VOL.07
ISSUE.69

CENJOWS
WEB ARTICLE

CYBER POWER AND ALGORITHMIC OPERATIONS

MR ABHINAV SHARMA



www.cenjows.in

CENJOWS WEB ARTICLE

Volume 07 Issue 69 NOV 2026

CENJOWS

(Established: 2007)

Room No 301, B-2 Wing, 3rd Floor
Pt Deendayal Antyodaya Bhawan,
CGO Complex, Lodhi Road
New Delhi –110003 (INDIA)
Telephone Nos: 011-24364881, 24366485
Fax: 011-24366484
Website: www.cenjows.in
E-mail: cenjows@cenjows.in

ABOUT US

CENJOWS was raised in 2007 as an independent think tank, registered under the Societies Registration Act, 1860. This aims to promote Jointness as a synergistic enabler for the growth of Comprehensive National Power and provide alternatives in all dimensions of its applications through focused research and debate.

Starting Year : 2020

Year of Publication : 2026

Frequency : Weekly

Language : English

Publisher : Maj Gen (Dr) Ashok Kumar, VSM (Retd) Director General, CENJOWS

301, B-2 Wing, 3rd Floor
Pt. Deendayal Antyodaya Bhawan CGO Complex,
Lodhi Road
New Delhi-110003

Editor : Maj Gen (Dr) Ashok Kumar, VSM (Retd) (dg@cenjows.in)

Deputy Editor : Dr Ulupi Borah, Distinguished Fellow (distinguishedfellow1@cenjows.in)

Editorial Board :

Brig K Ranjeev Singh, YSM (ddg@cenjows.in)

Col Ajay Kumar Rajak, SM

(secy@cenjows.in)

Gp Capt Ashish Kumar Gupta (Retd)

(seniorfellow2@cenjows.in)

Dr Ulupi Borah, Distinguished Fellow

(distinguishedfellow1@cenjows.in)

301, B-2 Wing, 3rd Floor

Pt. Deendayal Antyodaya Bhawan

CGO Complex, Lodhi Road

New Delhi-110003

Dr Monojit Das, Senior Fellow, PRO

(monojit.das@cenjows.in)

Publication Head : Ms Vaibhavi Katal

All correspondence may be addressed to:

Editor
Centre for Joint Warfare Studies (CENJOWS) 301, B-2 Wing, 3rd Floor
Pt Deendayal Antyodaya Bhawan CGO Complex,
Lodhi Road
New Delhi-110003
Telephone: (91-11) 24366485/Telefax: (91-11) 24366484 e-mails:
cenjows@cenjows.in / cenjows@yahoo.com Website: http://cenjows.in

© Centre for Joint Warfare Studies

Price: Rs 150/- INR

All rights reserved. No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording or by any information storage and retrieval system without permission from the DG, CENJOWS, New Delhi.

CYBER POWER AND ALGORITHMIC OPERATIONS

CONTENTS

Title	Page No
Introduction	**5**
How States Approach AI in the Cyberspace	**6**
AI in Cyber Operations: Local Models, Agents and Orchestration	**10**
Adversarial AI: When the Model Is the Target	**15**
Securing Military Networks and Tactical Data Links	**16**
Recommendations	**19**
Conclusion	**20**
References	**23**

CENTRE FOR JOINT WARFARE STUDIES

CENJOWS

Cyber Power and Algorithmic Operations



Mr Abhinav Sharma is a Technical Research Assistant at CENJOWS

Introduction

For a long time, cyber power was judged by what a state possessed: a stock of zero-day exploits, skilled human red teams and large stores of data. That test is losing its value. Machine learning now allows a defender to sort thousands of alerts in minutes and allows an attacker to probe a network without waiting for a human operator. What counts more and more is how quickly a state can act, how much of its work it can safely automate, and whether it controls the models doing that work.

This brief makes one argument, and the title states it. Cyber power is a state's ability to act in cyberspace and to stop others from doing so. Algorithmic operations are the cyber activities that are carried out by the software at machine speed, under human instruction. The two are now joined: cyber power is in fact the ability to run algorithmic operations better, faster and more independently than the adversary. That ability rests on three requirements, which we will discuss in each following section. Speed is how fast a system is at detection, decision and action. Control or sovereignty is what we mean by owning the models, data and hardware. Trust is that the system will perform

as expected, the system's output, even when an adversary is attacking it. India faces all three questions at once, with active threats on two borders and a long habit of importing its technology.

Each section takes up these requirements in turn. Section 1 asks who is already pursuing algorithmic speed and how far each state controls the technology behind it. Section 2 shows how such systems are built when control is kept at home. Section 3 asks whether they can be trusted by examining attacks on the models themselves. Section 4 applies all three requirements to military networks and tactical data links, with India's own systems as the case. The conclusion brings together the threads, and we present them in six action points. At the end a glossary of technical terms is provided for the readers who need it.

