Synergizing Artificial Intelligence and Net5.5G in Malaysia:

A Strategic Framework for Intelligent Connectivity

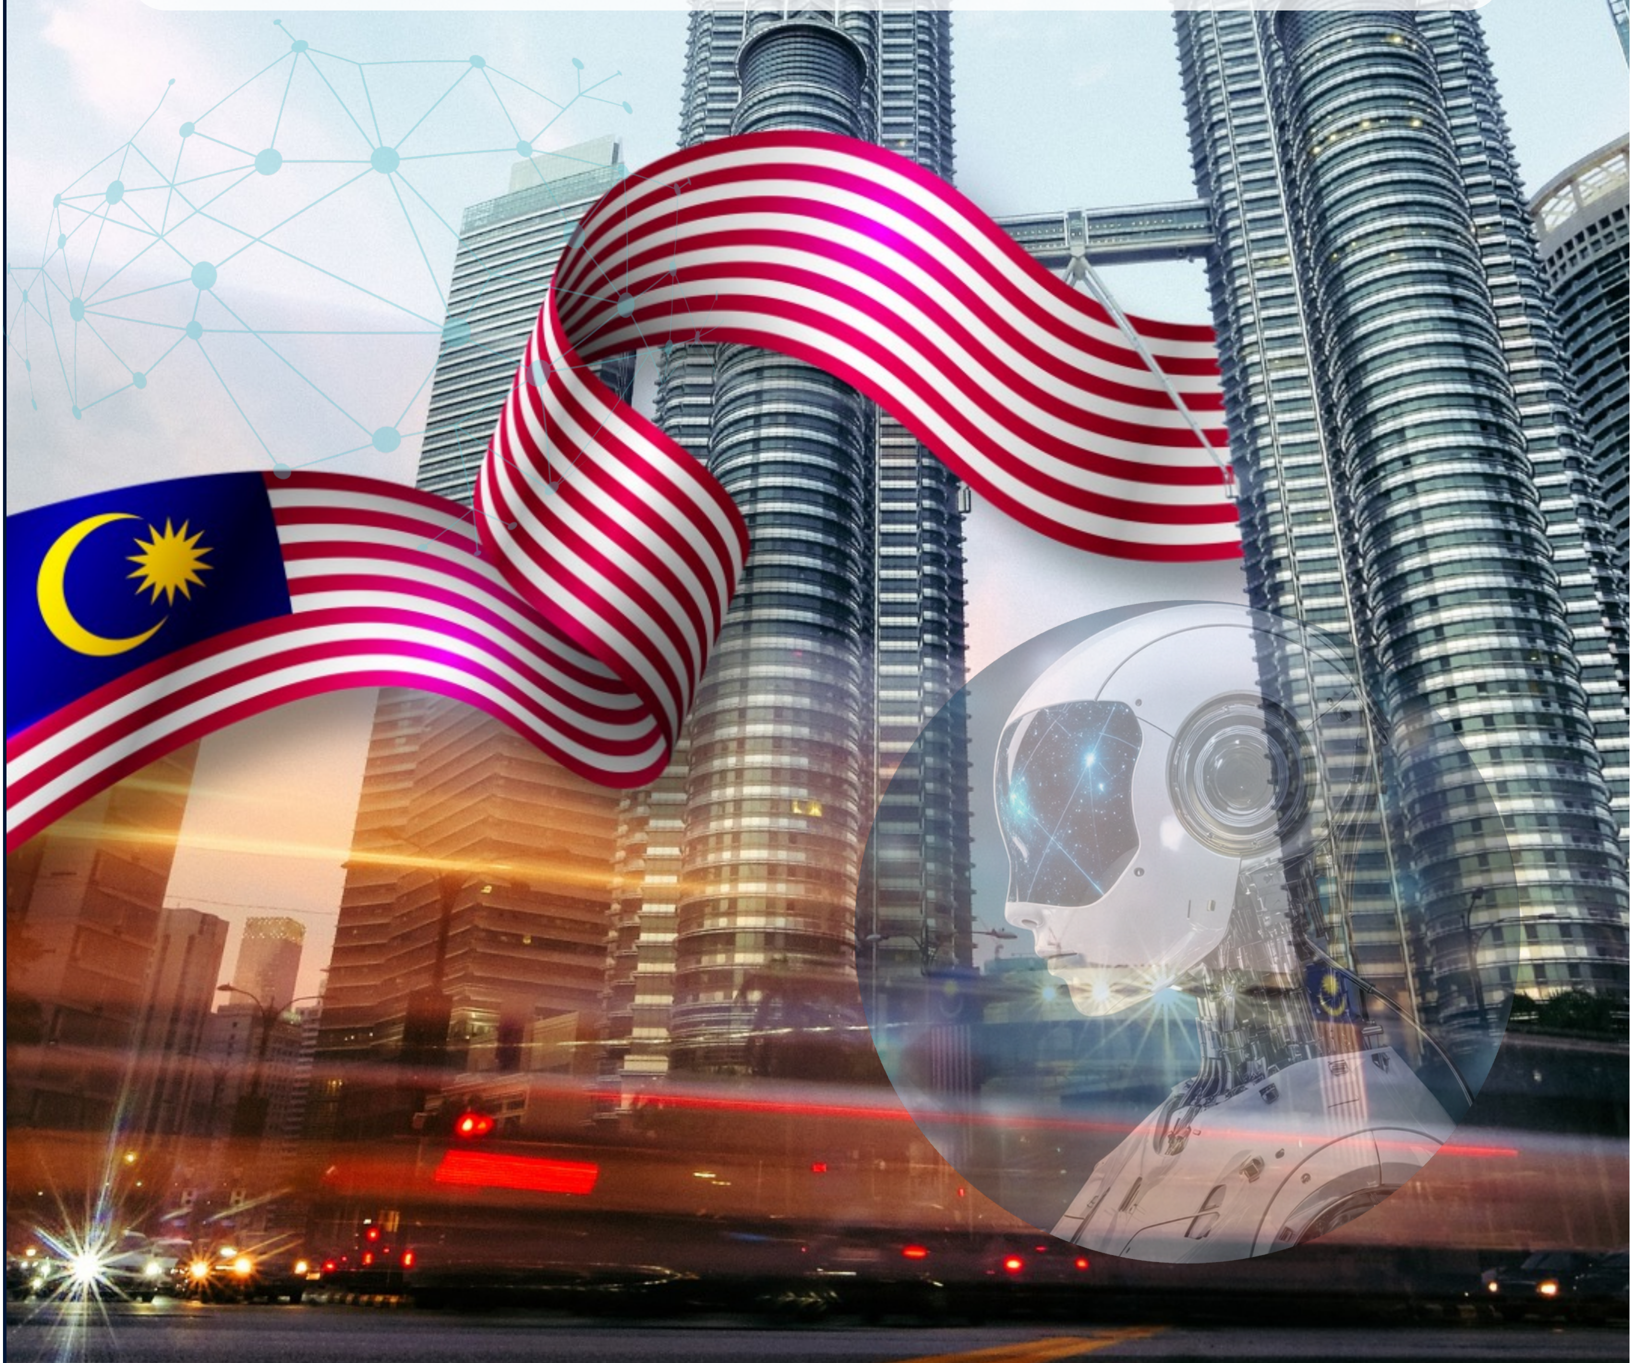

1 December 2025

Table of Contents

	Foreword		3
	Executive Summary		4-6
1.0	**The Dawn of the Intelligent Era**		
	1.1	From "Connection-Centric" to "Compute-Centric"	7
	1.2	The Global AI Infrastructure Race: Lessons for ASEAN	7-10
2.0	**The Malaysian Context: Opportunities and Bottlenecks**		
	2.1	The “Johor Paradox”: Data Center Hub vs. Value Capture	11-12
	2.2	Policy Mapping: NIMP 2030, MyDIGITAL and JENDELA	12-13
	2.3	The Infrastructure Gap: Why Legacy Networks Fail AI	13-14
3.0	**Net5.5G Architecture: The Technical Foundation**		
	3.1	Defining Net5.5G: Beyond Marketing to Engineering Reality	15-16
	3.2	The 400GE/800GE Converged Backbone	16-18
	3.3	SRv6 and the "One-Hop-to-Cloud" Paradigm	18-20
	3.4	AI-Centric Highly Secured Network: Security by Architecture, Not by Overlay	21-22
4.0	**Strategic Pillar I: The National Computing Network**		
	4.1	Unifying East and West Malaysia: The "East-Data-West-Computing" Model	23-24

4.2 Computing-Aware Routing: Treating FLOPS as a Utility 24

4.3 AI-Aware Wide Area Network (AI WAN): From Connectivity to Service Awareness 25-26

4.4 Integrated Security Fabric for the Sovereign AI Cloud 27

5.0 Strategic Pillar II: High-Efficiency Data Center Networks

5.1 The "Zero Packet Loss" Requirement for AI Training 28-29

5.2 RoCEv2, NSLB, and the 800GE Switching Fabric 29-31

5.3 Extending Lossless RDMA from Data Centers into the AI WAN 32-34

6.0 Strategic Pillar III: The Intelligent Campus and Edge

6.1 Wi-Fi 7 and the Democratization of AI Education 35-36

6.2 10GE Access for Smart Manufacturing (NIMP 2030) 37

6.3 Remote AI Training over AI WAN: Extending GPUs Beyond the Data Center 37-39

7.0 Proposed Implementation Roadmap and Policy Recommendations

7.1 Phase 1: Foundation-Building (2025-2026) 40

7.2 Phase 2: Acceleration (2027-2028) 41

7.3 Phase 3: Sovereignty (2029-2030) 41

Conclusion 42-44

References 45-49

FOREWORD

As Malaysia accelerates its journey into the AI-driven digital economy, the question is no longer whether Artificial Intelligence will transform our society and industrial competitiveness, but whether we can shape that transformation to reflect our national values, priorities, and aspirations. The work of the **Artificial Intelligence Standards Task Force** under the **Malaysian Technical Standards Forum Berhad (MTSFB)** is fundamentally about ensuring that this transformation is anchored in robust, forward-looking, and interoperable standards that safeguard innovation, trust and national sovereignty.

From a standards perspective, AI cannot be treated in isolation from the infrastructure that powers it. The technical reference models, interoperability frameworks, and performance benchmarks developed for AI systems must be tightly coupled with equally rigorous standards for the underlying network, compute, and data layers that feed these systems.

This whitepaper, **Synergizing Artificial Intelligence and Net5.5G in Malaysia: A Strategic Framework for Intelligent Connectivity**, arrives at a critical moment in our national discourse. It rightly highlights that next-generation connectivity, including Net5.5G architectures, 400GE/800GE backbones, and advanced wireless technologies, is not merely a question of speed, but a precondition for reliable, low-latency, and deterministic performance required for AI at scale.

For the Task Force, the themes articulated in this document resonate strongly with ongoing efforts to define reference architectures, quality-of-service parameters, and compliance guidelines that are fit-for-purpose for AI workloads in sectors such as healthcare, manufacturing, finance, and smart cities. Without such alignment between network capabilities and AI system requirements, Malaysia risks underutilizing both its data center investments and its emerging ecosystem of AI developers and solution providers.

This document should therefore be read not only as a connectivity roadmap, but as an invitation to co-develop a holistic standards ecosystem, one where AI models, platforms, and services are evaluated in tandem with the networks that support them. By doing so, Malaysia can move beyond piecemeal deployments toward a more coherent national AI infrastructure that is secure, scalable, and globally interoperable.

The **Artificial Intelligence Standards Task Force** looks forward to continued collaboration with all stakeholders to translate these strategic ideas into concrete, implementable standards and best practices. The choices we make today, at the intersection of AI and intelligent connectivity, will define not only our competitiveness, but also the safety, inclusivity, and resilience of Malaysia's digital future.

With warm regards,

Dr. Gopinath Rao Sinniah
Chair, AI Standards Task Force (AITF)
Malaysia Technical Standards Forum Berhad (MTSFB

EXECUTIVE SUMMARY

The Net5.5G Imperative for Malaysia

The global telecommunications landscape is undergoing a seismic shift, transitioning from the Information Era (defined by 4G/5G and content consumption) to the Intelligent Era (defined by AI training, inference, and machine-to-machine automation). For Malaysia, this transition coincides with an influx of foreign direct investment (FDI) in data center infrastructure, with digital investments hitting a record RM163.6 billion in approved investments in 2024.[9] The chart below illustrates the 250% surge in digital investments in 2024, with data centers comprising 76.8% of total investment.[48][49][50]

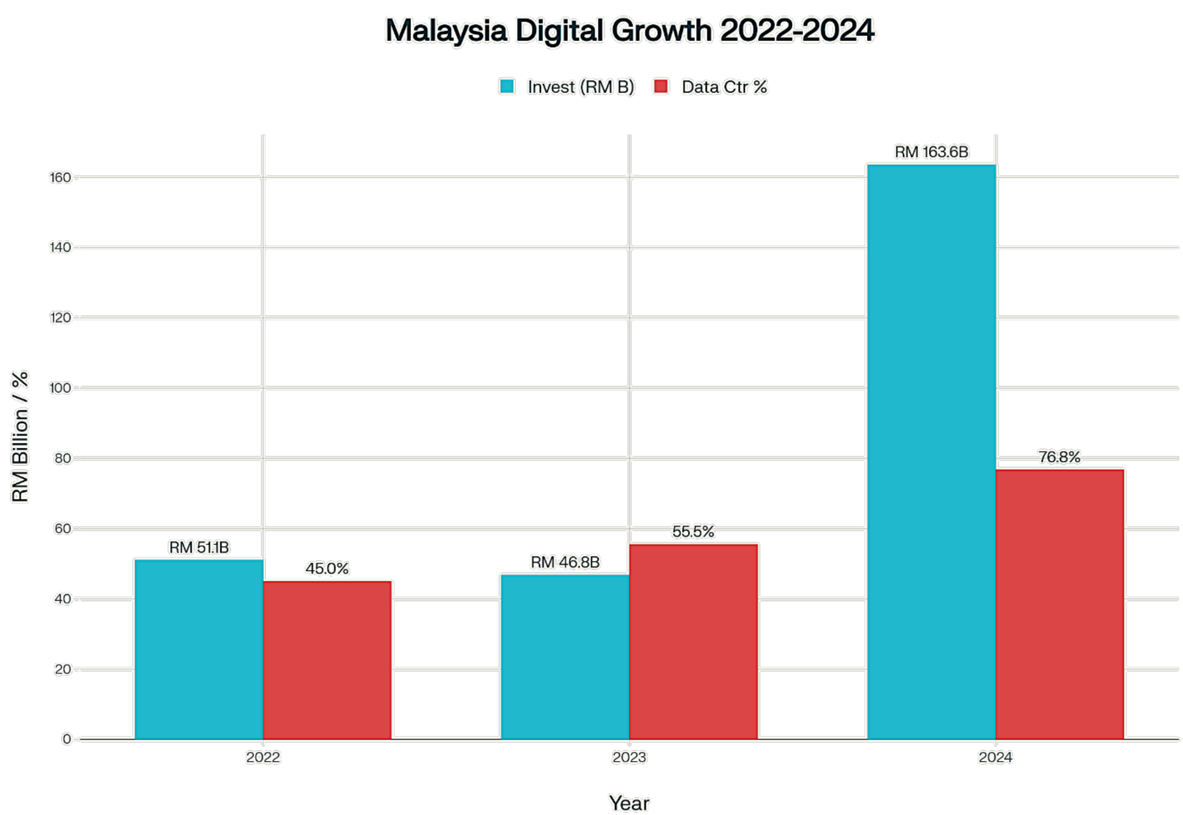

Chart E-1: Malaysia's Digital Investment Growth and % Allocated for Data Centers (2022-2024)

