of your data.

The built-in security test, reported in STIX 2.1 13 CHECKS · MITRE ATT&CK

The test runs on the server with nothing extra installed. Each check tries one thing an attacker would, records what held, and maps it to its ATT&CK technique and the M.AI control that stops it. The result exports as a STIX 2.1 bundle your security team reads in its own tools.

T1078	Using it without signing in	T1110	Guessing the passphrase	T1539	Stealing a session
T1189	Framing the page	T1083	Looking outside its folder	T1105	Path tricks on upload
T1190	SQL injection	T1133	Reaching an internal service	T1041	Exfiltration through the page
T1070	Editing the audit log	T1485	Wiping the data	T1499	Exhausting resources
T1005	Reading data at rest				

Known limits, stated plainly. The test is a set of dated checks on one version, not a certificate. Run it after installation with production settings (disk encryption and TLS on); development settings are expected to fail several checks.

NO MEMORY IN THE MODEL

Each answer starts clean

The M.AI LLM keeps nothing between questions. It is handed the records for this question and nothing else, then lets go.

The thread lives in your vault

Follow-ups are read from the answers already saved in the encrypted vault, never from the model's memory.

History finds sources, not facts

An earlier answer can point to a document. It is never treated as evidence for a new one.

Your data trains nothing elsewhere

Nothing goes to NexCommand, a cloud or another customer. Model updates arrive as signed packs, installed at your choice.


LEARNING INSIDE YOUR BOUNDARY

Reviewer decisions, kept local. Confirmed compliance mappings, corrected lexicon terms and rated answers stay on your server and improve retrieval for your users only.

A national adapter, at your bench. Where a nation wants its own adapter, it is trained on its own hardware and must pass the same release gate before it reaches users.



Security test and STIX 2.1 export are in the M.AI core (Field Kit v1.0.0.108; the Mainframe's core carries the same modules). Results shown are categories, not a pass report.



Prove it on your own data.

A short discovery, a demonstration with your own persona, and a pilot at one site, measured against agreed targets.

2-3 WEEKS

Discovery

Data sources, permissions, edition, success measures.

1 WEEK

Demonstration

Your persona on a slice of your real data.

8-12 WEEKS

Pilot

One site or fleet; read-only; agreed measures.

3-6 MONTHS

Expansion

More fleets, sites, users; approved write-backs.

ONGOING

Steady state

Signed updates, editions, languages, Field Kits.

ROLE-BASED TRAINING, IN YOUR LANGUAGE, ON YOUR DATA

AUDIENCE	TRAINING	LENGTH
Senior leaders	Capabilities, limits, decision authority	1 hour
Staff officers, case managers	Briefs, COA analysis, sources and confidence	1 day
NCOs, technicians, supply	Daily questions, work orders, SDR drafts, FED LOG	Half day
Administrators, J6	Install, access, updates, audit review	3 days
Train-the-trainer	Teaching M.AI in schools and NCO academies	2 days

FOURTEEN NATIONAL EDITIONS

- United States
- United Kingdom
- Australia
- Taiwan
- Saudi Arabia
- UAE
- Bahrain
- Kuwait
- Qatar
- Egypt
- Nigeria
- Philippines
- Kenya
- Tanzania



Request a demonstration
Read the M.AI white paper and the Field Kit brochure.
nexcommand.ai