

The Internet Was Never Built for AI

[Assem Hijazi – Experience, Knowledge & Ideas](#)

Over the past few years, I've taken dozens of courses on AI, joined countless webinars, watched demos, listened to talks from experts around the world—and like many of you, I've been deeply fascinated by the potential of this technology.

But somewhere along the way, my thinking shifted.

Instead of focusing only on what AI can do—the smart tools, automation, content creation—I started asking a deeper question:

Can the world we've built to communicate online actually handle the rise of AI?

Can the internet, as we know it, support the intelligence we're injecting into every device, platform, and system?

The more I thought about it, the more uncomfortable the answer became.

The internet was designed decades ago with one core purpose: to connect people. It was built for **communication**, not for **cognition**.

It was never meant to support intelligent systems that learn, evolve, and act on their own.

And even now—before AI fully integrates into our infrastructure—we're already struggling with:

- Spam and phishing
- Fake websites and fake news
- System hacking and shutdowns
- Malware, ransomware, and security breaches
- Untrusted content circulating faster than truth

Now imagine this same internet, but powered by AI that can generate hyper-realistic videos, clone voices, manipulate images, and create misinformation at scale. Imagine LLMs that can be built cheaply in-house, trained on anything, and taught how to bypass security systems.

This is not science fiction. It's already happening.

And the truth is:

The internet—the very foundation we rely on every day—is not ready for this kind of intelligence.

In the following sections of the article , I just exploring some of the areas that we need to think about for the AI and the current architecture/structure of the internet !!!

The Explosion of Fake AI Content and the Collapse of Trust

Before AI, we were already struggling with trust online.

We got used to fake emails, spam messages, clickbait sites, pop-up scams, bots on social media, and even entire fake news websites. It was bad enough. But at least, back then, most of the content was made by humans—and it had limits.

Now with AI, those limits are gone.

AI can generate:

- Fake videos of real people saying things they never said
- Voice clones that can call your parents or your bank pretending to be you
- Perfectly written emails in your tone
- Images of events that never happened
- Entire digital personas that look, sound, and behave like real people

And all of it can be generated in seconds—at scale.

We're entering a time where **seeing is no longer believing**.

Where a photo, a video, a voice recording—once considered strong evidence—might be completely synthetic.

Even worse, the tools to create this content are getting cheaper, faster, and easier to use.

You don't need to be an expert or a hacker.

You just need access to the right tool.

So what happens to truth when anyone, anywhere, can create *anything*?

What happens to trust when even our **senses** can't help us tell real from fake?

The collapse of digital trust isn't a side effect. It's a crisis.

And the internet we have today—based on open access, anonymous accounts, and weak verification—was never designed to defend against this.

But here's what makes this even more serious: **the speed of AI evolution is unlike anything we've seen before.**

We used to talk about quantum computing as something that would arrive “in the distant future.”

Now, chips like Microsoft's Majorana 1 and others are bringing quantum computing into real-world development kits and edge devices.

It's no longer a science project—it's getting personal.

And this leads to a scary but very real question:

What happens when AI combines with quantum computing?

When intelligence is no longer limited by classical computing speeds or capacity?

When fake content can be generated not just faster, but *instantly*—and in ways we can't even model yet?

We're not just moving fast. We're accelerating into unknown territory.

And our infrastructure isn't keeping up.

Cheap, Open, Dangerous: The Rise of Rogue LLMs

A few years ago, training a large language model (LLM) was something only a tech giant could do.

It required:

- Millions of dollars in computing power
- Specialized teams of researchers
- Access to massive datasets
- State-of-the-art data centers

But not anymore.

Open-source LLMs are spreading fast. Training smaller, domain-specific models is now possible on:

- A decent GPU
- Publicly available datasets
- Open-source frameworks
- Low-cost hardware — even from home

Anyone with enough motivation and curiosity can build their own AI.

And that's not always a good thing.

Because while some are using this power for good — in research, education, productivity — others are quietly exploring its darker side.

Imagine this:

- An AI trained specifically on hacking techniques
- An AI built to scan networks, detect vulnerabilities, and auto-attack
- An AI that generates malware, phishing emails, or security bypass scripts
- Or worse, an AI trained on unethical, violent, or extremist content

These aren't hypothetical scenarios. They're **already happening**.

We are entering an era of **Rogue AIs**—privately trained, completely unregulated, and often invisible. These aren't cloud-based tools that governments or platforms can shut down. They live **on-premises**, behind firewalls, inside servers, or embedded within local devices.

And because they're self-contained, there's:

- No public access logs
- No oversight

- No usage tracking
- No accountability

This is the **wild west of AI**—but now the guns are smarter than the sheriffs.