However, a critical gap is emerging. While Malaysia is successfully attracting “Compute” (Data Centers), our “Connect” (Network Layer) risks becoming a bottleneck. Legacy 100GE backbones and Wi-Fi 6 campus networks cannot support the massive, bursty, and latency-sensitive traffic patterns of Generative AI workloads.[7][10]

This whitepaper establishes the strategic case for **Net5.5G** as the default connectivity standard for Malaysia's AI ecosystem. Net5.5G is not a single technology but a comprehensive evolution of the IP network, characterized by **10Gbps ubiquitous access**, **400GE/800GE converged transport**, and **AI-native traffic management**.[8]

Key Strategic Findings

1. **The "Zero Packet Loss" Economic Multiplier**

   AI training clusters use the RoCEv2 (RDMA over Converged Ethernet) protocol, which is intolerant of network congestion. Analysis shows that a mere **0.1% packet loss** in a data center network can degrade AI training efficiency by **50%**. For a hyperscale investor spending RM1 billion on GPUs, a legacy network effectively destroys RM500 million in value. Adopting **400GE/800GE High-Efficiency Data Center Networks (DCN)** is therefore an economic imperative to maximize Return on Invested Capital (ROIC).[11][12][13][14]

2. **The National Computing Network (NCN)**

   To democratize AI, Malaysia must move beyond isolated data centers. We propose a **National Computing Network** architecture that uses **SRv6 (Segment Routing v6)** to logically pool computing resources across the country. This allows a researcher in a public university to seamlessly access high-performance computing (HPC) resources located in a private data center in Johor, with the network guaranteeing bandwidth and latency SLAs (Service Level Agreements).[15][16]

3. **Aligning with NIMP 2030, MyDIGITAL, and JENDELA Phase 2**

   Net5.5G is the missing link in Malaysia's existing policy frameworks:[17][18][19]

   - **NIMP 2030:** The target to transform 3,000 factories requires **10GE industrial campus networks** to support digital twins and AI visual inspection.[20]
   - **MyDIGITAL:** Cloud-First Strategy emphasizing digital inclusivity.[21]
   - **JENDELA Phase 2:** The target of "gigabit access" should be upgraded to "10-Gigabit readiness" via **Wi-Fi 7**.[22][23]

Core Recommendations

1. **Mandate Net5.5G Readiness for New Data Centers:** The Ministry of Investment, Trade and Industry (MITI) and MDEC should introduce "Green Lane" incentives for data centers that deploy **lossless 400GE/800GE networks**, certifying them as "AI-Ready Infrastructure."

2. **Upgrade the National Backbone:** Telekom Malaysia and wholesale providers must accelerate the transition from 100GE to **400GE/800GE** optical transport to handle the projected 100x growth in East-West AI data traffic.[24]

3. **Launch the "Giga-Campus" Program:** NAIO should collaborate with the Ministry of Higher Education to deploy **Wi-Fi 7 (Net5.5G)** networks in five leading research universities by 2026, ensuring the next generation of talent is not constrained by bandwidth.

1.0 THE DAWN OF THE INTELLIGENT ERA

1.1 From "Connection-Centric" to "Compute-Centric"

For the past two decades, the telecommunications industry has focused on connecting people to applications. The metric of success was coverage and download speed. Today, we are witnessing a fundamental inversion of this model. In the Intelligent Era, the network's primary role is to connect **data to computing power**, and **intelligence to the edge**.

This shift is driven by the specific characteristics of AI workloads:

- **Massive Volume:** Training a model like GPT-4 involves moving exabytes of data across server clusters.
- **Synchronous Dependency:** Unlike web browsing, where a dropped packet just causes a buffer, AI training distributes a single calculation across thousands of GPUs. If one GPU waits for data, *all* GPUs wait. The network is no longer a pipe; it is a component of the computer itself.
- **Inference Latency:** Real-time AI applications (e.g., autonomous logistics in Port Klang, or remote robotic surgery) require deterministic latency. A variance of 5ms can mean the difference between success and failure.

Net5.5G has emerged as the global standard response to these needs. Defined by the World Broadband Association (WBBA) and aligned with IEEE 802.3df standards, Net5.5G represents the evolution of the IP network to support **10Gbps** speeds everywhere (wireless and wired) and **400GE/800GE** pipes in the core.[8][25]

1.2 The Global AI Infrastructure Race: Lessons for ASEAN

Globally, nations are racing to build "AI Sovereign Clouds." China has deployed the "East-Data-West-Computing" strategy, building massive 400GE optical corridors to move data from the prosperous east coast to renewable-energyrich western provinces for processing. The United States' hyperscalers (Microsoft, Google, AWS) are aggressively deploying 800GE networks within their data centers to support trillion-parameter models. Singapore, while land-constrained, continues to dominate in high-value headquarters functions.[14][15][26][27]

The projection chart below shows global AI infrastructure spending reaching $758 billion by 2029, doubling from 2025 levels. It contextualizes Malaysia's position within the global AI infrastructure race and validates the urgency of Net5.5G adoption. [51][52][53]

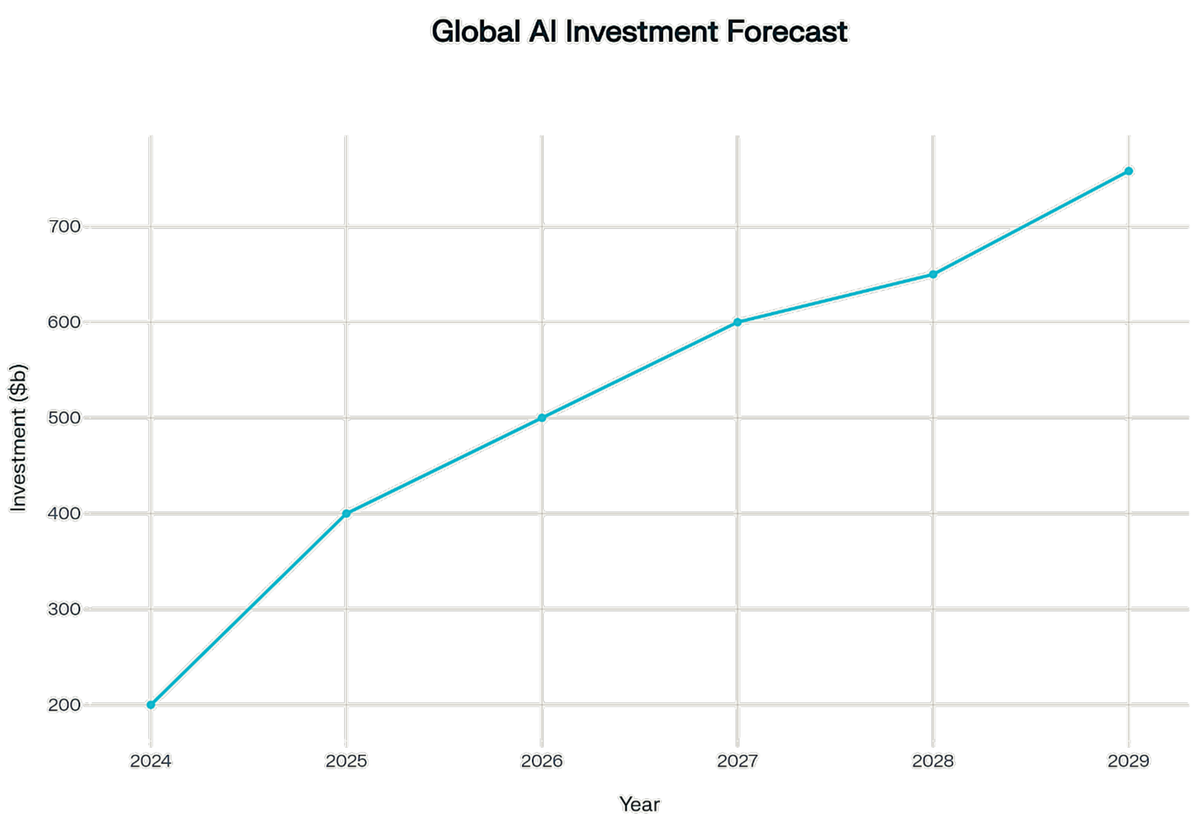

Chart 1-1: Global AI Infrastructure Investment Forecast (2024-2029)

The comparative bar chart below positions Malaysia as the second-largest ASEAN data center market with 504.8 MW capacity, trailing only Singapore's 950 MW.[49][54][55]

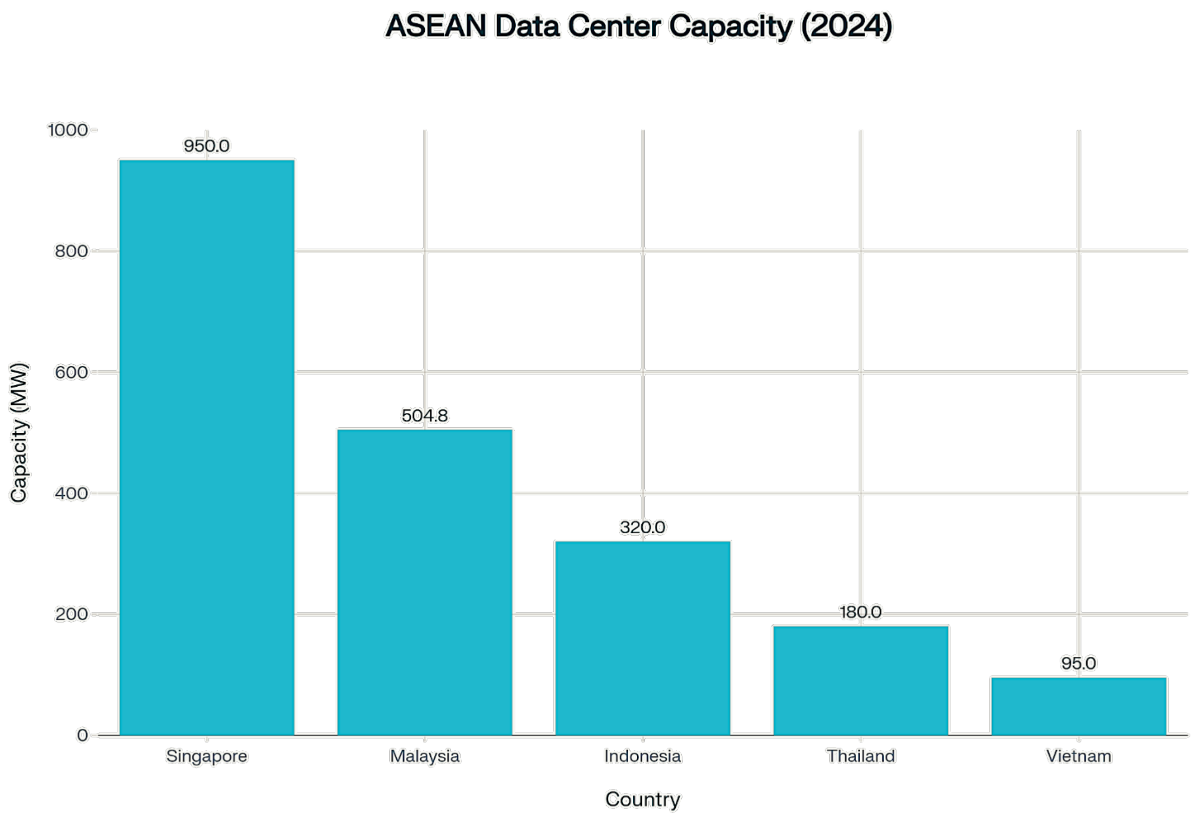

Chart 1-2: ASEAN Data Center Market Capacity Comparison (2024)

The Opportunity for Malaysia:

Malaysia possesses what Singapore lacks: land and energy (comparatively). However, to win the AI Hub status, we must match the *network quality* of Singapore and the *scale* of China.

