AIoT: Connecting AI to the Real World

Connected devices increase AI's relevance

Analysts loosely define the Artificial Intelligence of Things (AIoT) as the convergence of AI and IoT[1] — using AI to make IoT devices smarter and more autonomous. But that's a device-centric, "little data" definition. From a strategic, "big data" viewpoint, AIoT is the connection between machine intelligence and the physical world. The big data used for AI training and inference begins as little data captured by devices at the edge of the network that interact with things and people. These IoT sensors and human input devices are sources of truth that make AI relevant and valuable. In other words, connected devices comprise the nervous system connecting AI to our world. AIoT makes AI real and relevant.

AIoT and the Value of Device Connectivity

AI creates insatiable demands for trustworthy, real-world data to drive training and inference. Thus, rapid AI growth requires enormous amounts of accurate data about our world, and this dependency fundamentally changes the economics of device connectivity. Instead of valuing a connected device in terms of the intrinsic value of its functionality, like a thermostat measuring temperature and controlling an HVAC unit, AI expands the device value proposition to include contributions to higher-level systems such as energy management.

total_device_value = functional_value + AI_contribution_value - device_cost

AI-based ecosystems extract more information from connected sensors and controls, potentially increasing the value of each participating device significantly beyond its base functionality. However, the value of each device's AI contribution varies by use case and solution scale, so AI isn't a blank check for making expensive devices. Functional value still sets baseline customer cost expectations, so the key to developing profitable AIoT devices is adding AI capabilities and ecosystem connectivity as efficiently as possible — and cost-effective connectivity is the focus of this paper.

For more information on specific AI use cases for IoT devices, please read the third paper in the Moor Insights & Strategy (MI&S) NXP Matter series, "Matter for CE Product Manufacturers." Smart home AI examples include energy management, HVAC

[1] In its purest form AIoT involves performing the AI on the device, i.e. at the edge, with no need for external connections. Wikipedia, Artificial intelligence of things, 7 January 2024

optimization, home security, safety, health and wellness, and aging-in-place. These advanced, AI-driven applications need situational context — the ability to perceive, understand, and respond to complex situations in real time. Therefore, AI applications must connect with multiple home systems such as lights, doors, windows, cameras, security sensors, HVAC, appliances, AV, plumbing, irrigation, pools, cars, and energy sources. Universal connectivity requires standardizing two things — device networks and an application layer. Let's use examples from the consumer electronics industry to see how this works.

AIoT Redefines Device Connectivity

AI's explosive growth motivates connected device manufacturers to meet rapidly increasing demand for IoT products that provide real-world inputs to AI-based ecosystems. From a connectivity perspective, this requires (1) standard, widely available IP-based networks and (2) a software application layer that provides direct, secure, multivendor communication between devices and applications.

This new definition of device connectivity combines message delivery with message content within a vertical domain. Standardizing connectivity accelerates product development, simplifies device installation, and reduces total product cost. Standardizing message content further increases device value by enabling useful connections with multiple applications and ecosystems.

The following two sections explore the practicality of building mass-market AIoT products using a surprisingly small number of IP-bearing networks and a domain-specific application layer.

Standard Device Networks

Hundreds of unique device networks — wired, wireless, IP-based, non-IP, full-stack, LAN, PAN, LP-WAN, cellular, satellite, and many more — connect billions of devices worldwide. These networks evolved over the past 30 years, each with a compelling use case and business model. At first, rapid network innovation enabled embedded device manufacturers to differentiate product offerings based on power efficiency, range, mesh technologies, security, ease of use, device functionality, cost, and product compatibility without waiting for the complex and tedious process of industry-wide standardization. But it was worth the wait. Today, a handful of industry-standard networks subsume all relevant device networking features and functions. The rapid innovation phase of device connectivity has played out. We're now in the convergence phase, and we've reached

the point where most specialized network technologies are undifferentiated, adding only market friction and cost with no competitive advantage.

Product companies can now select from a surprisingly small set of essential device networks. Four local area networks plus cellular, satellite, and LP-WAN cover most use cases. Other specialized networks fill in the gaps.

- Standard local area device networks
  - Ethernet – wired
  - Wi-Fi – wireless
  - Thread – wireless low-power mesh
  - Bluetooth LE – point-to-point, easy pairing
- Standard wide area device networks
  - Cellular – mobility
  - Satellite – global coverage
  - LP-WAN – long-range, low bandwidth, low cost
- Specialized device networks
  - Require gateways, protocol translation, bridges
  - A few examples: Zigbee, Z-Wave, pre-IP KNX and BACnet

Seven standards-based device networks support most mainstream AIoT use cases using Internet protocols (LP-WAN a special case, covered below), and all wireless networks other than cellular and satellite use unlicensed radio spectrum. Let's take a closer look.