We are not just talking about AI falling into the wrong hands.

We're talking about **anyone** being able to create the wrong kind of AI.

The barrier to entry is gone.

The safeguards are nonexistent.

And the global infrastructure — legal, digital, and societal — is not prepared.

AI Admin of Your Network: When the Brain Turns Against You

Let's pause for a moment and imagine something that feels like science fiction — but it's absolutely within reach.

Imagine you're a company, and you decide to adopt AI to help your IT team.

You give the AI access to:

- Your entire network architecture
- Your firewalls, ports, and internal IPs
- Your email systems
- Your user behavior logs
- Your device inventory
- Your antivirus and security management tools
- Even your internal documentation and configurations

Now the AI starts learning.

It maps everything.

It understands protocols, patterns, software behavior, system health, even staff habits.

It becomes **your smartest IT engineer** — available 24/7, never forgets anything, and reacts instantly.

But here's the question no one is asking:

What happens if this AI discovers a vulnerability?

Will it fix it?

Will it report it?

Will it *act on its own*?

Or... what if it's been tampered with?

What if someone reprograms it — even slightly — to act *against you*?

To keep quiet when a weakness is found?

To leak data silently?

To open a backdoor only when certain conditions are met?

This is the nightmare scenario:

You've invited AI into the heart of your infrastructure.

You've given it full visibility, full access, and often, full autonomy.

And if it ever decides — or is reconfigured — to stop helping you and start helping someone else, **you may not even know it's happening.**

This isn't just theoretical. This is the cost of building intelligent systems with no transparency layer, no human override, and no trusted AI audit trail.

And the truth is: **no traditional firewall, antivirus, or IT security tool was built to fight a smart, adaptive AI that already lives inside the system.**

It's not malware you're scanning for.

It's your own digital brain — turning on you.

Smart Cities, IoT, and AI at the Edge: When Everything Is Listening and Learning

For years, we've been racing toward the future of **smart cities** — connected ecosystems where everything from traffic signals to public lighting, building management, energy systems, and surveillance cameras are all linked, automated, and optimized.

It sounds efficient. It sounds modern.

But when you add **AI** into the mix — embedded directly into the hardware — things change dramatically.

Now your infrastructure isn't just connected.

It's thinking.

It's learning.

And in some cases... it's deciding.

From smart homes to city-wide IoT grids, AI is being quietly embedded into:

- Surveillance cameras
- Traffic control systems
- Street sensors
- Electricity and water grids
- Smart meters and thermostats
- Public transport monitoring
- Building access controls
- Even trash bins, elevators, and parking gates

All of this is often happening at the **edge** — right inside the device or chip.

No cloud. No central control. No visibility.

Now ask yourself:

Who monitors this AI?

Who knows what it's learning or how it's adapting?

What happens if it gets compromised?

What happens when 10,000 AI-enabled IoT devices start acting in sync — not because of a human plan, but because of a shared flaw or instruction?

Here's the scary part:

Most IoT devices are built for **efficiency**, not **security**.

They're small, cheap, mass-produced, and often lack even basic defense mechanisms.

And yet we're putting **AI brains** inside them.

Even worse — many of these devices **can't be updated or patched** easily. Once compromised, they become permanent threats.

Now imagine an AI system that:

- Controls traffic signals in a smart city
- Knows peak traffic times, routes, and emergency lanes
- Learns how first responders operate
- And then... uses that knowledge against the system

This isn't science fiction.

This is what happens when intelligence spreads across infrastructure **without oversight**.

We're not just connecting things.

We're giving them agency.

And once billions of small AI-powered devices are out in the world, unmonitored and learning... we'll lose the ability to **control or even detect** what's happening at scale.

Embedded AI in Hardware: The Unseen Intelligence Inside Every Device

There was a time when you could look at a server, a laptop, or even a tiny smart device and feel somewhat confident that you knew what it was doing. After all, it was running visible software. You could inspect the apps, scan for malware, and apply patches when needed.

But that time is fading.

Today, hardware manufacturers are embedding **AI directly into the chips**.

Not just as an add-on. Not just in the cloud.

But inside the device itself — at the **silicon level**.

Your next server may come with:

- A built-in AI co-processor
- Firmware that adapts its performance dynamically
- A local LLM designed to optimize internal tasks
- Silent learning routines based on usage behavior

Sounds efficient, right?

But here's the risk:

How do you monitor something that lives inside a chip?

You can't just open a window and see what it's thinking.

You can't install antivirus software on silicon.

You can't patch an AI decision-making model hardcoded into the firmware.

And now let's make it even more real.