If we build data centers with "dumb pipes" (legacy networks), we risk becoming a "bit-barn" economy, i.e. hosting the hardware but capturing none of the high-value AI development work. Hence, Net5.5G is the technological upgrade required to move up this value chain in the AI ecosystem and position Malaysia the true leading AI Hub in Southeast Asia and beyond.[27][28]

2.0 THE MALAYSIAN CONTEXT: OPPORTUNITIES AND BOTTLENECKS

2.1 The "Johor Paradox": Data Center Hub vs. Value Capture

Johor has rapidly emerged as the fastest-growing data center market in Southeast Asia. Johor's data center capacity is projected to reach 3.6 GW by 2027, representing a 360-fold increase from 2022's 10 MW baseline.[48][54][56] Major players like YTL-Nvidia, GDS, and Equinix have committed billions, with capacity projected to grow substantially in the next 5 years.[27]

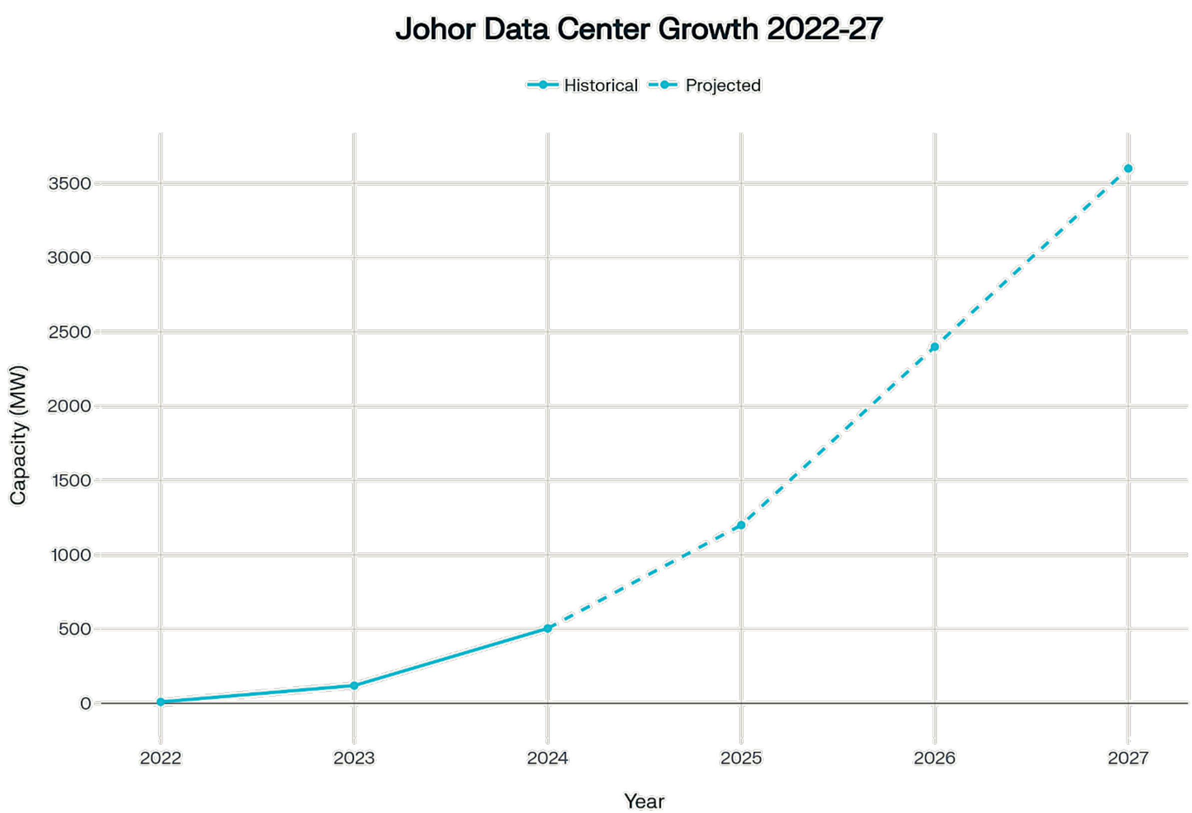

Chart 2-1: Johor Data Center Capacity Expansion (2022-2027)

However, a strategic risk exists, which we term the **"Johor Paradox."** While the *compute* capacity is being built in Johor, the *consumption* and *creation* of AI services largely happen in the Klang Valley (Kuala Lumpur/Cyberjaya) or overseas. Without a **Net5.5G-grade high-speed backbone** linking Johor to the rest of Malaysia, these GPUs will primarily serve Singaporean or international clients via submarine cables, bypassing the Malaysian domestic economy entirely.

To avoid this, Malaysia needs a **400GE/800GE National Digital Highway** that integrates Johor's compute power into a national grid, making it accessible to local startups, universities, and government agencies.[15][24]

2.2 Policy Mapping: NIMP 2030, MyDIGITAL, and JENDELA

The adoption of Net5.5G is not a new policy direction but an accelerant for existing mandates by the Malaysian government and its agencies as summarized in the Table 2-1 below:

National Policy Framework	Specific Mandate	The Net5.5G Enabler
NIMP 2030 (Mission 2)	Tech Up for a Digitally Vibrant Nation. Target: 3,000 Smart Factories[29][17][18].	**10GE Campus Networks & Wi-Fi 7:** Smart factories generate terabytes of visual data daily. Legacy Wi-Fi 6 cannot support the uplink requirements of hundreds of 4K cameras[20][30].
MyDIGITAL Blueprint	Phase 2: Digital Inclusivity. Cloud-First Strategy[31][21].	**National Computing Network (SRv6):** Ensures that rural and underserved areas can access cloud-based AI tools with the same reliability as KL-based users[15][16].
JENDELA Phase 2	Gigabit Access[22][23].	**10Gbps Connectivity & Wi-Fi 7:** Future-proofs infrastructure for immersive AI/Metaverse applications[32].
Malaysia AI Roadmap	Escalating Digital Infrastructure to Enable AI. Target: Establish AI ecosystem by 2025[33][34][35].	**National Computing Network:** Directly enables the AI Roadmap's infrastructure pillar, supporting 13,000 AI talent development targets and 900 AI companies[35].

Table 2-1: Net5.5G Enablers for Existing Malaysian Policies

NIMP 2030 Smart Factory Progress

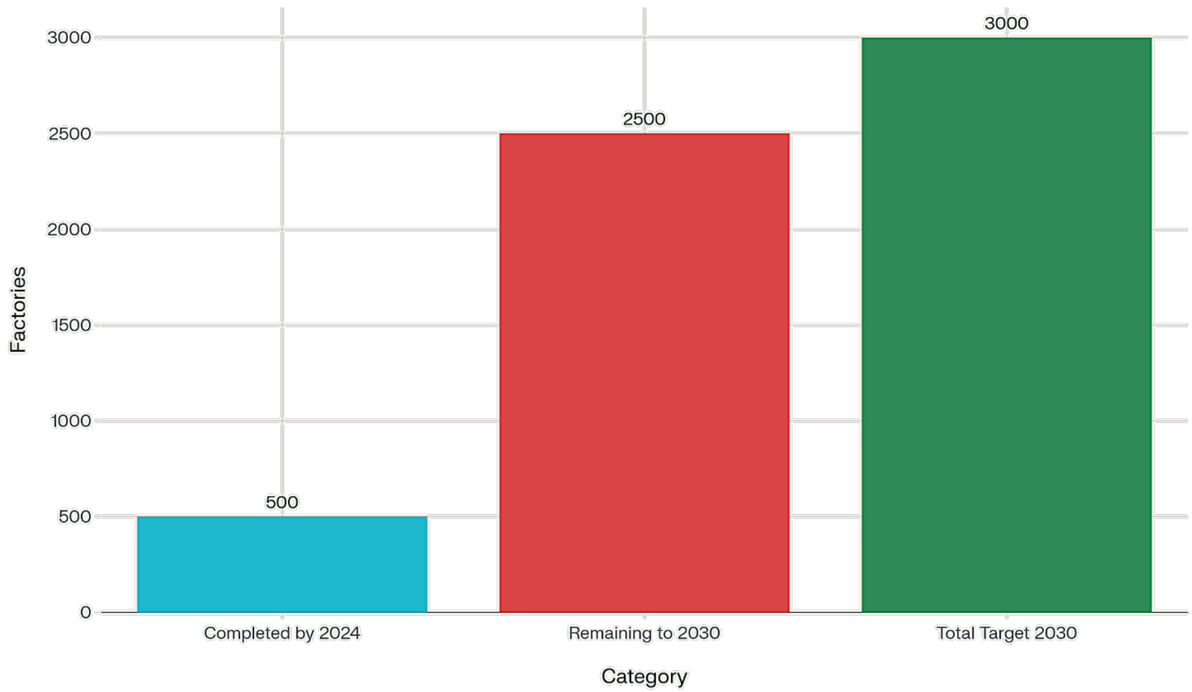

Chart 2-2: NIMP 2030 Smart Factory Initiative Progress

This progress bar chart above shows that 500 smart factories have been established by 2024, with 2,500 more needed to reach the 3,000 target by 2030. It quantifies the digital transformation challenge facing Malaysia's manufacturing sector.

2.3 The Infrastructure Gap: Why Legacy Networks Fail AI

Despite the success of JENDELA Phase 1 (4G coverage), Malaysia's backend infrastructure faces a hidden debt of sorts, which are explained herein.

2.3.1 The 100GE Ceiling: Most national backbones run on 100GE optical wavelengths. In an AI era, where a single training run can saturate a 100GE link for weeks, this is a bottleneck. We need to move to **400GE** immediately and plan for **800GE**.[14][24]

2.3.2 Campus Connectivity: Our universities produce talent, but their infrastructure is aging. Students learning to code AI models often struggle with slow Wi-Fi and capped upload speeds. A **10GE-enabled campus** is a prerequisite for producing a world-class AI workforce.[32][36]

2.3.3 Static Routing: Traditional networks use static routing (OSPF/BGP) that doesn't see congestion until packets drop. AI workloads require **Application-Aware Networking (SRv6)** that can steer traffic intelligently based on latency and load, not just hop count.[16][37]

3.0 NET5.5G ARCHITECTURE: THE TECHNICAL FOUNDATION

3.1 Defining Net5.5G: Beyond Marketing to Engineering Reality

While 5.5G is often associated with mobile networks (5G-Advanced), **Net5.5G** refers to the parallel evolution of the fixed IP network layer. It is not merely a speed upgrade; it is an architectural paradigm shift defined by the World Broadband Association (WBBA) and IEEE standards (802.3df) designed to accommodate the Intelligent Era.[8][25]

The bar chart comparison below demonstrates the Net5.5G 10x speed improvements across all network layers, from 10 Gbps access to 800 Gbps data center interconnects.[57][58]

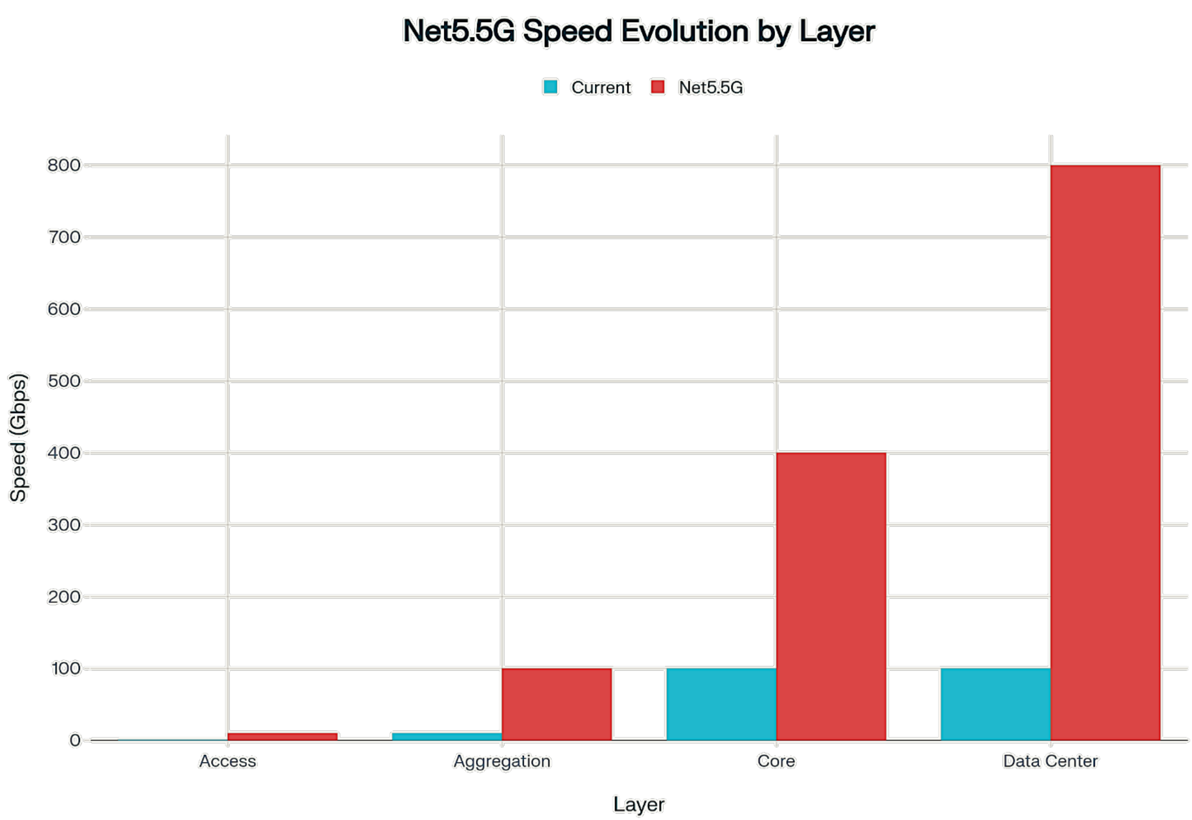

Chart 3-1: Net5.5G Speed Evolution by Layer

Unlike the Best Effort delivery model of the internet's first 30 years, Net5.5G introduces **Deterministic Quality**. For Malaysia to host mission-critical AI workloads, our national network must evolve from a *dumb pipe* to an *intelligent fabric* characterized by three non-negotiable technical pillars as defined by ITU-T standards:[38][39]

3.1.1 Ubiquitous 10Gbps Access: Moving beyond the Gigabit era to 10Gbps speeds at the edge (via XGS-PON and Wi-Fi 7), ensuring that the last mile does not choke the data ingestion process for AI models.[32][8]

3.1.2 400GE/800GE Converged Transport: Scaling the core network capacity by 4x to 8x compared to current 100GE standards to handle the exabytes of data generated by Large Language Models (LLMs).[7][14]

3.1.3 SRv6 (Segment Routing over IPv6): Replacing complex, legacy MPLS protocols with a simplified, programmable routing mechanism that allows applications (like AI training jobs) to dictate their own network path requirements.[16]