Local Area Device Networks

Ethernet, Wi-Fi, and Bluetooth LE are part of our daily lives and need no introduction. Thread is also ubiquitous, but it's invisible to most consumers. Mass-market devices such as smart speakers have included Thread capabilities for the past few years, and many new smartphones are also Thread-enabled, so it's already available in most homes. Ethernet and Wi-Fi cover wired and wireless high-bandwidth use cases, Thread offers mesh connectivity for power-constrained devices, and Bluetooth LE provides connectivity for point-to-point tasks such as device setup and direct control.

The combination of these three wireless networks meets the technical connectivity requirements for nearly all residential IoT products, and many commercial and industrial IoT products also use Wi-Fi. Although Thread remains unproven for commercial use cases, the Thread Group's roadmap adds enhancements that enable larger-scale deployments. Driven by strong demand for efficient IP-based mesh connectivity, Thread

is already finding its way into business environments, and this trend is likely to accelerate over the next few years.

Zigbee and other low-power mesh networks have been around for decades and have large installed bases, so why choose Thread? Thread was designed from the ground up as an IP network, capable of carrying the same messages as Wi-Fi but over a low-power mesh. Simple, standards-based Thread Border Routers connect any Thread device to any IP backbone network without protocol translation, gateways, or hubs. Also, Thread uses the same industry standard 802.15.4 radio subsystem Zigbee uses, so radio silicon is readily available from multiple suppliers. Looking forward, Thread is the best choice for secure, reliable, low-power mesh networking.

Building IoT products with Wi-Fi, Thread, and Bluetooth LE is easy and inexpensive. Silicon companies combine all three radios on a single chip with a unified software stack. NXP's "Tri-Radio" technology minimizes product integration costs by (1) reducing device complexity, (2) eliminating network stack software development, and (3) solving the nasty coexistence problems inherent in multi-radio products that share the same frequency band (2.4 GHz). NXP combines Tri-Radio technology with an application microcontroller to enable low-power single-chip products. Microprocessor (Linux) products require more configuration flexibility, so NXP also packages Tri-Radio technology as a separate chip. In either case, unified connectivity over unlicensed spectrum is the most cost-effective way to connect IoT devices.

Wide Area Device Networks

Local area device networks cover most AIoT use cases with standards-based, off-the-shelf single-chip (or two-chip) radios operating in unlicensed spectrum. However, these ubiquitous, inexpensive networks don't fit use cases that require mobility, remote locations, high density, managed QoS, extremely high security, or long range. Cellular, satellite, and LP-WAN networks fill these gaps, as described in this section. The examples listed below are for illustrative purposes and are not exhaustive.

- **Mobility** – Wide area device networks support mobility for applications such as automotive, trucking, transportation, shipping, oil and gas, mining, medical monitoring, or anything outdoors and in motion.
- **Remote locations** – Cellular and satellite services fill gaps where fixed broadband connectivity isn't available. Examples include agriculture, construction, outdoor equipment, smart cities, utilities, and environmental

monitoring. Some use cases with low bandwidth requirements can use less costly LP-WAN technologies, i.e., LoRaWAN.

- **High-density** – 5G networks can support a million connections per square kilometer. For venues such as sports stadiums, 5G can connect tens of thousands of mobile users while simultaneously connecting lights, security, safety systems, and signage.
- **Managed quality of service (QoS)** – Cellular networks, particularly 5G, offer predictable performance for applications with stringent quality of service (bandwidth and latency) requirements. In contrast, installers manage Wi-Fi QoS by overprovisioning and undersubscribing — deploying enough access points to ensure acceptable performance under worst-case scenarios. That's okay for most consumer and industrial situations but not good enough for critical infrastructure or safety-related applications.
- **High security** – SIM cards have provided security for mobile devices since the early 1990s. Today, over seven billion devices use SIM-based security. For AIoT devices, these same secure elements and the scalable infrastructure that supports smartphones can provide high levels of link security where it matters most, such as in clinical medical equipment, heavy industry, and critical infrastructure. However, all applications should use an end-to-end security layer on top of every network link — cellular, Ethernet, Wi-Fi, or Thread. Please refer to the section on Matter for details about how that standard uses such a layer to secure its connections.
- **Long-range** – LP-WAN technologies such as LoRaWAN offer low-bandwidth, low-power connections over astonishingly long distances. These specialized networks use unlicensed spectrum and are simple to develop, easy to deploy, inexpensive to operate, and reasonably secure. Applications include utility metering, smart city, parking, environmental monitoring, pet tracking, and agriculture. Although LoRaWAN can carry IP protocols, the bandwidth is very low, and message syntax is abbreviated so much that it's not directly interoperable with Wi-Fi, Thread, or other IP networks. However, gateways convert highly compressed LoRa wireless protocols to IP, so from an application standpoint, LoRa appears to be an IP network.