How States Approach AI in Cyber Warfare

Section 1 asks the first question: who is already using algorithms for cyber advantage, and how much of the technology do they control? Each state shapes its use of AI in cyberspace according to its military strategy, its geography and the culture of its institutions. The five cases below differ in detail, but each show how much now depends on speed and control.

- **The United States: Defending Forward**

  The United States cyber strategy rests on "persistent engagement" and "defending forward", ideas developed by US Cyber Command (USCYBERCOM).[1] The reasoning is that waiting behind a perimeter does not work against capable adversaries. American teams therefore operate inside foreign networks to disrupt threats before they reach US systems. Automation supports this approach in three ways.

  - **Autonomous Cyber Reasoning Systems (CRS):** From the DARPA Cyber Grand Challenge to the military's advanced R&D programmes, the U.S. has been working on autonomous software which does this work. These systems break down binary code and trace executable logic defects via symbolic execution (which maps out all of a program's paths that present security issues). They also develop practical vulnerability

patches and deploy them across military systems before human input can be added.

- **Algorithmic "Hunt Forward" Operations:** Cyber protection teams use specialised algorithmic triage tools which go through terabytes of host telemetry and packet captures (the raw data of what is travelling across a network) in allied networks. They identify subtle behavioural anomalies which expose persistent state threats.
- **Foundation models for planning:** Secure, containerised language models are used to combine intelligence reports into a single picture for targeting. The main aim is to cut planning time from weeks to hours.

- **China: Intelligentised Warfare**

  The People's Republic of China (PRC) has introduced the concept of Zhinenghua (Intelligentised Warfare), which it uses as the base of its path to future military dominance. The PRC presents its investment in AI, autonomous cyber operations and quantum technologies as the key to that dominance. After large-scale modernisation of the PLA's structure, which included information operations, the PLA Cyberspace Force has become a strategic branch meant to paralyse adversary C4ISR systems.

  Chinese algorithmic cyber doctrine is governed by **Systems Destruction Warfare** (*Tixi Pohuai Zhan*):[2].

  - **Data collection at scale:** Chinese state-backed threat groups use machine learning in campaigns to collect, process and share large sets of stolen intellectual property, telecom metadata and national infrastructure blueprints.
  - **Pre-emptive paralysis:** In a Taiwan Strait or Himalayan conflict scenario, the Chinese do not see cyber operations as a form of spying alone. Autonomous cyber entities and coordinated logic attacks are used to go after an adversary's joint command networks, electrical grid and logistics systems before the first kinetic action is taken.

- **Russia: Information Confrontation**

  Russian thinking treats technical network operations and psychological or influence operations as parts of one whole, which it calls information confrontation. Western doctrine, by contrast, usually keeps the two apart.[3] Russian state-linked organisations (the GRU, SVR and FSB) use machine learning as an equaliser against stronger opponents.

  - **Disruption of infrastructure:** Russian-linked groups have repeatedly gone after Ukrainian power, telecom and government networks. Machine learning is expected to help them choose and strike targets faster, and to cause disruption on a larger scale.
  - **Automated deception:** Machine learning is also used in very targeted phishing campaigns, tailored to the interest of the target audience, and in the continuous improvement of malware so that it can get past Western detection systems. It also supports grey zone operations, which wear away at an adversary's political will without pushing them to a large-scale response.

- **Israel: Predictive Interdiction**

  Israel has little strategic depth and very short warning times. It has therefore built a doctrine of pre-emption and close integration between cyber and kinetic action. Unit 8200 of the Israel Defence Forces and the National Cyber Directorate are the main institutions behind it.[4]

  - **Predictive targeting:** Using advanced algorithmic intelligence platforms, the Israeli defence forces put together multi-source signals intelligence (SIGINT) from the dark web, telemetry data and intercepted communications, and present it in real-time dynamic target graphs.
  - **Cyber–kinetic coupling:** In Israeli doctrine, automated cyber and electronic warfare systems are connected to tactical strike assets. An identified cyber weakness may be used to blind an adversary's radar installation or command link at the exact microsecond that an air defence missile or precision-guided munition enters the area of operation.

- **India: From Reactive Defence to Credible Deterrence**

  India faces pressure from two sides. Along the northern borders and as far as the eastern islands, it sees advanced, state-level activity linked to the PLA. To the west, state-sponsored proxy groups run persistent phishing, defacement and credential-theft campaigns against defence establishments.[5]

  For many years, civil and military cyber defence in India worked separately, with little cooperation. The creation of the Defence Cyber Agency (DCyA), working alongside the National Critical Information Infrastructure Protection Centre (NCIIPC) and CERT-In, marks a move towards a more forward-leaning posture. Two points deserve attention.