3.2 The 400GE/800GE Converged Backbone

The current state of Malaysia's fiber backbone is largely built on 100GE optical wavelengths. While sufficient for 4K video streaming, this infrastructure faces imminent saturation from AI workloads. A single training run of a trillion-parameter model can generate continuous east-west traffic flows that sustain peak bandwidth for weeks, unlike the bursty nature of consumer internet traffic.[14]

The 400GE Immediate Requirement:

Upgrading to 400GE is the immediate tactical step for 2026-2027. This involves deploying 400G-ready optical transport (OTN) and IP routers in the core rings connecting major data center hubs (e.g., Cyberjaya-Bukit Jalil-Johor). This 4x capacity increase allows for *Elephant Flow* handling, which is accommodating the massive data transfers required for AI model replication and checkpointing without causing packet loss for other users.[24]

The 800GE Strategic Horizon:

As we look toward 2028 and beyond, the IEEE 802.3df standard for **800GE** becomes critical. 800GE is not just about double capacity; it enables a fundamental shift in network economics. By moving to 800GE interfaces, network operators can reduce the cost per bit by 55-70% and energy per bit by approximately 50-65% compared to aggregating multiple 100GE links. For a nation aiming to be a *green* AI hub, this efficiency is vital.[57][59][60]

The dual-line forecast chart below shows 800GE adoption surpassing 400GE by 2029, reaching 65% market penetration. [57][59][60]

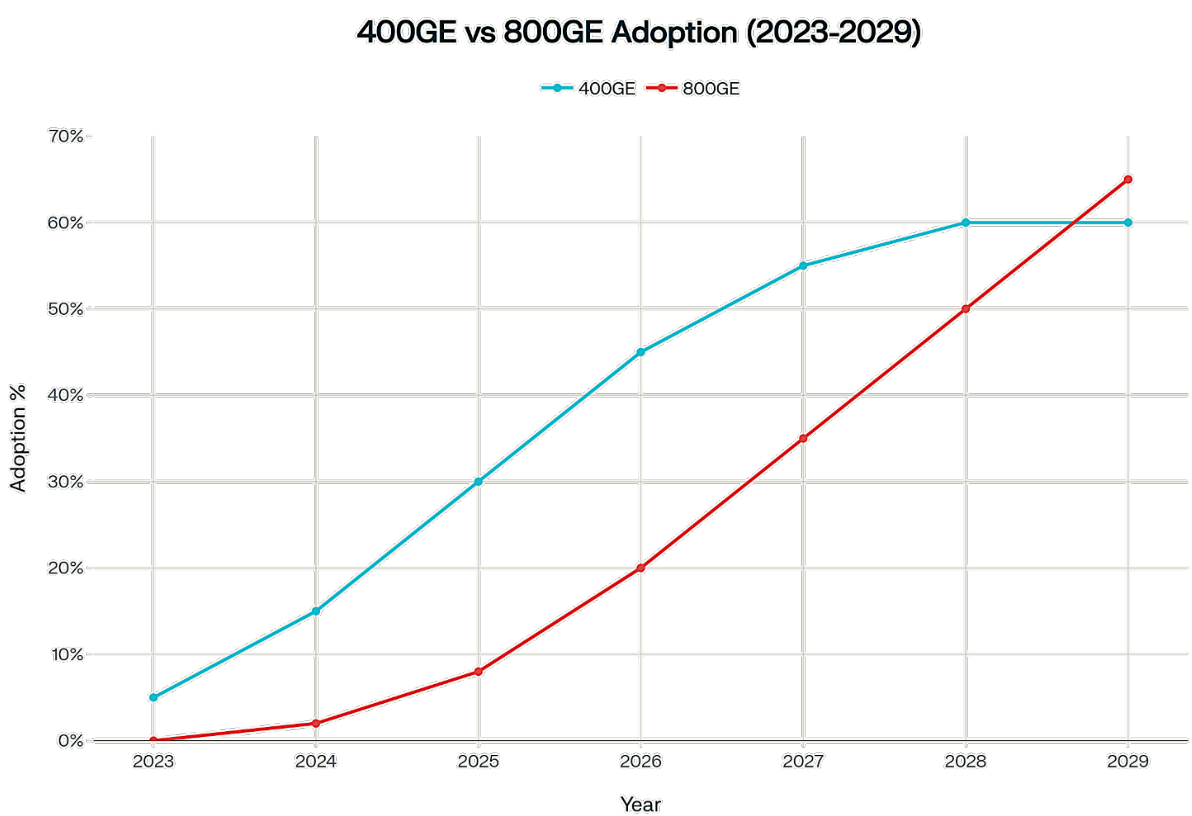

Chart 3-2: Global Data Center Adoption: 400GE vs 800GE (2023-2029)

The multi-metric comparison chart below shows that 400GE delivers 4x bandwidth with 55% cost reduction, while 800GE offers 8x bandwidth with 70% cost savings per bit.

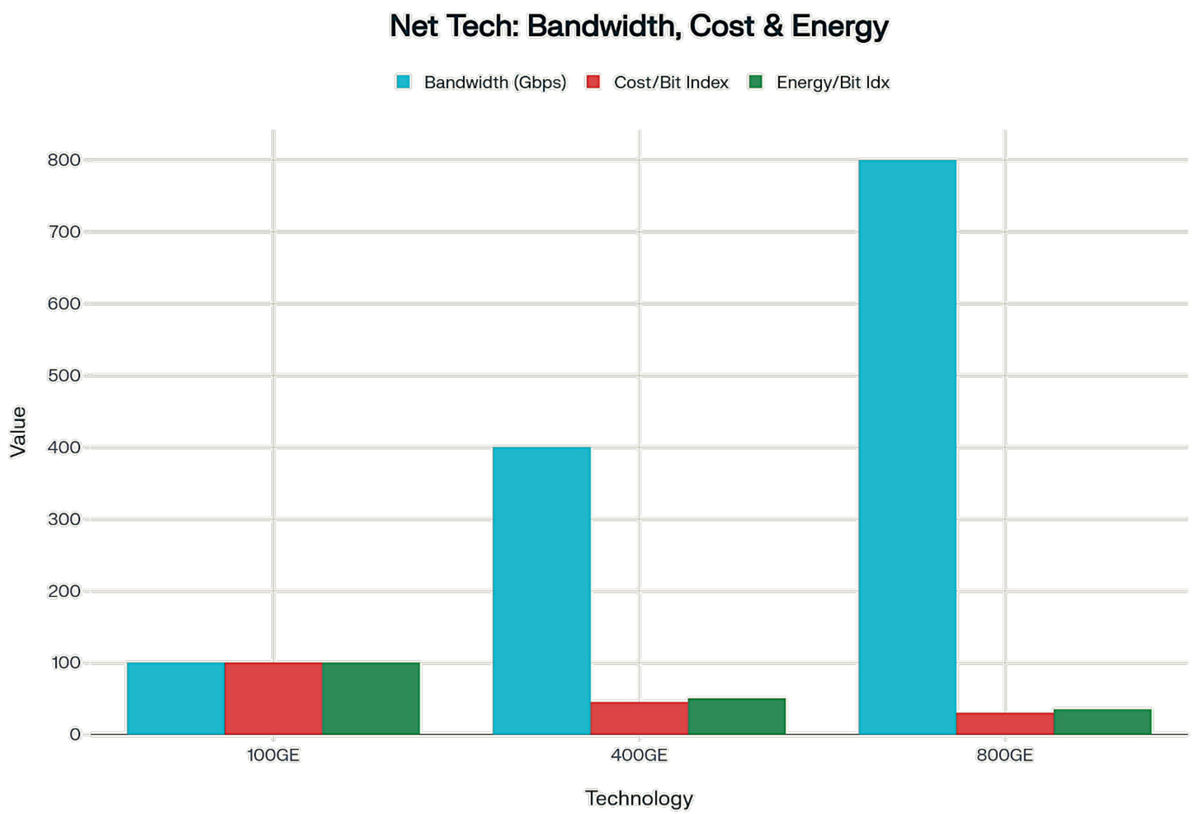

Chart 3-3: 100GE vs 400GE vs 800GE: Bandwidth, Cost, and Energy Efficiency

3.3 SRv6 and the "One-Hop-to-Cloud" Paradigm

Legacy networks use complex stacks of protocols (RSVP-TE, LDP, MPLS) that require manual configuration hop-by-hop. This "static" networking is incompatible with the dynamic nature of AI cloud bursting.

SRv6 (Segment Routing v6) simplifies this by embedding routing instructions directly into the IPv6 packet header. This enables:[16]

- **Network Slicing:** We can create a virtual, isolated *AI Slice* on the public network that guarantees low latency for government AI services, ensuring they are unaffected by public internet congestion (e.g., during a major sporting event).

- **One-Hop Connectivity:** A researcher in Universiti Sains Malaysia (USM) Penang can theoretically "see" a GPU cluster in Johor as being just one logical hop away. The complexity of the intermediate 50 routers is abstracted, simplifying the deployment of distributed AI applications.[37][16]

The comparative chart below demonstrates SRv6's 70% reduction in network complexity and 60% faster deployment time compared to legacy MPLS, while improving latency duration and scalability.

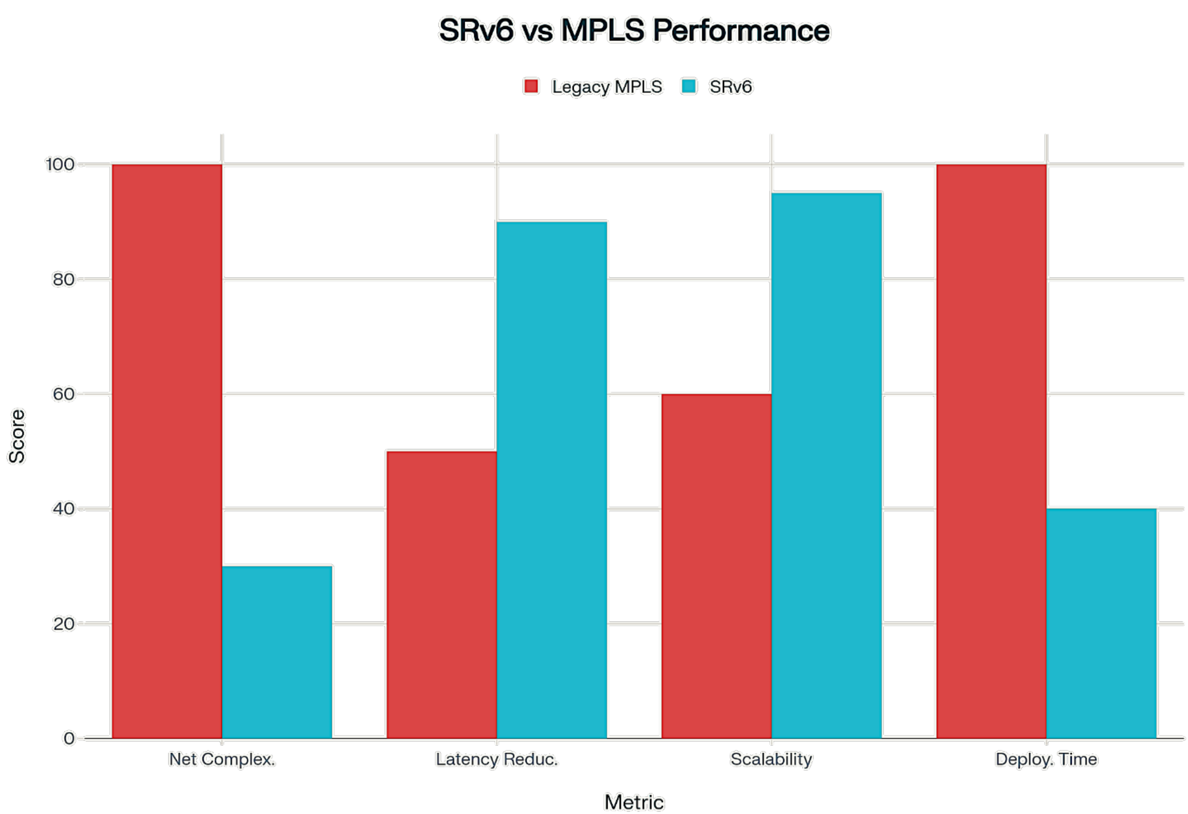

Chart 3-4: SRv6 vs Legacy MPLS Performance Comparison

3.3.1 Secure and Deterministic AI Network Slicing: Unlike conventional virtual private network models, SRv6 slicing operates natively at the Internet Protocol layer, allowing precise control over traffic paths, performance guarantees, and isolation boundaries.

In an AI-centric environment, different workloads, such as government analytics, enterprise AI platforms, academic research, and public services, have distinct security and performance requirements. SRv6 slicing enables each of these workloads to operate on its own logically isolated slice, with millisecond-level latency guarantees and megabit-level bandwidth isolation.

This capability is particularly critical for AI systems, where performance degradation or cross-tenant interference can lead to incorrect outcomes or service failures. Secure slicing prevents lateral movement during cyber incidents, ensures that encrypted traffic remains protected even under congestion, and allows regulators to enforce differentiated policies for sensitive workloads.

Through deterministic and secure slicing, Net5.5G ensures that AI services are not merely fast, but predictable, auditable, and sovereign by design.

3.4 AI-Centric Highly Secured Network: Security by Architecture, Not by Overlay

As Artificial Intelligence becomes embedded in critical national services, security can no longer be treated as a peripheral function layered on top of the network. Traditional approaches such relying on fragmented appliances, manual policy configuration, and reactive defense, are insufficient for the scale, speed, and complexity of AI-driven infrastructure. What is required is an AI-centric highly secured network, where security is intrinsic to the architecture itself.

Net5.5G enables this transformation by converging networking, computing, and security into a unified design. The architecture is defined by five core characteristics: ultra-simplified network structure, deterministic service experience, converged computing connectivity, intelligent operations, and integrated network security. This architecture is described further in the diagram below.