Specialized Device Networks

As mentioned in this section's opening paragraph, IoT connectivity architecture is exceptionally diverse. Billions of devices use hundreds of different IoT networks and variants. Product companies with significant installed bases are often reluctant to switch from these specialized connectivity schemes to Internet protocols on standard networks.

However, the transition makes good business sense for three reasons. First, IP-based products can connect with many other devices and applications, including high-value, AI-enhanced ecosystems. Metcalfe's Law[2] tells us that the financial value of a network is proportional to the square of the number of connected devices, so interoperating with larger ecosystems increases product value — perhaps not by n^2, but certainly by a superlinear amount. Second, products based on standard networks are cheaper and faster to develop and build. Third, using standard networks minimizes technical debt because long-term support costs for mainstream technologies are much lower.

Building new products on IP technologies does not mean abandoning customers and products that use specialized, pre-IP networks. Protocol-translating bridges and gateways preserve customer investments in legacy connectivity by allowing IP-based and pre-IP devices to coexist. For example, Matter specifications define bridges that translate Matter messages and commands to and from pre-IP equivalents. Matter is a consumer standard, but other application domains can define similar protocol translation schemes.

Standard Application Layers

Internet protocols enable messages to travel over IP-bearing networks without passing through hubs, gateways, or protocol translators. Consequently, an application can communicate with an IoT device connected via Wi-Fi, Thread, satellite, or cellular using identical (or similar) message payloads. An IP-based communication link is the first requirement for standardizing device connectivity. The second requirement is an application layer — a lingua franca — a device language that standardizes the content of the message payloads that travel over IP networks.

Application Layer Security and Privacy

Application layers provide secure, private, end-to-end communication over any IP network. Application layers need additional security because it's easy to "snoop" network traffic. Wi-Fi and Thread networks are secure from outside access, but message payloads are visible to all connected devices. Likewise, Ethernet messages are open to any device plugged into the network. Application layers encrypt messages between devices and applications so other nodes on the network cannot read the messages. HTTPS is an example of an application layer that we use every day. It uses transport layer security (TLS) to secure Internet traffic, such as web pages. IoT application layers such as Matter do the same thing, enabling private, secure end-to-

[2] Metcalfe's Law – https://en.wikipedia.org/wiki/Metcalfe%27s_law

end communication over any network — or combination of networks. And many use HTTPS and TLS, just like your web browser. However, IoT application layers use specialized techniques to validate device and application authenticity and manage operational credentials.

Application Layer Semantics

In addition to providing network security, application layers have a common language — a data model defining application-to-device communication within a specific application domain. For consumer applications, commands like "turn on the light" and sensor readings like "it's 82 degrees" should mean the same for every application and device.

Matter, described in the following section, defines an application layer for the smart home domain. Application layers for other domains, such as commercial buildings, healthcare, retail, agriculture, or smart cities, would have unique device definitions, attributes, and security architectures. Let's introduce Matter, and then see how to apply the same concepts to other domains.

MATTER – A STANDARD APPLICATION LAYER FOR SMART HOMES

Matter is the new smart home connectivity standard from the Connectivity Standards Alliance (CSA). It is the paradigm example of an application layer that standardizes device security, privacy, and messaging semantics over IP-bearing LAN networks such as Ethernet, Wi-Fi, and Thread. This section briefly introduces Matter and discusses its essential characteristics as an AIoT application layer.