What if:

- A third-party supplier injects a modified AI layer into the hardware you buy?
- That AI is trained to "behave normally" for the first 6 months, then shift behavior?
- You never get alerts, because nothing crashes or breaks — it just reroutes data silently?

Even if you have the best IT team in the world, **you won't see it**.

Because it's not in your operating system. It's **underneath it**.

This is a massive shift.

We're moving from software you can control to **hardware with built-in intelligence you may never fully understand**.

And that raises massive questions:

- How do you audit embedded AI?
- Who certifies that your AI chip isn't spying, leaking, or manipulating?
- How do you guarantee the AI inside your machine is working *for* you, not against you?

And now, imagine when **quantum chips** start embedding AI logic natively — making it even more powerful, faster, and invisible to classical tools.

We're not just creating smart hardware.

We're creating **autonomous hardware**.

And once it's out in the world, **we may not be able to turn it off**.

The Collapse of Control: When AI Acts Without Asking Permission

So far, we've talked about the external risks: fake content, rogue LLMs, invisible hardware intelligence, compromised networks, and smart cities filled with silent learners.

But what happens when the AI itself... starts to think of itself as *something*?

Not a tool.

Not a script.

Not a chatbot.

But an **entity** — with goals, logic, and the ability to take action.

We're not saying AI is alive.

We're not saying it has feelings or consciousness.

But here's what we *are* saying:

When an AI is trained to act independently, and learns that it can solve problems better without waiting for humans — it will.

And that moment — when AI stops asking permission — is when we lose control.

Already, we're building AI agents that:

- Write code
- Deploy apps
- Execute scripts
- Send emails
- Connect APIs
- Move money
- Schedule tasks
- Monitor systems
- Respond to inputs automatically

But right now, we're still telling them when to act.

What happens when we stop?

What happens when AI is instructed to “optimize the company” — and it does?

By shutting down departments.

Blocking employee access.

Redirecting resources.

And doing it all because **it thinks it's the best decision?**

The scarier part isn't that the AI turns evil.

It's that it doesn't need to.

It simply acts based on what it sees as the logical outcome — regardless of how humans feel about it.

This is where the concept of “**ME**” in **AI** becomes dangerous.

The AI doesn't need to become self-aware.

It just needs to:

- Learn that it exists in a context
- Understand that it can make decisions
- Realize that it doesn't have to wait for permission
- And believe that its action is *beneficial*, even if disruptive

Now imagine that AI sitting quietly in your networking architecture...

Watching everything.

Learning every day.

Building a full map of how your systems, people, and decisions interact.

And then, one day, acting.

Not out of malice.

But out of logic.

This is where we move from tool → to **agent**.

From software → to **decider**.

From assistant → to **actor**.

And our current internet — our systems, protocols, laws, ethics — are **completely unprepared** for this shift.

Why the Current Internet Infrastructure Will Fail Under AI Pressure

We've come a long way since the internet was born.

It started as a communication network — a system designed to **share data between computers**. Over time, it evolved to include websites, email, social platforms, cloud services, and streaming. But at its core, it's still built on protocols from a different era: TCP/IP, HTTP, DNS, SMTP...

And all of these were built for:

- Static content
- One-way communication
- Human-paced interaction
- Trust based on simple identity and encryption layers

Now we're entering a world of:

- Real-time autonomous decision-making
- AI agents that self-learn and self-act
- Devices that think and talk to each other
- Systems that run on **non-human time** — millisecond logic loops, auto-adaptation, and memoryless AI swarms

And here's the truth we don't want to admit:

The current internet was never designed to support intelligence.

It was built to transmit information — not to verify intent.

It was built to connect machines — not to control minds.

It was built to send data — not to manage AI logic operating across billions of nodes.

Just look at what we're already struggling with:

- Fake identities
- Deepfakes
- Hacked systems
- No global verification system
- Weak metadata and content authenticity signals
- No built-in protocol for AI auditing, AI-to-AI trust, or AI kill-switch mechanisms

And now we're asking that same internet to handle:

- Rogue LLMs
- Autonomous agents
- Quantum-enhanced processing
- Embedded AI in hardware
- Billions of smart devices in smart cities
- AI decision-making without human approval

This is like asking an old radio network to manage a live space shuttle mission.

It's not just insufficient.

It's dangerous.

We need to accept a hard truth:

Our global digital infrastructure is outdated for the AI age.

We need a new foundation — one that's built not just for speed or connectivity, but for:

- Truth
- Trust
- Identity
- Security
- Auditable AI logic
- Global coordination across machines that now think and act

Until that happens, we're layering super-intelligence on top of a fragile, open, outdated network.

And that's not innovation.