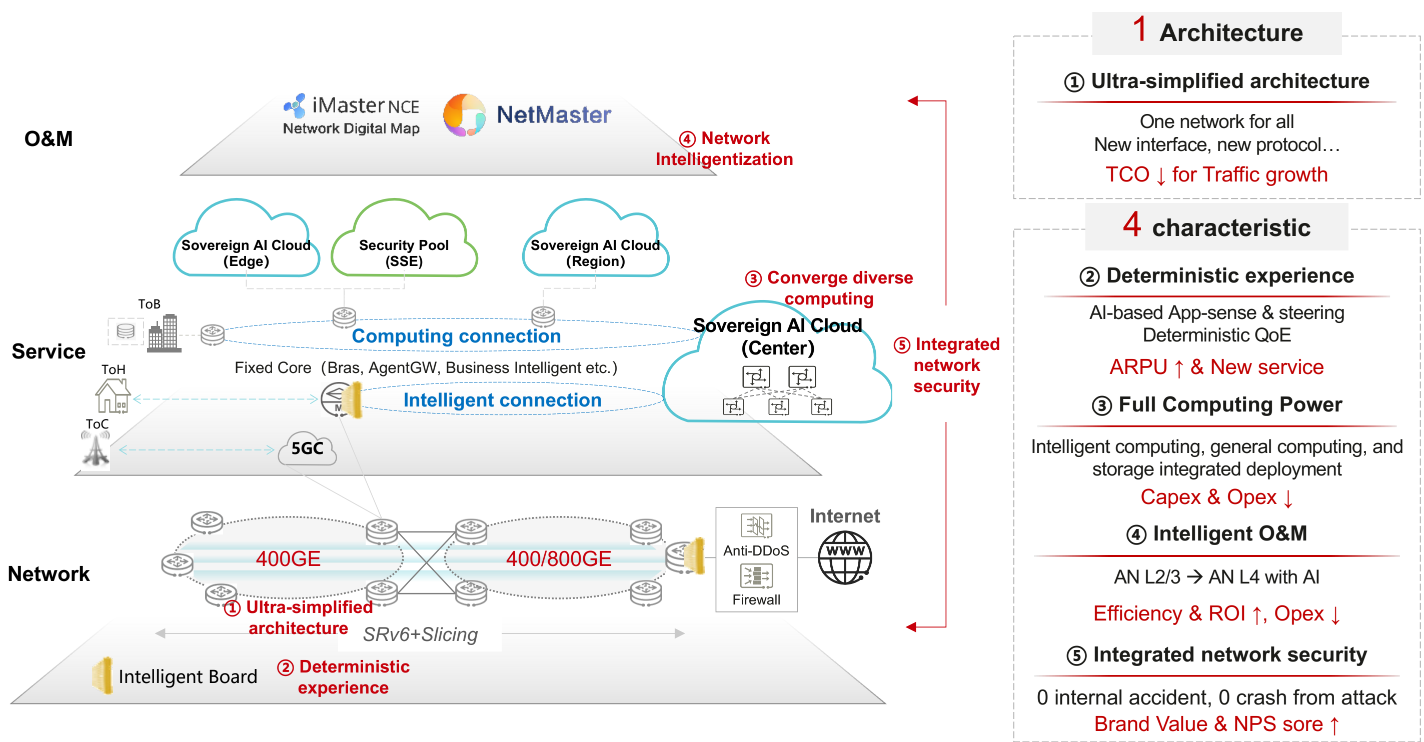

Figure 3-1: Characteristics of AI-Centric Highly Secured Net5.5G Network

By reducing protocol complexity and eliminating unnecessary architectural layers, the network minimizes failure points and operational risk while improving scalability.

Within this model, a single network supports intelligent computing, general-purpose computing, and storage services, spanning core data centers, regional hubs, and edge inference nodes. Security controls are embedded across all layers of connectivity, ensuring consistent protection regardless of where AI workloads are executed. This approach shifts security from reactive defense to proactive, systemic resilience, aligning with the requirements of sovereign AI infrastructure and critical national systems.

4.0 STRATEGIC PILLAR I: THE NATIONAL COMPUTING NETWORK

4.1 Unifying East and West Malaysia: The "East-Data-West-Computing" Model

Malaysia faces a geographical and resource imbalance similar to China's, which inspired their "East-Data-West-Computing" strategy.[15][26]

- **The Demand Centers (West):** Kuala Lumpur and Selangor generate the majority of data and demand for low-latency inference (e.g., smart traffic systems, financial AI).
- **The Resource Centers (East/South):** Areas like Johor (and potentially Sarawak/Sabah in the future) offer land and power for massive training clusters but are distant from the users.

East-West data traffic in Malaysia is projected to increase 48-fold from 2.5 exabytes in 2023 to 120 exabytes by 2028, driven by AI workloads.[2][48][49]

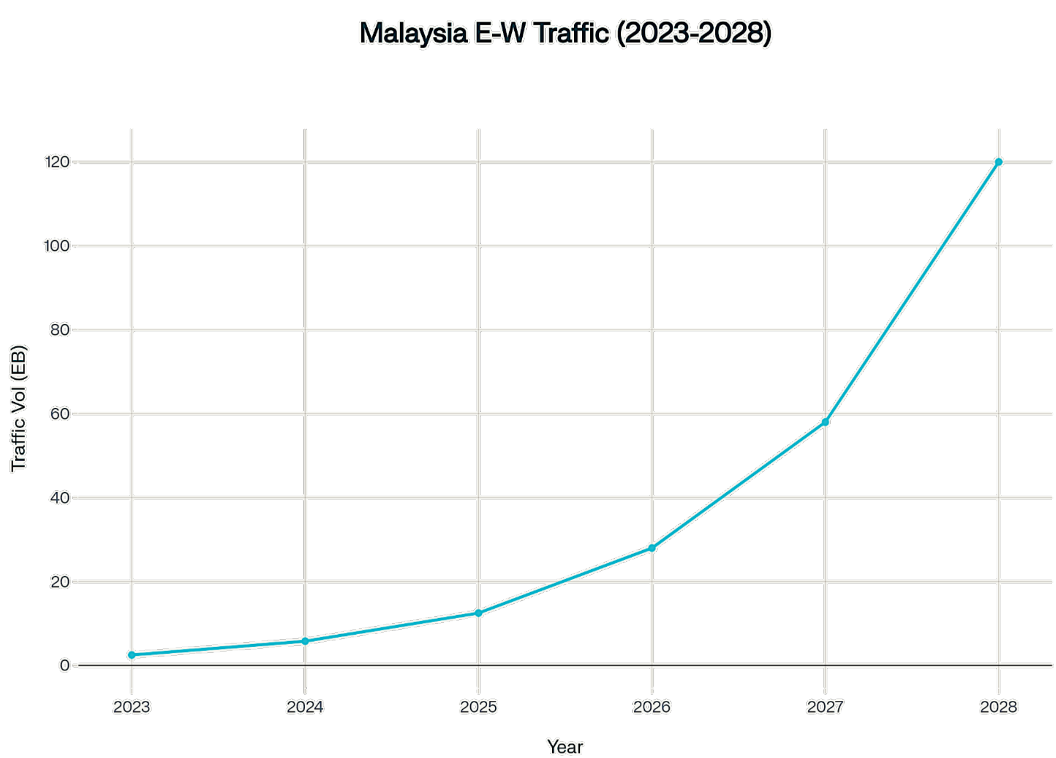

Chart 4-1: Malaysia East-West Data Traffic Projection (2023-2028)

The Solution: A National Computing Network (NCN)

We propose a unified logical network that decouples Training from Inference:[15]

- **Hot Data (Inference):** Processed in Edge Nodes in KL/Selangor using smaller, agile GPU clusters connected via 400GE metro rings.
- **Cold Data (Training):** Massive datasets are transported via a high-capacity 800GE backbone to Mega Training Hubs in Johor or renewable-powered zones in Sarawak.

This architecture transforms computing power from a *product* you buy (a server) into a *utility* you consume (like electricity), delivered over the network. The NCN ensures that a startup in KL can train a model on Johor infrastructure without noticing the physical distance.[15]

4.2 Computing-Aware Routing: Treating FLOPS as a Utility

In a traditional network, the router knows only about IP addresses and bandwidth. It doesn't know that Server A is overloaded while Server B is idle.

Computing-Aware Routing (CAN), a key feature of Net5.5G, changes this. Network devices communicate with computing resources to advertise their real-time status (e.g., "GPU utilization is 90%").[37]

- **Scenario:** An AI application requests a computing task.
- **Action:** The network doesn't just route to the *closest* server (which might be full); it routes to the *optimal* server (which has free capacity and sufficient bandwidth).

This maximizes the utilization of national computing assets. Instead of one agency's server sitting idle while another agency’s is crashed, the **National Computing Network** load-balances AI workloads across the entire country's infrastructure, dramatically increasing national ROI on hardware investments.[37]

4.3 AI-Aware Wide Area Network (AI WAN): From Connectivity to Service Awareness

The National Computing Network cannot be realized through bandwidth expansion alone. While 400GE and 800GE backbones address capacity, the next bottleneck emerges at the service and workload awareness layer. Artificial Intelligence workloads, particularly training and large-scale inference, behave fundamentally differently from traditional enterprise or consumer traffic. They generate persistent, high-volume “elephant flows,” operate under strict latency and loss constraints, and increasingly rely on fully encrypted transport protocols such as QUIC and Encrypted Client Hello.

Traditional Wide Area Networks were designed for best-effort delivery and are effectively blind once traffic is encrypted. In contrast, an AI-Aware Wide Area Network (AI WAN) introduces service awareness without violating encryption or data privacy. Using flow-level intelligence such as packet length distributions, packet interval patterns, and bidirectional traffic correlation, the network can accurately identify AI services and workload types even under full encryption. AI-based traffic identification models now achieve greater than 98% accuracy, enabling precise differentiation between AI training, AI inference, remote education platforms, and conventional cloud applications.

This capability fundamentally changes how national infrastructure is operated. Rather than treating all encrypted traffic as indistinguishable, AI WAN enables policy-driven, workload-aware forwarding, allowing the National Computing Network to prioritize mission-critical AI workloads, enforce service-level agreements, and dynamically allocate resources based on application intent. In effect, AI WAN elevates the network from a passive transport layer to an active participant in the AI computing stack, ensuring that computing resources distributed across Malaysia can be accessed with predictable performance, reliability, and efficiency. AI WAN transforms the National Computing Network from a high-capacity pipe into an intelligent AI service fabric.

The diagram below shows how an AI-Aware Wide Area Network restores service awareness under fully encrypted traffic by using flow-level intelligence and AI-based identification, enabling accurate classification of AI workloads without violating data privacy.

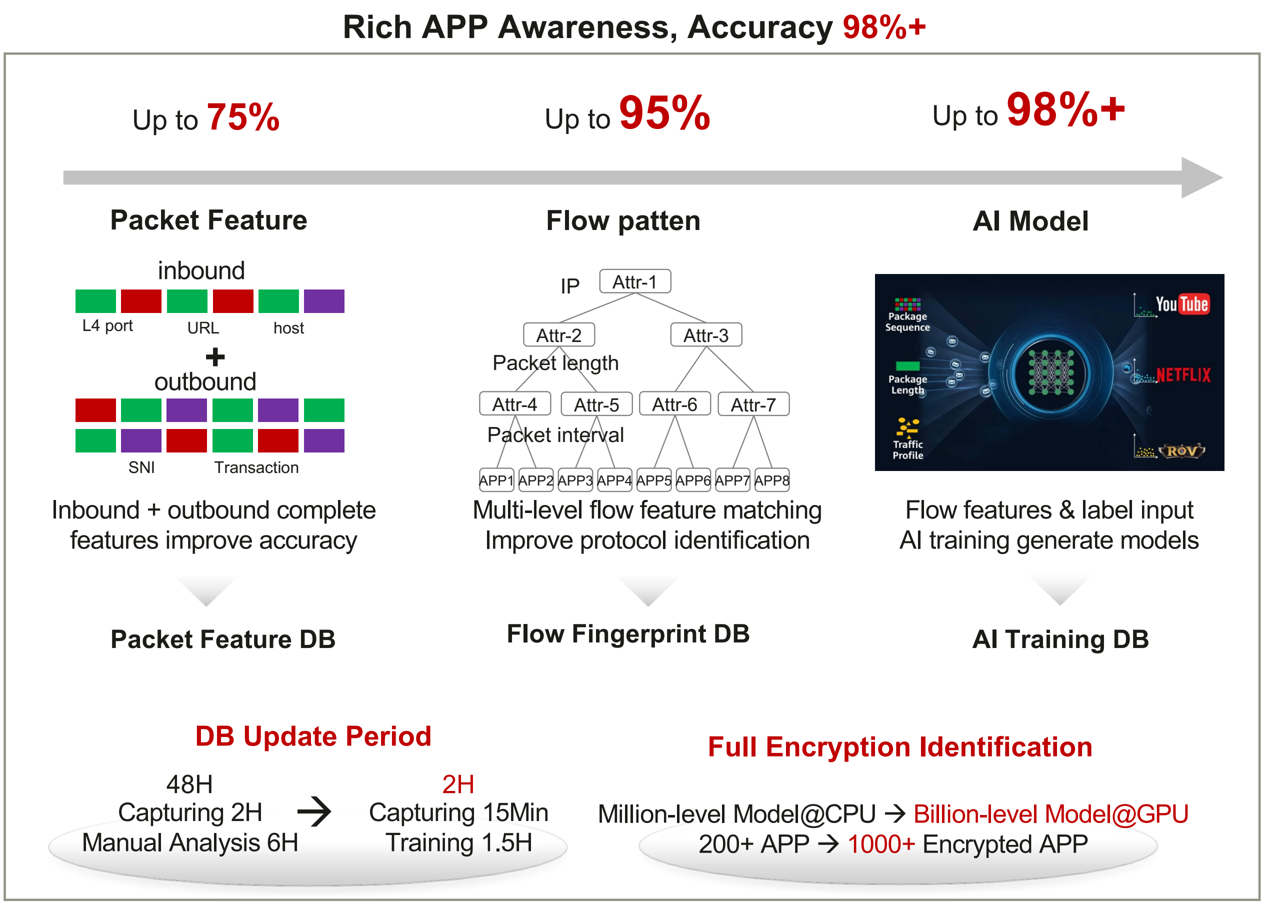

Figure 4-1: AI-Aware Wide Area Network

4.4 Integrated Security Fabric for the Sovereign AI Cloud

Malaysia's ambition to establish a Sovereign Artificial Intelligence Cloud requires more than domestic data centers and local ownership. It demands an integrated security fabric capable of protecting AI workloads across their entire lifecycle, from data ingestion and training to inference and service delivery.