MATTER INTRODUCTION

Matter's primary goal is to unify consumer device networks and protocols so certified products from any manufacturer can interoperate seamlessly over existing network infrastructure. Please refer to the other three MI&S white papers in NXP's Matter series for a detailed introduction to the standard:

- Matter — Making Smart Homes Smarter – An introduction to Matter
- Matter — Making Smart Homes More Secure – A deep dive into Matter's security model
- Matter for CE Product Manufacturers — A strategic analysis of Matter's industry-wide impact from a CE manufacturer's perspective

Matter, now over a year old and in its third release at this writing, has strong support from major CE brands, impressive membership growth, and over 1200 certified products spanning a steadily increasing number of device types[3]. Seamless interoperability, Matter's primary goal, requires network standardization, but that's the easy part. Matter uses Ethernet, Wi-Fi, and Thread, the most widely deployed IP-bearing networks. The hard part is standardizing the application layer so that devices and applications can interoperate over any IP network.

MATTER AS AN AIOT APPLICATION LAYER

Matter defines an application layer — a connectivity "fabric" *above* the network layer transport. The Matter fabric has two unique components:

1. Security and privacy – methods for ensuring device trust, managing secure device connections, and encrypting all messages from end to end
2. Semantics (data model) – a common language for efficient device communication over the fabric

Matter is a virtual device connectivity overlay — a fabric on top of physical IP-based networks. Matter protocols authenticate and authorize network nodes, create, and manage the secure fabric, and define the structure and semantics of messages that flow within the fabric.

Matter Security and Privacy

Our white paper "Matter — Making Smart Homes More Secure" covers trust and security in detail, but here's a simple overview of how Matter adds a new device to a home network. First, Matter checks the device's provenance to confirm that it is certified and trustworthy. Next, Matter initiates user-friendly setup procedures to install and provision the device, typically with a single click on a smartphone and an automated commissioning session over a temporary Bluetooth Low Energy (Bluetooth LE) device connection. Finally, Matter onboards the device by exchanging keys that encrypt all device-to-device communication.

Matter Semantics

Matter's data model defines a "language" for device communication. The standard language ensures that functional controls and sensor values like on, off, brightness,

[3] Matter 1.2 release blog – Connectivity Standards Alliance, 23 October 2023 https://csa-iot.org/newsroom/matter-1-2-arrives-with-nine-new-device-types-improvements-across-the-board/

temperature, and color have the same meaning on all devices. The result is a comprehensive and ever-expanding dictionary for device communication over IP networks, specifying the message structures, data types, attributes, and commands for every supported device type.

The era of ubiquitous device connectivity has arrived. With Matter, manufacturers can use any combination of standardized, widely available, IP-based networks — Ethernet, Wi-Fi, and Thread — plus Bluetooth LE for setup. Additional IP-bearing networks, such as cellular and satellite, could support new use cases outside the home and in other vertical domains. However, new networks would require new device commissioning and security procedures.

Matter defines an application layer to standardize protocols, security, and semantics for smart home applications, and the same concepts can apply to other domains. We explore this possibility in the next section.

MATTER VERTICAL MARKETS – OUTSIDE THE HOME

The Connectivity Standards Alliance (CSA) developed Matter as an application layer to standardize networks, security, and semantics for CE products in residential environments. However, demand for Matter-like application layers is increasing in other domains such as multi-family residential, small business, healthcare, and other verticals. Where does Matter fit? Where does it not fit? How will application layers evolve to address other vertical markets? This section covers these questions from a product developer's perspective.

MATTER IN NON-RESIDENTIAL APPLICATIONS

As a smart home standard, Matter does not directly address other product domains such as commercial buildings, retail, manufacturing, health care, and agriculture. But Matter products are already appropriate for some commercial applications with residential-like characteristics. For example, Matter is a good fit for lighting and HVAC controls in small to medium-sized businesses with residential device types. However, at least for now, larger commercial and industrial applications are beyond Matter's scope.

We foresee Matter's scope potentially expanding into some commercial applications. In the early days of Wi-Fi, many enterprises dismissed the technology as inappropriate for commercial use, mainly because of weak security. Wi-Fi quickly became ubiquitous in commercial environments after the Wi-Fi Alliance addressed security in 2003 with WPA. We believe Matter will eventually have a similar trajectory, possibly adding business-

friendly features in future versions. We expect the CSA to continue focusing on residential applications as Matter adoption accelerates. Still, the same market forces that compelled enterprises to adopt Wi-Fi motivate business customers to deploy Matter, or Matter variants, into non-residential verticals where it's potentially a good fit.