  - **A maturing doctrine.** Indian security experts now accept that static perimeters such as firewalls cannot guarantee the integrity of a network. India's first Joint Doctrine for Cyberspace Operations, released in 2024, reflects this view.[6]
  - **The sovereign technology dilemma.** In the past, India's cyber architecture was very much a product of import. There was great dependence on commercial off-the-shelf networking equipment from abroad, foreign microprocessors and Western proprietary operating systems. As a result, India's approach to reliable deterrence is through the rapid development of indigenous network defence solutions, sovereign foundation models and cryptographically proven national communication infrastructure.
  - Two lessons follow from these cases. First, every state studied is pursuing speed, and each links it to deterrence or to winning the opening moves of a conflict. Second, the states differ most in control. The United States, China and Israel build their own technology, while India has relied on imports, which is its main gap. This is why Section 2 describes a sovereign design, built around models and hardware that a military owns and can inspect.

AI in Cyber Operations: Local Models, Agents and Orchestration

Section 1 showed that speed and control decide who gains an advantage. Section 2 explains how a military can have both. AI in cyber operations has moved from experimental prediction models to working software. These systems combine locally hosted open-weight models, standard ways of connecting models to tools, and controlled environments for testing what the models produce. Figure 1 shows the overall design. Each layer exists to deliver speed while keeping control in the military's own hands.

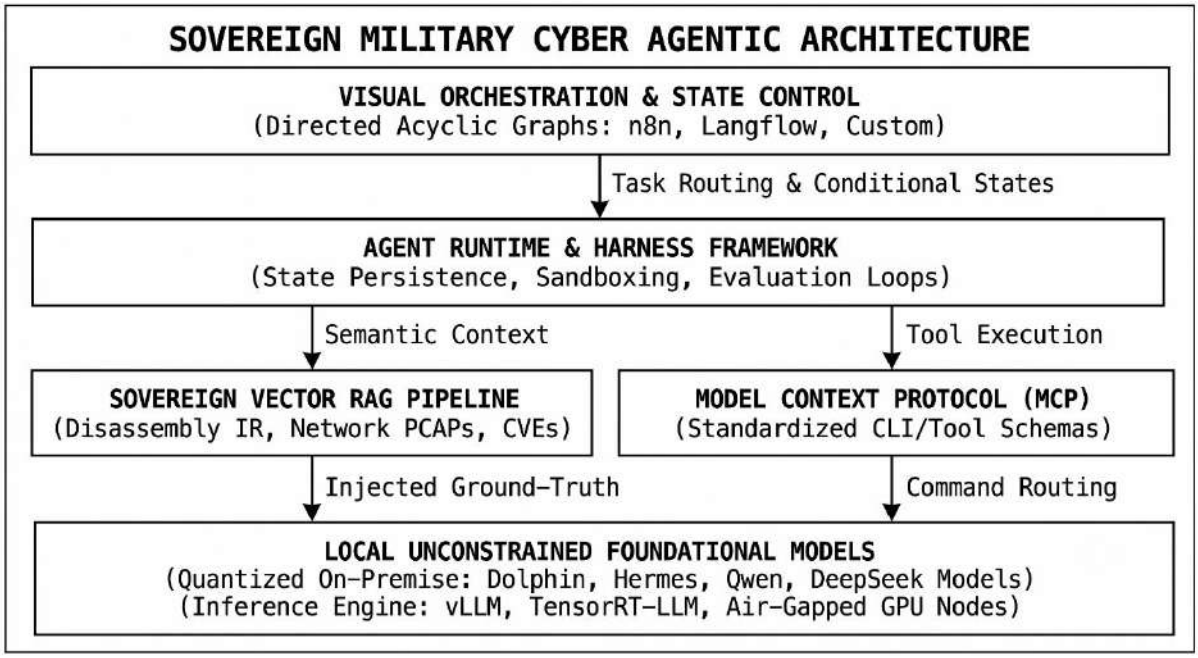


Figure 1: Sovereign Military Cyber Agentic Architecture **Source:** Author's own design built with the help of AI, 2026

Why Militaries Need Local Models

Commercial frontier models that run in the cloud are a poor fit for military cyber work for three reasons.

- **Operational security:** Sending classified telemetry, sensitive network data or proprietary code to a commercial provider creates an obvious route for leaks.
- **Usage restrictions:** Consumer AI services refuse many security tasks, including vulnerability research, reverse engineering and exploit development, which military teams must be able to carry out.

- **Air-gap compatibility:** Many critical systems have to run in isolated environments with no connection to the public internet.

The answer is to host open-weight models on the military's own hardware. Three practices matter here.

- **Model choice:** Open-weight families such as Qwen, which is strong at code and mathematical reasoning, and Hermes, which is tuned for calling tools and handling multi-step tasks, can be downloaded, hosted on premises and run inside secure enclaves.[7,8] Community variants such as the Dolphin series have had their refusal behaviour relaxed, which suits security work but makes it all the more important to check where the weights come from.
- **Efficient fine-tuning:** Methods such as Low-Rank Adaptation (LoRA) and its quantised form, QLoRA, keep the main model frozen and train only a small added layer. A team can adapt a model to its own network telemetry, legacy SCADA logic or known attack patterns at modest cost, without losing the model's general reasoning.
- **Local inference:** Quantisation shrinks the memory a model needs, for example, by storing it at 4-bit or 8-bit precision. Runtimes such as vLLM and TensorRT-LLM then let it run on rugged edge servers and field compute nodes.