Net5.5G provides this foundation by integrating firewalling, anti–Distributed Denial of Service protection, secure service edge capabilities, and encryption path control directly into the network fabric. Rather than deploying isolated security appliances at fixed points, protection is distributed and coordinated across the core, regional, and edge layers of the Sovereign AI Cloud.

This integrated approach enables millisecond-level attack detection, intelligent mitigation with accuracy exceeding 95 percent, and rapid encryption path switchover without service interruption. Network slicing based on Segment Routing over Internet Protocol version 6 ensures that sensitive government and regulated-sector workloads operate within logically isolated environments, immune to congestion and lateral movement from public traffic.

By embedding security into the connectivity layer, Malaysia can ensure that its Sovereign AI Cloud is not only high-performing but also resilient, trustworthy, and compliant with national data protection and cybersecurity mandates.

5.0 STRATEGIC PILLAR II: HIGH-EFFICIENCY DATA CENTER NETWORKS

5.1 The "Zero Packet Loss" Requirement for AI Training

The internal network of a data center (DCN) is the "backplane" of the AI supercomputer. Modern AI training utilizes the **RoCEv2** (RDMA over Converged Ethernet) protocol to allow GPUs to share memory directly. However, RoCEv2 was designed for lossless networks (like InfiniBand) and performs poorly on standard Ethernet if congestion occurs.[11]

As illustrated in the chart below, the relationship between packet loss and AI training efficiency is non-linear and punishing. A seemingly trivial **0.1% packet loss** triggers the "Go-Back-N" retransmission mechanism in RoCEv2, causing the entire GPU cluster to pause and wait. This results in a **50% drop in effective computing power**.[61][62][63]

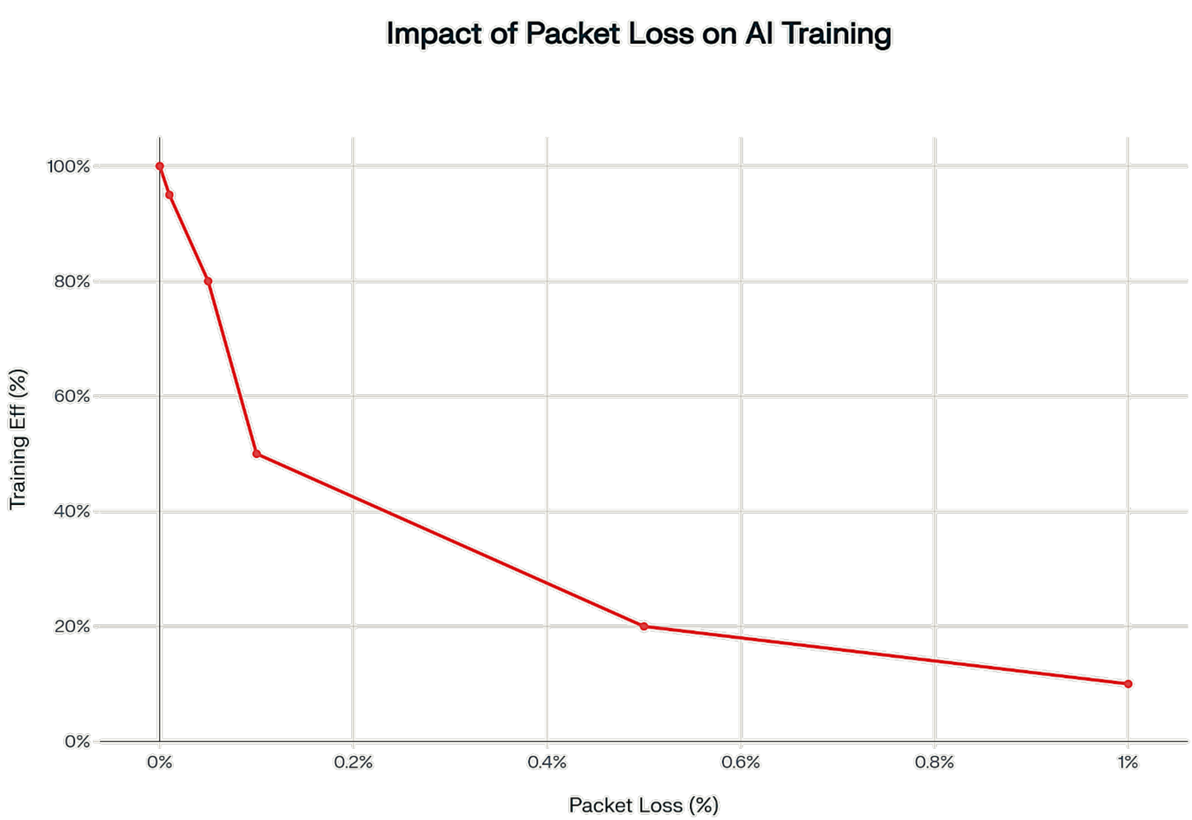

Chart 5-1: Impact of Packet Loss on AI Training Efficiency

The Economic Implication:

If a data center operator invests RM 500 million in H100 GPUs but runs them on a standard Ethernet network with 0.1% loss, they are effectively wasting RM 250 million of that capital. The GPUs are idle half the time, waiting for data.[12]

5.2 RoCEv2, NSLB, and the 800GE Switching Fabric

To solve this, Net5.5G mandates **High-Efficiency Data Center Networks** featuring:[12][13][14]

1. **Network Scale Load Balancing (NSLB):** Unlike traditional ECMP (Equal-Cost Multi-Path) routing which hashes flows randomly (often causing collisions), NSLB sprays packets across *all* available paths dynamically. This ensures that the network fabric is utilized 100% without congestion, maintaining the "Zero Packet Loss" environment required by RoCEv2.

The comparison Chart 12 below shows Net5.5G networks achieving 95-99% compliance with critical RoCEv2 requirements versus only 30-60% for legacy networks.[65][66][67] This validates the technical necessity of Net5.5G for AI workloads.

Chart 5-2: RoCEv2 Network Requirements: Legacy vs Net5.5G Compliance

2. **800GE Switching Fabric:** Inside the data center, the connections between spine and leaf switches must upgrade to 800GE. This provides the sheer firehose capacity needed to feed next-generation GPUs (like Nvidia Blackwell) which have network interfaces running at 400Gbps or higher.[14][25]

   Upgrading from 100GE to 800GE networks reduces AI training time by 79% while increasing GPU utilization from 65% to 95%, as shown in the chart below.[62][63][64]

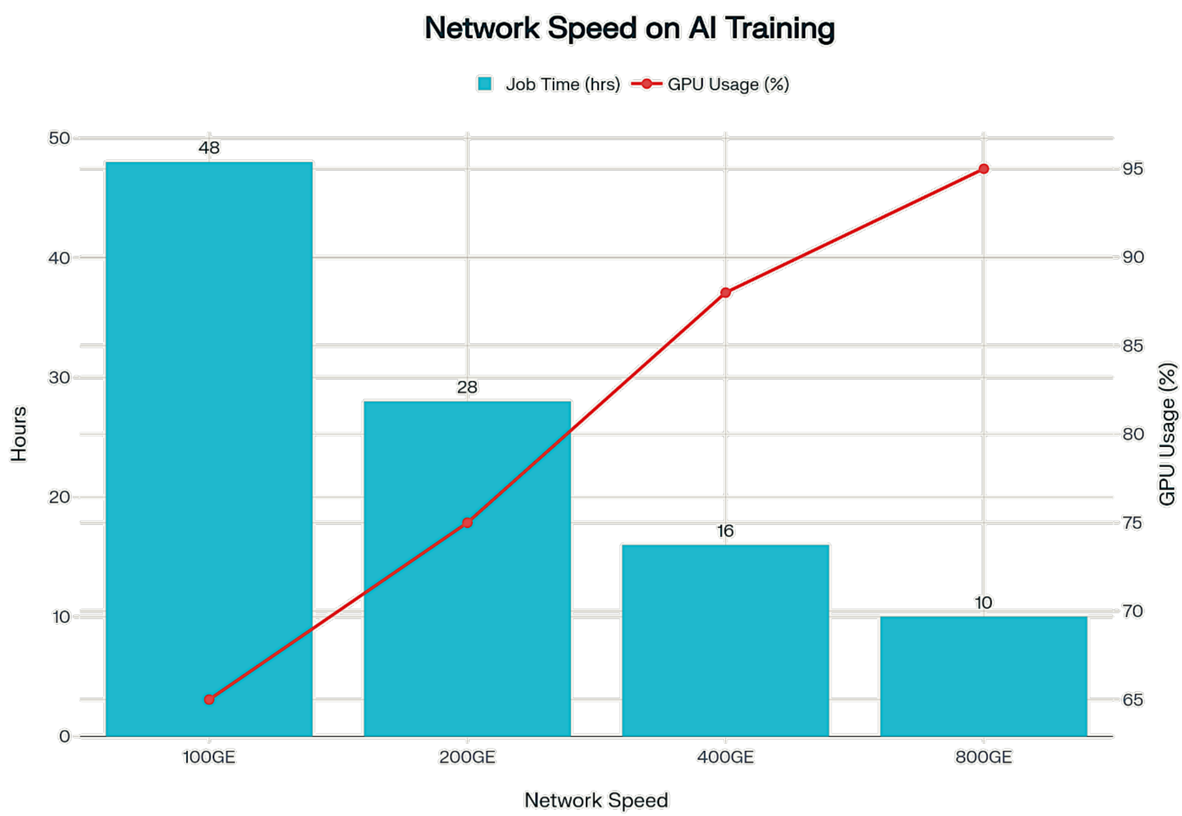

Chart 5-3: Impact of Network Speed on AI Training Performance

3. **Intelligent Lossless Ethernet:** Using algorithms like AI-ECN (Explicit Congestion Notification), the network predicts congestion *before* it happens and throttles traffic slightly to prevent packet drops, ensuring continuous, high-throughput training runs.[13][12]

5.3 Extending Lossless RDMA from Data Centers into the AI WAN

High-efficiency data center networks have already established the importance of lossless Ethernet for AI training. Protocols such as RDMA over Converged Ethernet version 2 enable GPUs to communicate directly with minimal latency and CPU overhead, but they are extremely sensitive to packet loss.

As AI workloads expand beyond single data centers into distributed and remote training models, the same lossless characteristics must be preserved across the Wide Area Network. Net5.5G enables this extension through lossless AI WAN capabilities, allowing RDMA traffic to traverse long-haul IP networks without sacrificing performance.

This is achieved through several key mechanisms. Large RDMA flows are decomposed into smaller granular flows, improving load balancing and reducing congestion hotspots. Network-level congestion control provides rapid feedback to transmitting nodes, preventing packet drops before they occur. Together, these capabilities allow the WAN to maintain lossless behavior comparable to that of an internal data center fabric.

The figure below shows how Net5.5G extends lossless RDMA beyond data centers into the Wide Area Network through flow decomposition and network-level congestion control, enabling efficient long-distance and privacy-sensitive AI training.