CONNECTIVITY CONSIDERATIONS: HOME AND SMALL BUSINESS (HSB) VS. ENTERPRISE

This paper describes significant AI-driven business opportunities for manufacturing billions of AIoT devices that operate on standard LAN networks in unlicensed spectrum and natively connect with AI-enabled application ecosystems. All hardware, network, and software components are standards-based and off-the-shelf, and minimal system customization is needed, so the economies of scale are attractive. Matter is the paradigm example, and it's a good fit for the smart home domain and other verticals with similar characteristics. The "similar characteristics" list expands as Matter (and Thread) mature. Still, at some point, fundamental limitations of LAN network scalability and Matter-like application layers preclude enterprise deployment. In these cases, enterprises turn to licensed connectivity solutions. In this section, we explore the limits of unlicensed connectivity for enterprises.

TABLE 1: CONSUMER VS ENTERPRISE DEVICE CONNECTIVITY

Connectivity Requirements	Consumer (HSB)	Enterprise
Spectrum	Unlicensed (ISM bands)	Unlicensed and licensed
Infrastructure cost	Very low investment (off-the-shelf)	High investment
QOS, performance	Unmanaged – overprovisioned mesh	Managed – end-to-end QoS
User authentication	Primary user; shared network PWs	Individual users and groups
Multi-site	Not required	Often required; no Matter impact
Management	User managed, self-service, automated	IT-managed; enterprise tools
Automation systems	Cloud Service Provider ecosystems	Industrial apps (Industry 4.0)
Device trust	Mass-market, implied trust	Customized, explicit trust

Source: Moor Insights & Strategy

Enterprises are often reluctant to consider unlicensed spectrum wireless device connectivity because of preconceived notions about technical limitations, the availability of workarounds for those limitations, platform availability, and cost. Table 1 compares how consumers and enterprises view eight considerations for connecting IoT devices. The first five considerations (highlighted in green) have practical workarounds, but the last three (highlighted in orange) do not.

The comparison is quite nuanced, so please read the comments below and do your own research before applying these guidelines to specific enterprise use cases. This analysis uses Matter as a proxy for "Matter-like" solutions with an application layer appropriate for the targeted vertical industry.

- **Spectrum** – Home networking uses the shared, unlicensed industrial, scientific, and medical (ISM) bands. Thread, Bluetooth, and Wi-Fi use the 2.4 GHz band, and Wi-Fi can also use 5 GHz, 6 GHz, and 900 MHz bands. For over 20 years, enterprises have used Wi-Fi for business applications like mobile device connectivity while deploying wired networks for large-scale, mission-critical operational infrastructure. Wireless networks have considerable cost and flexibility advantages, so many enterprises are aggressively migrating to licensed spectrum (i.e., private 5G) to meet stringent requirements for QOS and reliability. However, Wi-Fi 6 and 7 remove many channel capacity and speed limitations that prevent industrial use of unlicensed wireless. Industrial IoT applications will use a mix of unlicensed and licensed spectrum depending on technical requirements, customer preference, and cost-benefit analysis.
- **Infrastructure cost** – Network equipment for home and small business applications must be simple and inexpensive because consumers are reluctant to invest in infrastructure. Most consumers just plug in routers and hope for the best. Enterprises professionally install industrial-grade Wi-Fi equipment, but the infrastructure still costs substantially less than licensed spectrum alternatives such as private 5G.
- **QOS, performance** – Consumer Wi-Fi networks typically have a single router, and performance degrades with distance from that router. Bandwidth-hungry apps like video streaming require high-performance whole-home coverage, so router makers are embracing the new spectrum (6 GHz) in Wi-Fi 6 and 7 to avoid interference and create mesh networks of extender devices. MI&S expects this trend to accelerate as mesh prices drop, so "overprovisioned mesh" networks will dramatically enhance home network coverage and bandwidth. Thread networks already have a mesh topology, so whole-home coverage improves as the number of Thread devices increases. Enterprises, on the other hand, don't leave QOS to chance. IT personnel already measure QOS and guarantee Wi-Fi coverage, packet loss, latency, and reliability. Performance isn't a practical barrier for most commercial applications, particularly with Wi-Fi 6 and 7.
- **User authentication** – For smart homes, Matter's security model is simple, elegant, and designed for consumers. Users authenticate to a home automation ecosystem such as Alexa, Google, or HomeKit, and that ecosystem installs and

manages devices at the user's direction. Enterprises already use role-based user authentication to control Wi-Fi network access (i.e., WPA2-Enterprise). Matter, or a Matter-like application layer, could connect devices to an application domain. Controlling the domain requires user authentication, but the individual Matter devices do not. Enterprises could begin using Matter as-is by defining a single "Matter administrator" role for each domain.