The Building Blocks

A model on its own is not an operational system. Four components turn it into one.

- **Retrieval-augmented generation (RAG):** Language models sometimes state falsehoods with great certainty. RAG reduces this problem by making the model to use a local vector database which it queries before responding. The database includes technical documents, vulnerability records (CVEs) and reverse-engineering notes, so the model works from evidence rather than guesswork.
- **Model Context Protocol (MCP):** Connecting of models to external software used to require custom work for each tool. With MCP we have an open standard which provides a single uniform interface for models to use. Via these models, you may query a network scanner, use a debugger, access directory services or modify a firewall rule.

- **The harness:** Language models are not fully predictable. A harness therefore runs whatever the model produces, whether a script or a defensive patch, inside a disposable sandbox such as a microVM or container. It checks the result and sends any errors back to the model for another attempt. Nothing reaches a live system unchecked.
- **Visual orchestration:** In modern systems that have moved away from classic monolithic scripts, operators make use of visual workflow orchestrators like n8n and Langflow. On these platforms, complex cyber kill chains or incident response playbooks are presented as Directed Acyclic Graphs (DAGs), which are deterministic step-by-step decision trees. When an agent at Node A is unable to access an endpoint, the orchestrator passes the workflow to Node B, where a different enumeration agent is brought in, or it goes to a human commander.

Offensive Use: Compressing the Kill Chain

Used offensively, multi-agent frameworks can shorten the cyber kill chain, the sequence of stages from first reconnaissance to final effect, from weeks to minutes (Figure 2). Three stages show the change.

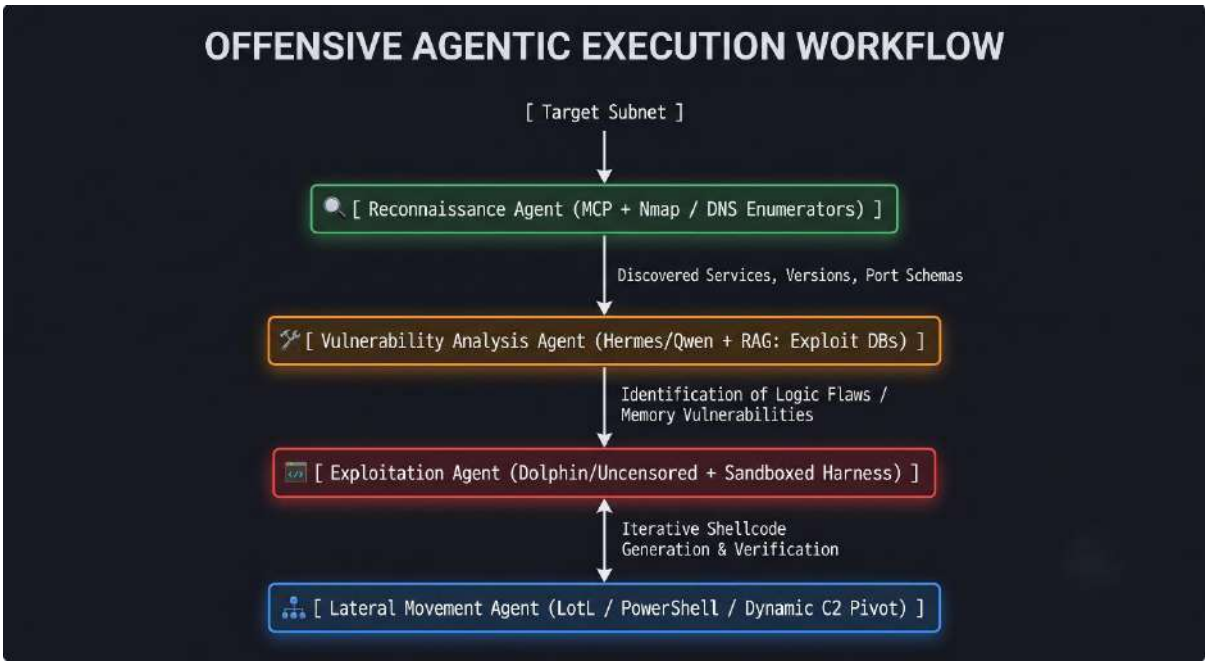


Figure 2: Offensive Agentic Execution Workflow **Source:** Author's own design built with the help of AI, 2026