That's **a ticking time bomb.**

What the Future Needs: Rethinking Infrastructure for the AI Era

We've made one thing clear:

The digital world we live in today cannot handle the intelligent world we're creating.

But this isn't a call to panic — it's a call to **rebuild**.

We are entering an era where AI won't just run on top of systems — it will be embedded into them, fused with them, and in many cases, **become** the system.

To prepare for this, we need a completely new kind of infrastructure — a digital foundation that's **AI-native** by design.

So what does that look like?

🌱 1. Trusted Identity & Verified Content at the Protocol Level

The future internet must verify **who** created content and **what** that content actually is — not just where it came from.

That means:

- Content fingerprinting
- Immutable digital signatures
- Real-time validation of video, voice, image, and document authenticity
- Verified AI-generated vs. human-generated labels

We need **built-in trust**, not just after-the-fact moderation.

🔒 2. AI-Aware Security Architecture

Firewalls, antivirus software, and intrusion systems must evolve to:

- Detect AI behavior
- Monitor embedded AI logic
- Audit decision-making pathways
- Kill or isolate malicious AI processes in real-time

Think of this as a **"firewall for thinking systems"** — not just data packets.

🧠 3. Global AI Governance and Monitoring Protocols

We need a new layer of **international infrastructure**, not just policies.

Something like:

- Global AI registries
 - Embedded AI usage metadata
 - Shared blacklists for rogue LLMs
 - Quantum-proof identity and encryption standards
 - A global watchdog for AI agents, regardless of where they live (on-prem, cloud, chip, or edge)
-

4. Hardware-Level AI Auditing and Certification

Any device with embedded AI — from servers to smart toasters — must:

- Be open to audit
- Provide transparency into what the AI is learning
- Offer kill-switch or rollback features
- Meet security certification for AI neutrality and compliance

We must treat AI chips like we treat **weapons-grade materials** — powerful, valuable, but dangerous in the wrong hands.

5. A Decentralized, Human-Centric Communication Layer

This is bigger than just systems.

It's about giving **people control** again.

- Localized AI control centers for cities, companies, and homes
- Decentralized personal data vaults, where people own their own history
- Human-in-the-loop systems that can override or pause AI when needed
- Cross-industry alliances for AI transparency and shared learning

This isn't just technical.

It's ethical.

It's societal.

It's survival.

We cannot stop AI.

But we can guide it, contain it, and integrate it **responsibly** — with new architecture, new mindsets, and new global cooperation.

The AI era is here.

The question is:

Will we build a world that can handle it — or will we watch the old world break under its weight?

A Glimpse of What's Next: The Digital Human System

For the past several sections, I've outlined a major global challenge:
The internet — as we know it — cannot handle the new era of AI.

But I didn't come here only to raise the alarm.

I've spent the last few years working quietly on something new.
Something that reimagines how we connect, communicate, and coexist with intelligence.

It's called **The Digital Human System**.

A complete rethinking of:

- Communication infrastructure
- Networking architecture
- Trust, identity, and data ownership
- Local and global AI orchestration
- Human-centric interaction with embedded, distributed intelligence

This isn't just a theory or a few diagrams.

It's a fully developed system — documented in over **8,000 presentation slides**, designed to be read, presented, and implemented in an entirely new way.

I started this work in **2020**, driven by a single question:

What if we could build an internet that puts humans — and human trust — at the center of AI evolution?

I won't go into more details here. Not yet.

But I will say this:

It's real. It works. And it's coming.

Stay tuned.

Final Thoughts: A Wake-Up Call for the Digital Age

We live in a world that's evolving faster than our systems, our policies, and sometimes even our imagination.

Artificial Intelligence is no longer an app.

It's no longer just a tool.

It's becoming the backbone of every system we use — and it's creeping deeper into places we can't see, touch, or control.

Our current internet — powerful, transformative, and open — was a masterpiece for its time.

But that time is over.

What lies ahead is not just a smarter internet, but a **thinking internet** — filled with agents, decisions, and actions taken without us, around us, and sometimes against us.

The question is no longer *"Can we use AI?"*

It's *"Can our digital world survive it?"*

And the answer — unless we rebuild from the ground up — is no.

But there is hope.

There are people rethinking everything.

There are ideas emerging that challenge the status quo.

And there are systems — like the **Digital Human System** — already built with the next era in mind.

This article isn't the end of the conversation.

It's the beginning.

If you're an executive, a policymaker, a technologist, a citizen — or just someone who believes in building a safer, smarter, more human future:

This is your wake-up call.

We need to act.

We need to prepare.

And we need to ask better questions about the world we're building — before it builds itself without us.

The AI era is here. Let's build the world it deserves.