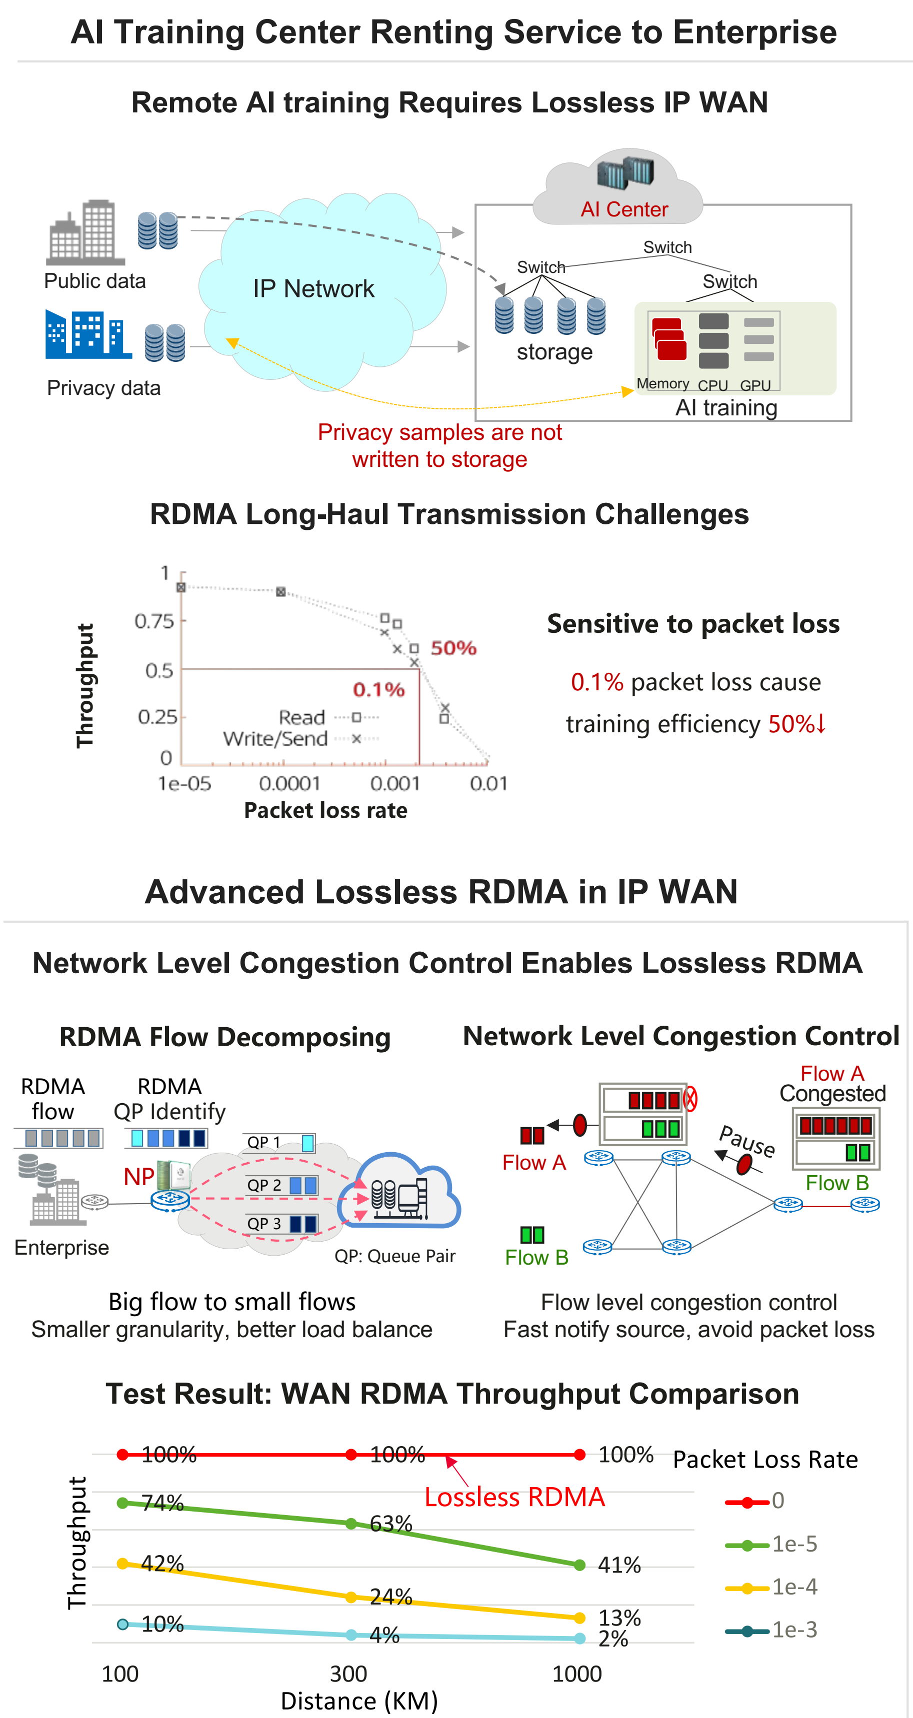

Figure 5-1: Net5.5G Extends Lossless RDMA Into Wide Area Network

The result is a national infrastructure capable of supporting remote and privacy-sensitive AI training, including regulated enterprise workloads and government applications. Organizations can securely train models using centralized GPU resources while keeping sensitive datasets local, significantly reducing data duplication risks. At a national level, this enables new service models such as AI training center rental services and shared sovereign AI infrastructure, maximizing utilization of Malaysia’s growing data center investments.

6.0 STRATEGIC PILLAR III: THE INTELLIGENT CAMPUS AND EDGE

While Chapters 4 and 5 addressed the national backbone and data center cores, the ultimate value of Net5.5G is realized at the edge, where users, devices, and applications interact. For Malaysia, this means transforming our university campuses and manufacturing floors, among others.

6.1 Wi-Fi 7 and the Democratization of AI Education

The Malaysian government has rightly identified talent as the critical bottleneck in the AI ecosystem. The AI Roadmap specifically targets 13,000 AI talent development by 2026 and the establishment of 900 AI companies. However, the physical infrastructure of learning remains stuck in the Gigabit era.[35]

The Challenge:

Modern AI education involves more than just coding. It requires:[32][36]

- **Immersive Learning:** AR/VR headsets for medical students at Universiti Malaya to visualize 3D anatomy.
- **Cloud Coding:** Students accessing remote Jupyter Notebooks running on GPUs in Cyberjaya.
- **Massive Downloads:** A typical AI dataset (e.g., ImageNet) is over 150GB. On a shared Wi-Fi 6 dormitory network, downloading this for a class assignment can take hours.

The Net5.5G Solution: Wi-Fi 7 (IEEE 802.11be)

Wi-Fi 7 is not just faster; it is deterministic wireless:[41][32]

- **320 MHz Channels:** Doubles the bandwidth of Wi-Fi 6, allowing a lecture hall of 500 students to simultaneously stream 4K content without buffering.

- **Multi-Link Operation (MLO):** Allows a device to use 5GHz and 6GHz bands simultaneously. If one band is congested, data flows instantly to the other. This reduces latency by 80%, making remote cloud desktop experiences feel local.

  Budget 2026 allocates RM635 million for public universities for upgrading infrastructure and expanding internet coverage in all public Institutions of Higher Learning, and an additional RM120 million for MCMC to improve internet access.[42]

  Recommendation: The Ministry of Higher Education (MOHE) should launch a **"Giga-Campus 2.0"** initiative, mandating Wi-Fi 7 deployment in the computer science and engineering faculties of all 20 public universities by 2026.[32] Wi-Fi 7 delivers up to 46 Gbps speeds with 320 MHz channels and 5ms latency, representing a 4.8x speed improvement over Wi-Fi 6/6E[68][69][70], shown below:

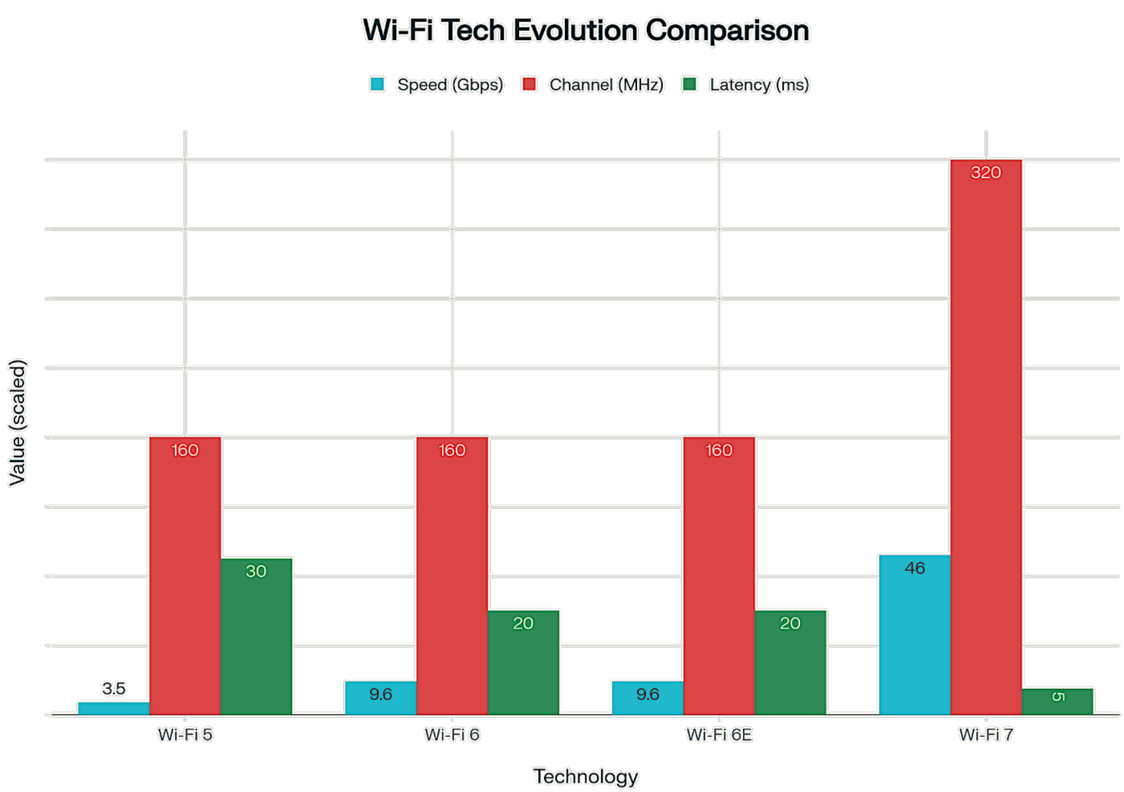

Chart 6-1: Wi-Fi Technology Evolution Comparison

6.2 10GE Access for Smart Manufacturing (NIMP 2030)

Mission 2 of the NIMP 2030 aims to "Tech Up" Malaysia by creating 3,000 smart factories. The limiting factor for these factories is often the uplink speed.[20]

- **Scenario:** A semiconductor factory in Penang uses AI optical inspection (AOI) to detect defects on wafers. This requires hundreds of 4K cameras streaming video to a local edge server 24/7.
- **The Bottleneck:** Traditional GPON networks have slow upload speeds (asymmetric). Wi-Fi 6 struggles with interference in metal-heavy factory environments.
- **The Net5.5G Fix: 10Gbps XGS-PON** and **Wi-Fi 7** provide symmetric, interference-resistant connectivity. This enables real-time Digital Twins, virtual replicas of the factory floor that sync in milliseconds, allowing managers to predict equipment failures before they happen.[30]

6.3 Remote AI Training over AI WAN: Extending GPUs Beyond the Data Center

AI talent development and innovation increasingly depend on the ability to access high-performance computing resources remotely. Modern AI education, research collaboration, and enterprise experimentation no longer occur solely within centralized data centers. Instead, students, researchers, and engineers are geographically distributed, while GPU clusters remain concentrated in specialized facilities such as hyperscale data centers in Johor or future renewable-powered zones.

Remote AI training places unprecedented demands on the Wide Area Network. A single training task can generate up to 60 terabytes of data movement per day, consisting of multiple concurrent elephant flows ranging from hundreds of megabits per second to gigabit-scale throughput. When such traffic traverses conventional metro or national networks, throughput efficiency can degrade not due to lack of bandwidth, but due to poor flow scheduling and congestion mismanagement.

AI WAN addresses this challenge through precise elephant-flow detection and intelligent flow scheduling. With built-in flow measurement engines capable of identifying elephant flows with 99 percent accuracy, the network dynamically balances traffic across available paths, ensuring that throughput closely matches allocated bandwidth. This transforms the Wide Area Network into a deterministic transport layer where “bandwidth equals throughput” for AI workloads. The diagram below shows how AI WAN enables high-throughput remote AI training by accurately identifying elephant flows and applying intelligent load balancing, ensuring that throughput consistently exceeds 90 percent of allocated bandwidth.

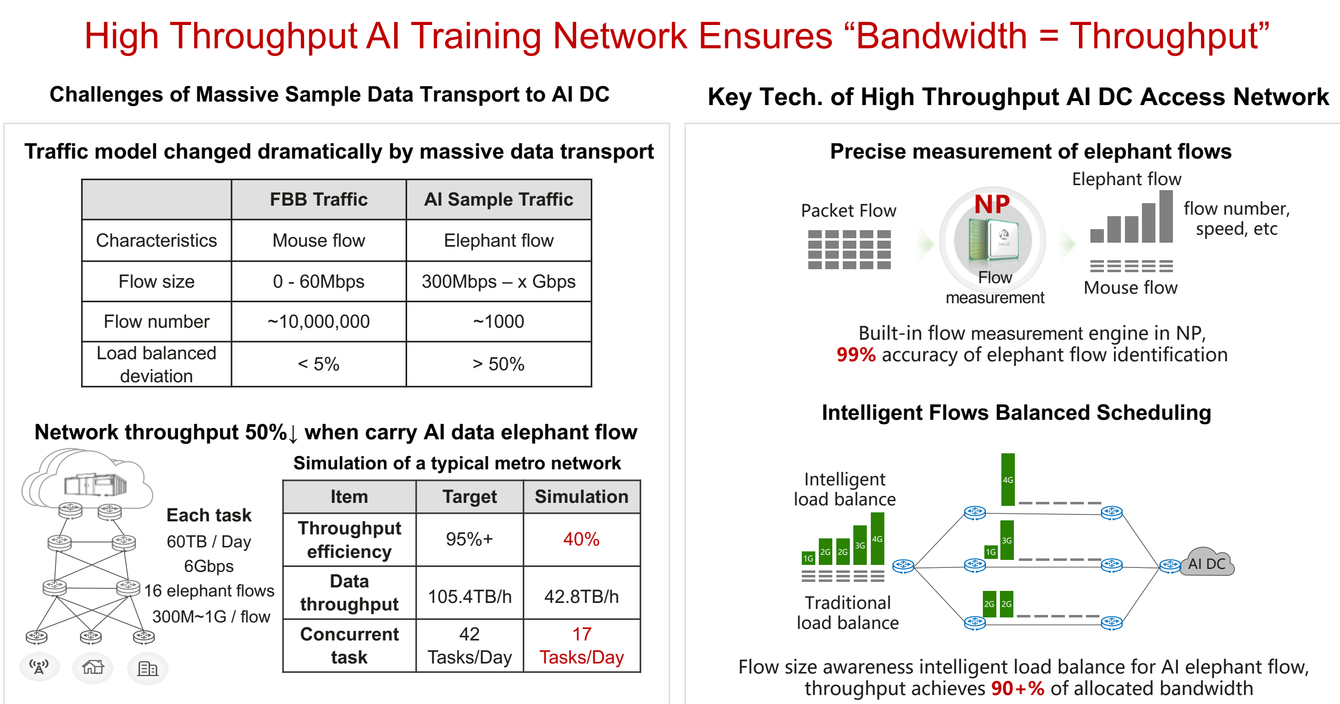

Figure 6-1: Remote AI Training over AI WAN

For Malaysia, this capability is strategically significant. It enables universities to establish remote AI laboratories, allows startups to train models using national GPU resources without physical relocation, and supports government-led AI talent programs at scale. By decoupling physical location from computing capability, AI WAN ensures that access to AI infrastructure is governed by policy and merit rather than geography, directly reinforcing the objectives of the Malaysia Artificial Intelligence Roadmap and MyDIGITAL Blueprint.