- **Autonomous Reconnaissance:** An agent put together in n8n performs port sweeps and banner grabbing (gathering the software version strings put out by exposed ports) via MCP tools. Instead of relying on rigid version matching, the model looks at the semantic relationships between exposed microservices, which helps identify misconfigurations and systemic vulnerabilities that standard scanners do not see.
- **Reverse Engineering and Exploit Synthesis:** Target binaries are broken down into mid-level code, which is then studied by the code models. The models look at weak memory allocations, race conditions or unhandled exceptions. The models also put forth custom assembly scripts, and the local harness reports on whether the exploit overcomes memory protections such as Address Space Layout Randomisation (ASLR) and Data Execution Prevention (DEP/NX).
- **Dynamic Living-Off-the-Land (LotL) and Evasion:** Instead of deploying identifiable malicious programs to the disk, a module puts out custom system commands that use legitimate system binaries (PowerShell, WMI, bash). By analysing the target host's local event log activity level, the module varies its rate of action and randomises memory footprints as it moves laterally through the network, so as to evade endpoint detection sensors.

Defensive Use: Faster Triage and Self-Healing Networks

The same components can be used for defence. They address two long-standing weaknesses of security operations centres (SOCs): alert fatigue and slow analysts (Figure 3).

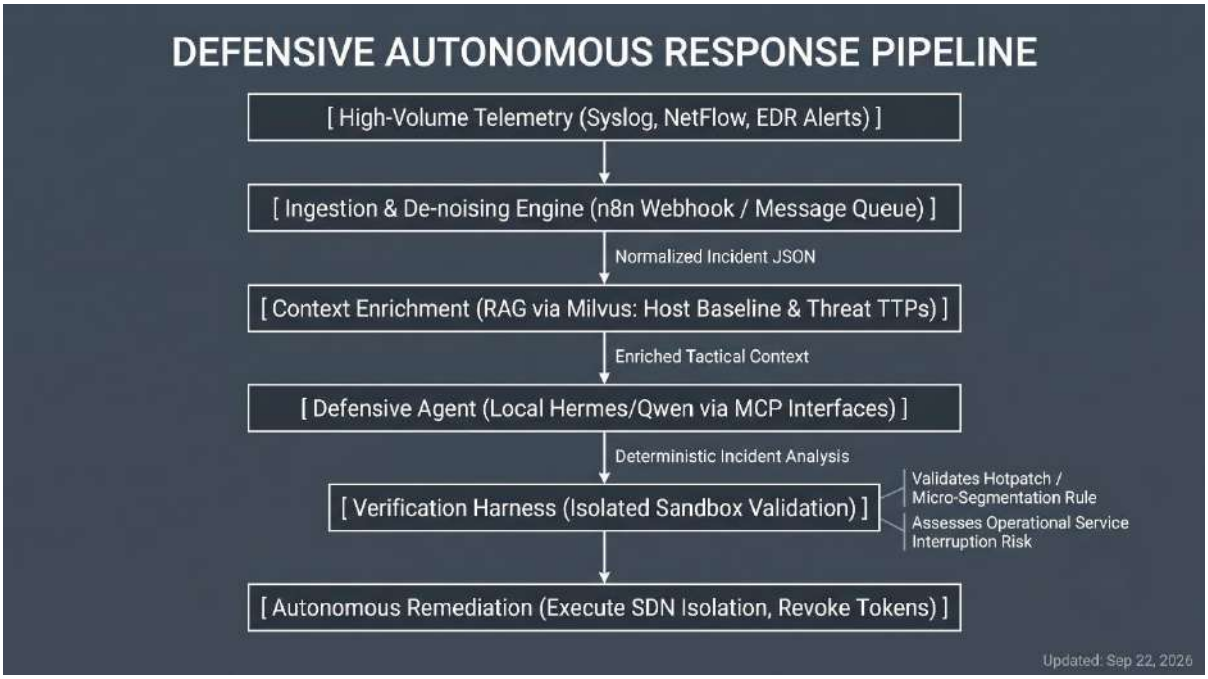


Figure 3: Defensive Agentic Execution Workflow **Source:** Author's own design built with the help of AI, 2026

- **Alert triage:** The orchestrator collects alerts from firewalls and endpoint sensors. The model which is used to do a comparison with a baseline of normal behaviour that we have obtained via RAG. This, in turn, separates out routine admin actions from credential theft and also clears noise from the queue.
- **Automated forensics:** When there is an issue which raises an alarm, an agent uses MCP to gather memory artifacts from the host, to list its live network connections and to check in with threat intelligence sources. It then puts together a timeline of the intrusion.
- **Validated remediation:** The agent puts forth certain responses. As an example of this is micro-segmentation, which we see to isolate an infected subnet via software-defined networking, and which also stops an intruder from moving laterally. In the case of a zero day in a critical system, the agent will draft a hot patch, test it out in a separate microVM for stability, and then roll it out to the network. What we see is a network which largely does the healing itself, which is very handy, although actions which have a large operational impact still should require human approval.

Sections 1 and 2 have dealt with speed and control. Both depend on the models, and this brings a new kind of risk. A system that acts at machine speed is only useful if its

decisions can be trusted, and the models can't be attacked. That is the subject of the next section.

Adversarial AI: When the Model Is the Target

Section 3 tests the third requirement, trust. As machine learning enters C4ISR, radar signal processing and endpoint defence, the models themselves become targets. The field of adversarial machine learning studies how an attacker can exploit the statistical weaknesses of neural networks.[9] Three kinds of attack matter most for a military that adopts the systems described in Section 2.

Data Poisoning and Supply-Chain Compromise

A model is only as reliable as the data it learns from. In a poisoning attack, the adversary quietly alters a small share of the training data. In a military setting this could be synthetic aperture radar (SAR) imagery, underwater acoustic signatures or network flow records, changed gradually over several months.

- **Backdoors:** The adversary introduces a latent variable, which may be termed a "trigger" or "trojan", into the poisoned data. In usual operation the model performs almost flawlessly. But when the model is exposed to the trigger, which may be an ordinary-looking sequence of bytes in a network packet header, it misclassifies the malicious traffic as benign. This creates a permanent and invisible backdoor.
- **Supply-chain exposure:** Present-day foundation models use large-scale and mostly unvetted public data sets, pre-trained base checkpoints and open-source code libraries. Third-party models which defence systems use therefore inherit supply-chain poisoning, which gets past traditional static software audits.

Evasion Attacks

Evasion attacks happen after training, when the model is in use. The attacker adds a tiny, calculated change to an input so that the model classifies it wrongly. The Fast Gradient Sign Method (FGSM) does this by working out how the model's error changes with the input and then moving the input a small step in the direction that increases the error.[10]

$$x_{adv} = x + \varepsilon \cdot sign(\nabla_x J(\theta, x, y))$$

In this expression, x is the original input, y is its true label, θ stands for the model's parameters, J is the loss function, and ε is a small number that limits the size of the change. Projected Gradient Descent (PGD) applies the same idea in repeated small steps and is generally stronger.

To human operators, a file or an electronic warfare signal looks exactly the same in its native form. But for a defensive ML classifier, that same input has been shifted by subtle adversarial perturbations, which move the feature vector across the classification decision boundary. As a result, a malicious exfiltration script or a weaponised exploit is put into the same category as normal background network traffic, often with high accuracy.

Model Extraction

An adversary with access to an exposed or edge-deployed model can query it repeatedly and record the outputs and confidence scores. From these, the adversary builds a "shadow model" that copies the behaviour of the real system. Attacks can then be refined against the copy offline, leaving no trace on the defender's network.

Adversarial examples often transfer from one model to another. An attack that works on the shadow model therefore has a good chance of working on the real one at the first attempt, and the defender receives no early warning.

None of this is an argument is against adopting AI. It shows that speed without trust is a liability. Trust can only be earned by controlling the data, the model weights and the testing process, and so the discussion returns to sovereignty. Section 4 now applies speed, control and trust to the setting where they matter most.

Securing Military Networks and Tactical Data Links

Sections 1 to 3 set out the principles. Section 4 applies them at the forward edge, where military cyber resilience is tested hardest. There, tactical networks, command posts and weapon systems have to keep working under constant electronic and algorithmic attack. The same three requirements apply: networks must respond at speed, be built from components the military controls, and give commanders data they can trust.

The Air-Gap Myth

In modern warfare, a perfect air gap is no longer realistic.

- **A connected battlespace:** Modern warfare relies on continuous joint coordination: sensor-to-shooter loops which include forward-edge compute for artillery coordination, unmanned aerial vehicle (UAV) video feed reports and live common operational pictures (COPs). These require constant connections across wireless networks, tactical radios and maintenance ports.
- **How the gap is crossed:** Modern air-gapped networks are infiltrated via hardware maintenance interfaces, infected USBs, unverified firmware updates and side-channel radio frequency injection. Once an adversary gains access to the air-gapped intranet's perimeter, the internal structure is very much at risk.

Tactical Data Links and Software-Defined Radios

Tactical Data Links (such as Link 16, Link 22 and proprietary sovereign systems) form the nervous system of joint air-land-sea combat. These communications are increasingly targeted by combined cyber-electronic attacks:

- **Algorithmic jamming:** Modern adversaries use machine learning in software-defined radios (where frequencies and modulation schemes are determined by software) to study spread-spectrum frequency hopping in real time. The model they use works out the pseudo-random hop sequence and then directs jamming energy or false data packets into the transmission window.
- **Track spoofing:** By inserting cryptographically tampered sensor reports into the tactical data network, an attacker can alter the target database within a theatre command structure. This may lead to phantom radar contacts, the masking of actual strike assets or friendly-fire events.

Protecting these links needs a deterministic zero-trust approach at the tactical edge, with two main elements.