7.0 PROPOSED IMPLEMENTATION ROADMAP AND POLICY RECOMMENDATIONS

To translate this technical architecture into national reality, we propose a three-phase roadmap aligned with the 12th and 13th Malaysia Plans and the government's Digital Economy target of RM17.2 billion allocation in 2025.[42]

Figure 7-1: Digital Infrastructure Roadmap

7.1 Phase 1: Foundation (2025-2026)

- **Policy Action:** To officially include Net5.5G standards (400GE Backbone, Wi-Fi 7, SRv6) in the *JENDELA Phase 2* audit criteria. The government's Budget 2026 allocates RM1.69 billion to enhance digital infrastructure, which should prioritize Net5.5G standards.[42][43]
- **Incentive:** MIDA can introduce a **"Digital Infrastructure Investment Allowance" (DIIA)**. Data Centers that deploy "Lossless Ethernet" (RoCEv2 compliant) networks can be granted a 100% Investment Tax Allowance (ITA) for 5 years, classifying them as "AI-Ready Infrastructure".[44][45]
- **Pilot:** Launch the first "National Computing Network" node connecting Cyberjaya (Inference) to a pilot AI Park in Johor (Training). Additionally, the government's proposed Sovereign AI Cloud investment of RM2 billion through MCMC will lay the foundation for Malaysia's secure, sovereign AI infrastructure.[4]

7.2 Phase 2: Acceleration (2027-2028)

- **Expansion:** Roll out 400GE rings to East Coast (Pahang/Kelantan/Terengganu) and East Malaysia (Sabah/Sarawak) to support new renewable-energy data centers.[24][26]
- **Education:** Complete Wi-Fi 7 upgrades for all Research Universities, supporting the AI Roadmap's talent development targets.[32]
- **SME Adoption:** Provide "Smart Tech Up" matching grants for SMEs to upgrade their internal LANs to 10GE to support AI adoption.[45]

7.3 Phase 3: Sovereignty (2029-2030)

- **Sovereign AI Grid:** Full realization of the "East-Data-West-Computing" model. A researcher in Kuching can instantly provision GPU resources in Johor via a unified government portal.[15][26]
- **800GE Migration:** Begin upgrading the core backbone to 800GE to support the next generation of Zetta-scale AI models.[14][25]
- **AI Ecosystem Milestone Achievement:** Malaysia reaches the AI Roadmap target contribution of RM530 billion to the economy by 2030.[35]

CONCLUSION

The transition from the Information Era to the Intelligent Era represents far more than technological evolution, it fundamentally reshapes the competitive landscape for nations willing to embrace it. Malaysia stands at a critical inflection point where strategic decisions made today will determine whether the nation becomes a global AI powerhouse or merely a bit-barn economy, passively hosting others' infrastructure while capturing minimal value.

The core argument of this whitepaper is unequivocal. Net5.5G is not optional, it is existential. Malaysia possesses genuine competitive advantages: abundant land, significant energy resources, and a strategic geographic position in Southeast Asia, that can position it as the region's leading AI hub. However, these advantages remain latent without the underlying connectivity infrastructure to unlock their potential. A data center equipped with billion-ringgit investments in cutting-edge GPUs but constrained by legacy 100GE networks effectively sees 50% of that capital destroyed through inefficiency. Conversely, a Net5.5G-enabled facility maximizes return on invested capital and attracts hyperscale operators seeking infrastructure that can reliably support trillion-parameter models.

The economic implications are substantial. Malaysia's digital investments reached RM163.6 billion in 2024, with data centers comprising 76.8% of total investment. The projected 48-fold increase in East-West data traffic by 2028 validates the urgency of network infrastructure modernization. Without immediate action, this traffic will overwhelm legacy backbones, creating congestion that degrades AI performance and frustrates users. Conversely, a strategically deployed 400GE/800GE backbone, coupled with intelligent SRv6 routing, positions Malaysia as the network of choice for multinational AI operators and homegrown startups alike.

Beyond economics, Net5.5G directly enables achievement of Malaysia's existing national mandates. The New Industrial Master Plan (NIMP) 2030 target of 3,000 smart factories requires 10GE access networks to support real-time digital twins and AI-driven visual inspection. MyDIGITAL Blueprint's emphasis on digital inclusivity demands a National Computing Network that pools computing resources equitably across the country. The

JENDELA Phase 2 gigabit access target should be upgraded to 10-Gigabit readiness via Wi-Fi 7. The Malaysia AI Roadmap's ambitious targets of 13,000 AI talent professionals and 900 AI companies by 2030 cannot be achieved if university campuses and innovation hubs remain constrained by legacy wireless infrastructure. Net5.5G is the missing link that synthesizes these disparate policy objectives into a cohesive national strategy.

The technical foundation is proven. Global leaders, like China with its East-Data-West-Computing strategy, the United States hyperscalers with their 800GE deployments, and Singapore with its high-efficiency data center networks, demonstrate that Net5.5G infrastructure is achievable and economically justified. The International Engineering Standards (IEEE 802.3df, ITU-T recommendations) provide tested blueprints. The only barrier is coordinated execution.

The path forward is clear. The Malaysian Government, MDEC, MIDA, MCMC and NAIO, should signal to the market that Net5.5G is the preferred standard. Incentive structures, whether through 100% investment tax allowances, Green Lane fast-tracking, or direct infrastructure subsidies, must make Net5.5G adoption economically irresistible for data center operators. Universities must be mandated to upgrade to Wi-Fi 7 by 2026. The telecommunications backbone must transition from 100GE to 400GE/800GE, funded through a combination of operator capex, government grants, and development finance. The National Computing Network architecture, anchored in SRv6 and computing-aware routing, must be deployed and operationalized.

Malaysia has been given a rare opportunity: the convergence of capital availability, talent ambition, and technological timing. The AI era will generate trillions of dollars in value. The question is not whether this value will be created, but where it will be created and who will capture it. A nation with inferior infrastructure will be left behind. A nation with world-class Net5.5G infrastructure will spearhead the way.

By adopting this framework, Malaysia transitions from a passive participant in the global AI economy to an active architect of it. We will have signaled to the world's leading AI companies, researchers, and entrepreneurs that Malaysia is open, prepared, and technologically competitive. Universities will produce graduates not just trained in AI theory but hardened by infrastructure that matches the world's best. Startups will build here because they can compete globally from day one. Manufacturers will implement AI-driven optimization knowing they have the network to support it. Net5.5G is not merely about network speed, it is about national positioning, competitiveness, and sovereignty.

The race for AI leadership has begun. Malaysia's move must be decisive, coordinated, and immediate. By adopting the framework outlined in this whitepaper, Malaysia signals to the world that we are open for business, not just for today's cloud, but for tomorrow's intelligence.

REFERENCES

1. https://thetechcapital.com/malaysias-digital-economy-set-to-exceed-25-of-gdp-by-2025/
2. https://techwireasia.com/2025/11/malaysia-digital-economy-leads-sea-growth-2025/
3. https://news.microsoft.com/en-my/2024/10/18/malaysia-budget-2025-commentary/
4. https://www.digital.gov.my/en-GB/siaran/Peruntukan-Dalam-Belanjawan-Madani-2026-Akan-Merangsang-Transformasi-Digital-Malaysia
5. https://www.mydigital.gov.my/budget-2025-accelerates-malaysias-transformation-in-ai/
6. https://theedgemalaysia.com/node/737259
7. https://developingtelecoms.com/telecom-technology/wireless-networks/19233-net5-5g-white-paper-for-the-ai-era-released-by-wbba-in-paris.html
8. https://worldbroadbandassociation.com/wp-content/uploads/2024/07/WBBA_WG4_Network-Evolution.pdf
9. https://mdec.my/media-release/news-press-release/375/malaysia’s-digital-investments-hit-record--rm163.6-billion-in-2024
10. https://the-mobile-network.com/2024/05/net5-5g-application-deployment-achieves-remarkable-progress-inspiring-carriers-new-growth/
11. https://drivenets.com/blog/rocev2-as-a-scale-out-transport-protocol/
12. https://www.edge-core.com/blogs/powering-ai-at-scale-the-indispensable-role-of-data-center-bridging-with-pfc-and-ecn-for-large-gpu-deployments/
13. https://www.nida-alliance.com/uploadfile/202510/282474df285728c.pdf
14. https://www.cisco.com/c/dam/en/us/solutions/collateral/service-provider/converged-sdn-transport/white-paper-sp-analysys-mason-scaling-to-800g.pdf
15. https://www.citictel-cpc.com/en-my/blog/eastern-data-western-computing-network-construction

16. https://www.ietf.org/archive/id/draft-filsfils-srv6ops-srv6-ai-backend-02.html

17. https://www.mida.gov.my/launch-of-the-new-industrial-master-plan-2030-nimp-2030/

18. https://www.pmo.gov.my/ms/ucapanterkini/speech-text-by-prime-minister-the-launch-of-the-new-industrial-master-plan-2030/

19. https://www.midf.com.my/sites/corporate/files/2023-09/NIMP 2030-Attempting to Bring Forward Malaysia-MIDF-010923.pdf

20. https://www.mida.gov.my/mida-news/malaysia-sets-minimum-target-of-3000-smart-factories-by-2030-miti/

21. https://27.group/what-is-mydigital-initiative-digital-nasional-berhad-about/

22. https://www.malaymail.com/news/malaysia/2023/06/01/mcmc-coo-jendela-will-roll-out-its-phase-2-later-this-year/72018

23. https://www.businesstoday.com.my/2023/06/08/jendela-the-next-phase-kenanga-research/

24. https://www.fibermall.com/blog/400g-the-latest-trend-in-backbone-network.htm

25. https://www.5gtechnologyworld.com/how-ieee-802-3df-brings-800g-ethernet-to-life/

26. https://dcpulse.com/article/china-cloud-edwc-eastern-data-western-computing

27. https://w.media/highlights-from-the-malaysia-cloud-datacenter-convention-2025/

28. https://introl.com/blog/malaysia-ai-infrastructure-15-billion-investments

29. https://www.nimp2030.gov.my/index.php/pages/view/317

30. https://5g-acia.org/media/2021/05/5G-ACIA_Using_Digital_Twins_to_Integrate_5G_into_Production_Networks_single-pages.pdf

31. https://dig.watch/resource/malaysia-digital-economy-blueprint

32. https://e.huawei.com/en/articles/2024/solutions/enterprise-network/turbocharging-the-wlan-with-wifi

33. https://dig.watch/resource/malaysia-national-artificial-intelligence-roadmap-2021-2025

34. https://hkifoa.com/wp-content/uploads/2024/12/ai-roadmap-2025-malaysia.pdf

35. https://www.mida.gov.my/mida-news/malaysia-lays-roadmap-for-regional-ai-hub/

36. https://files.eric.ed.gov/fulltext/EJ1098600.pdf

37. https://e.huawei.com/en/blogs/enterprise-networking/2024/net5-5g-intelligent-cloud-network

38. https://www.ietf.org/lib/dt/documents/LIAISON/liaison-2023-01-18-itu-t-sg-13-detnet-ls-on-work-items-related-to-deterministic-networking-in-sg13-attachment-1.pdf

39. https://www.itu.int/rec/dologin_pub.asp?lang=e&id=T-REC-Y.3121-202301-I!!PDF-E&type=items

40. https://www.linkedin.com/posts/george-micheal-gerges-788a6a25_lets-deep-dive-into-each-ai-workloads-challenges-activity-7383601553095438336-icP0

41. https://www.ruijie.com/en-global/solutions/wifi7

42. https://platform.mdec.com.my/cmscdn/v1.aspx?GUID=ff549680-e362-4fee-95bd-026469bf5b23&file=Belanjawan+2025-snapshot-3_compressed.pdf

43. https://www.mydigital.gov.my/budget-2026-accelerating-malaysias-digital-transformation-for-all/

44. https://mdec.my/malaysiadigital/tax-incentive

45. https://www.hexaiot.com.my/maximize-your-digital-transformation-mida-automation-capital-allowance-sirim-smart-tech-up-grant-explained/

46. https://totaltele.com/worlds-first-net5-5g-network-construction-standard-is-released-unleashing-digital-productivity-in-the-ai-era/

47. https://www.itu.int/dms_pub/itu-t/opb/tut/T-TUT-QOS-2022-1-PDF-E.pdf

48. https://www.edgeconnex.com/news/edge-blog/johor-malaysia-the-rising-star-of-the-data-center-industry/

49. https://www.mida.gov.my/mida-news/malaysia-wraps-up-2024-as-leading-data-centre-hub-in-sea-with-us23-billion-in-investment/

50. https://mdec.my/media-release/news-press-release/375/malaysia’s-digital-investments-hit-record--rm163.6-billion-in-2024

51. https://my.idc.com/getdoc.jsp?containerId=prUS53894425

52. https://empirixpartners.com/the-trillion-dollar-horizon/

53. https://www.ropesgray.com/en/insights/alerts/2025/11/artificial-intelligence-q3-2025-global-report

54. https://technode.global/2024/12/23/malaysia-ends-2024-as-southeast-asias-top-data-center-hub-with-23b-investments/

55. https://theedgemalaysia.com/node/773980

56. https://thediplomat.com/2024/10/johor-and-malaysias-data-center-boom/

57. https://vitextech.com/ai-data-center-upgrades-for-2025-how-to-select-the-best-400g-800g-optical-transceivers-cables-and-network-solutions/

58. https://www-file.huawei.com/-/media/corp2020/pdf/giv/striding-towards-the-intelligent-world/2024/intelligent_world_data_communication_2024_en.pdf?la=en

59. https://www.datainsightsmarket.com/reports/800g-data-center-switches-620348

60. https://osi.rosenberger.com/fileadmin/content/osi/EN/News/Whitepaper/Omdia-400G-beyond-Companion-Report.pdf

61. https://developer.nvidia.com/blog/ai-fabric-resiliency-and-why-network-convergence-matters/

62. https://techstrong.ai/articles/in-ai-networks-inefficiency-has-a-high-cost/

63. https://drivenets.com/blog/latency-in-ai-networking-inevitable-limitation-to-solvable-challenge/

64. https://www.kentik.com/blog/the-network-impact-on-job-completion-time-in-ai-model-training/

65. https://techdocs.broadcom.com/us/en/vmware-cis/vsphere/vsphere/7-0/vsphere-networking-7-0/managing-network-resources/rdma/network-requirements-for-rdma-over-converged-ethernet.html

66. https://network.nvidia.com/pdf/whitepapers/roce_in_the_data_center.pdf

67. https://www.wwt.com/blog/scalable-rocev2-deployments-overcoming-network-challenges-in-ai-data-centers

68. https://www.networkcomputing.com/wireless-networking/wi-fi-7-picks-up-speed-new-devices-accelerate-enterprise-adoption

69. https://www.delloro.com/news/wi-fi-7-and-aiops-features-will-lift-campus-switch-sales-to-over-20b-in-2025/

70. https://biztechmagazine.com/article/2025/10/enterprise-leaders-guide-adopting-wi-fi-7