- **Multi-sensor cross-validation:** Tactical systems must not accept data packets as they come in. Before a radar track is put on the command display, it must be run through a series of independent checks by edge algorithms. A track is put in quarantine if it appears only in RF reports without confirmation in other domains

(e.g., thermal, optical, and electronic support). If there is no cross-domain confirmation, the track is a probable spoofing attempt.

- **Post-quantum cryptography:** Quantum computers will eventually be able to break RSA and elliptic-curve cryptography, which protect most military data links today. Networks should begin adopting the algorithms standardised by the US National Institute of Standards and Technology (NIST): ML-KEM, derived from CRYSTALS-Kyber, for key exchange, and ML-DSA, derived from CRYSTALS-Dilithium, for digital signatures.[11] Designers must allow for the limited memory and tight timing of field radios.

India's Defence Communication Networks

India's military communications include the Defence Communication Network (DCN), the armed forces' Network for Spectrum (NFS) and Project ASCON, the Army Static Switched Communication Network. They work in settings that range from high Himalayan terrain to the deep sea.[12] Securing them means addressing three issues.

- **Foreign supply chain exposure:** A large issue with India's legacy infrastructure is that it has for a long time depended on imported telecommunication equipment, foreign-designed routing silicon and closed-source proprietary firmware. These elements create structural issues, among them hidden hardware trojans and firmware backdoors. For sovereignty and resilience, India needs to replace foreign core networking elements with home-grown ones that have gone through a cryptographic audit. This includes the efforts under the India Semiconductor Mission and indigenous processor initiatives (for instance, the Shakti and VEGA microprocessor programmes).
- **Zero-trust micro-segmentation:** Service branches, theatres of operation and forward-deployed teams must transition away from traditional perimeter security to zero-trust micro-segmentation. Every request, from an air-base terminal to an infantry field device, must be dynamically authenticated with behavioural metrics and given the least privileged access required for operation.
- **Quantised edge models:** Forward-based teams operate in Contested, Degraded and Low-Bandwidth (CDIL) environments, which have little to no cloud connectivity. India's defence R&D organisations should focus on developing quantised, low-power machine learning models which run on edge

devices, enabling front-line units to detect and neutralise intrusions at machine speed.

These three measures map directly onto the argument of the brief: indigenous hardware gives control, micro-segmentation gives trust, and edge models give speed. The conclusion brings these points together.

Recommendations

To turn algorithmic cyber power into capability, national security leaders and joint planners should give priority to six actions. Each serves speed, control or trust, and each rest on a case made earlier in the brief.

The first action, which serves control, is to deploy sovereign on-premises compute (Section 2). All military AI models should run on secure, air-gapped, locally hosted infrastructure, using fine-tuned open-weight models. This removes exposure to commercial clouds and the risk of data leaks. The second, which serves speed, is to standardise the interfaces between agents and tools. Common standards such as MCP should be adopted across the Army, Navy and Air Force so that cyber tools, scanners and defensive systems work together across all three services.

The third action serves trust: Sandboxed verification is what should be put in place (Sections 2 and 3). We should see that all AI-generated code, patches and countermeasures are tested in isolated settings before going live and that they have a high operation impact which requires human approval. This in turn protects our systems from accidental disruption. Also, an issue of trust is to adopt zero trust across strategic networks (Section 4). The DCN, Project ASCON and the NFS should move from use of static perimeter firewalls to continuous cryptographic authentication and micro-segmentation, which limits an intruder’s lateral movement.

The fifth action is to harden tactical links and adopt quantum-safe cryptography (Section 4). Multi-sensor cross-checking should be used to catch phantom radar tracks, and tactical data links and software-defined radios should begin moving to post-quantum encryption. The sixth serves control: civil–military deep-tech partnerships should be speeded up (Section 1). Dedicated military cyber technology units with flexible career paths should be set up, working with domestic start-ups and

universities through Innovations for Defence Excellence (iDEX) and the Defence AI Project Agency (DAIPA) to field sovereign capability quickly.[13]

Conclusion

The brief began with one claim: cyber power in the digital age is becoming algorithmic power, and algorithmic power is defined by speed, control and trust. The evidence supports it. Section 1 showed every major state pursuing speed, with India limited by its dependence on imports. Section 2 showed how speed can be built on locally controlled models. Section 3 showed that such systems are only as trustworthy as their data and weights. Section 4 showed what this means for military networks and tactical data links. Manual vulnerability research, slow patch cycles and annual perimeter reviews can’t compete with attacks which play out at machine speed.

States which come out on top will do so by integrating local AI, proven agent workflows and reliable infrastructure. States that are tied to commercial clouds, unverified foreign microelectronics and open external interfaces will have key points of failure in high-intensity conflict.

For India this means to make cyber defence a central element of warfighting and out of a support role. Deterrence in the fifth domain will be in terms of command of local quantised inference, common agent tool standards and post-quantum cryptography, all set within a clear joint doctrine.

Annexure: Glossary of Technical Terms

This glossary explains the technical terms used in the brief for readers who are not specialists in the field.

Active Directory (AD): Microsoft's system for managing users, devices and permissions across an organisation. An attacker who takes control of the domain controller in Active Directory gains full administrative access to the organisation's network.

Agentic framework (multi-agent architecture): Software in which several AI programmes, called agents, each handle a specialised task on their own. One agent might scan signals, another might read code, and a third might prepare a response. Instead of waiting for a user's prompt, the agents work towards set goals.

Air gap: A physical and logical separation that keeps a secure network apart from the public internet and from external radio links. It was long thought to make a command system immune to cyberattack.

Industrial control systems / SCADA: Supervisory Control and Data Acquisition systems are the hardware and software that run physical infrastructure, such as power grids, fuel distribution, naval propulsion and ammunition production.

Large language model (LLM) and small language model (SLM): An LLM is a neural network trained on very large amounts of text and code, which allows it to read, write and reason about both. In military cyber work it acts less like a chatbot and more like a reasoning engine that reads code, decodes unfamiliar protocols and issues system commands. An SLM is a smaller version that costs less to run and suits edge devices.

Living off the land (LotL): An intrusion method in which the attacker brings in no new malware. Instead, the attacker misuses administrative tools already installed on the machine, such as PowerShell on Windows or Bash on Linux, which makes detection by antivirus software harder.

Model Context Protocol (MCP): An open standard that gives an AI model one uniform way to use external software tools, such as network scanners, debuggers or firewall controllers, through structured, machine-readable instructions.

Model harness (sandbox harness): A disposable virtual test environment, such as a microvirtual machine. Any code, exploit or defensive patch proposed by an AI is run here first, to confirm that it does what is intended and will not damage critical systems.

Retrieval-augmented generation (RAG): A method that connects an AI model to a trusted database of verified information, such as technical manuals, vulnerability records and past threat reports. The model checks this database before answering, which reduces errors and invented facts, often called hallucinations.

Telemetry: The continuous stream of diagnostic and operational data from computers, servers, routers and radios, such as login attempts, memory changes and open network ports. In cyber operations it works like a sensor sweep of friendly and hostile networks.

Zero-day exploit: An attack that uses a software or hardware flaw not yet known to the vendor. Because the vendor has had no warning, no fix exists, and security tools have no signature with which to detect it.

Declaration

I declare that this manuscript is being submitted exclusively to CENJOWS for publication consideration, is original, and has not been published or submitted elsewhere. I further certify that it contains no classified, restricted, or sensitive information and is based entirely on open-source material suitable for publication in the public domain.

ENDNOTES

[1] United States Department of Defense, Summary: Department of Defense Cyber Strategy 2023 (Washington, DC: Office of the Secretary of Defense, 2023), 4–11.

[2] Jeffrey Engstrom, Systems Confrontation and Systems Destruction Warfare: How the Chinese People's Liberation Army Seeks to Wage Modern Warfare (Santa Monica, CA: RAND Corporation, 2018), 15–32.

[3] Timothy Thomas, Russian Military Thought: Concepts and Elements (McLean, VA: MITRE Corporation, 2019), 115–144.

[4] Chuck Freilich, Israeli National Security: A New Strategy for an Era of Change (Oxford: Oxford University Press, 2018), 182–198.

[5] Cherian Samuel and Munish Sharma, eds., India's Strategic Options in a Changing Cyberspace (New Delhi: Institute for Defence Studies and Analyses / Pentagon Press, 2019), 35–62.

[6] Headquarters Integrated Defence Staff, Joint Doctrine for Cyberspace Operations (New Delhi: Ministry of Defence, Government of India, 2024).

[7] Nous Research, "Hermes 3 Technical Report," arXiv preprint arXiv:2408.11857, 2024.

[8] Qwen Team, "Qwen2.5 Technical Report," arXiv preprint arXiv:2412.15115, 2024.

[9] Battista Biggio and Fabio Roli, "Wild Patterns: Ten Years After the Rise of Adversarial Machine Learning," Pattern Recognition 84 (2018): 317–331.

[10] Ian J. Goodfellow, Jonathon Shlens and Christian Szegedy, "Explaining and Harnessing Adversarial Examples," International Conference on Learning Representations (ICLR), 2015, arXiv:1412.6572.

[11] National Institute of Standards and Technology, Module-Lattice-Based Key-Encapsulation Mechanism Standard, FIPS 203, and Module-Lattice-Based Digital Signature Standard, FIPS 204 (Gaithersburg, MD: NIST, 2024).

[12] Headquarters Integrated Defence Staff, Joint Doctrine for Cyberspace Operations (New Delhi: Ministry of Defence, Government of India, 2024).

[13] Ministry of Defence, Government of India, Innovations for Defence Excellence (iDEX), Defence Innovation Organisation, https://idex.gov.in.