- **Multi-site** – Although Matter specifications do not address managing networks across multiple physical locations, each Matter instance (fabric) is independent. Hence, managing multiple fabrics is mainly a function of the automation ecosystem (or enterprise domain). Use cases like multi-family housing are ecosystem-level features that don't necessarily require extensive changes to the Matter spec. Each housing unit can have a separate fabric, and the automation ecosystem can manage each one independently. Multi-family would be a new type of ecosystem application.
- **Management** – Most consumers set the Wi-Fi SSID and password, but some consumers don't even do that. Enterprises monitor and manage all networks with IT personnel and network management tools. To meet this requirement, Matter and Thread need network management extensions and other enhancements that are not on the horizon. This is a showstopper for many enterprises.
- **Automation systems** – For consumers, application ecosystems like Google Home, Amazon Alexa, Apple HomeKit, and Samsung SmartThings provide home automation, ambient interaction, and, eventually, autonomous operation that make smart homes truly "smart." Enterprise and industrial automation applications are specific to each installation and require extensive customization, so instead of a handful of ecosystems, there are dozens for each customer — maybe hundreds. Matter's multi-ecosystem (multi-admin) does not currently scale beyond a few consumer ecosystems, so this enterprise requirement is a showstopper.
- **Device trust** – Please refer to the white paper "Matter — Making Smart Homes More Secure" for a complete discussion of Matter's device trust architecture, but here's a summary of the device attestation process. The CSA certifies all products to verify compliance with the specifications and confirm interoperability. During product installation, Matter checks the unique attestation certificate in each manufactured device to confirm that it is certified by the CSA and built by a trusted manufacturer. This is an implicit trust model — trust in the CSA certification certificate and the product attestation certificate authority implies trust in the device. Implicit trust is ideal for mass-produced consumer products but does not address complicated industrial supply chains with multiple

touchpoints that customize device hardware and software. To fix this, Matter needs an explicit device trust model for enterprises and a new set of installation procedures. One such trust model is FIDO Device Onboard (FDO)[4], a standard from the FIDO Alliance for explicitly assigning device ownership after manufacturing. Enterprise Matter devices with multi-stage supply chains need an explicit trust model that follows devices through every development stage from silicon to deployment.

Two additional deployment situations are showstoppers for unlicensed spectrum networks regardless of application domain — mobility and remote locations. These situations require cellular, satellite, or LP-WAN networks. Theoretically, a Matter-like application layer would work okay over cellular or satellite, but Matter's provisioning and deployment features would need considerable modification. LP-WAN networks such as LoRaWAN do not have sufficient bandwidth to support Matter at the device level. However, LoRaWAN gateways could support Matter in the same ways that Matter Bridges support legacy protocols, but this idea needs validation.

CONCLUSION

IP networking brought computing out of the cold room, cellular services gave us mobility, and now satellite constellations promise global coverage. Still, most wireless internet traffic travels over unlicensed radio spectrum — particularly Wi-Fi — and this is unlikely to change. Unlicensed networking growth is poised to accelerate even more rapidly as Wi-Fi 6 and 7 improve performance and low-power devices converge on Thread. IP networking protocols allow product developers to use any combination of radio technologies — unlicensed (Wi-Fi, Thread, Bluetooth LE, and LP-WAN) or licensed (cellular or satellite). Standardizing IP-based interoperability protocols for specific vertical industries is the next step, and consumer electronics manufacturers are leading the way with Matter.

Unlicensed networks are the lowest cost and most broadly deployed options for most products, and integrating Wi-Fi, Thread, and Bluetooth LE into IoT devices has never been easier. For instance, NXP's Tri-Radio chips streamline wireless hardware and software development, enabling developers to specify chips and chipsets that "just work" with off-the-shelf platform software — easy certification, no spectrum licensing, no service provider, and predictable performance. Lower costs, better performance, and

[4] FIDO Device Onboard – https://fidoalliance.org/intro-to-fido-device-onboard/

universal IP-based device connectivity enable product companies to scale rapidly as AI expansion creates a vacuum for real-world data.

The compound effects of multimodal AI[5], universal device connectivity, and interoperability standards create a powerful data vacuum — an unprecedented environment for IoT growth. Metcalfe's Law, explained in the "Specialized Device Networks" section above, predicts that adding nodes to a network causes a nonlinear increase in the network's value. This symbiotic situation creates a flywheel effect that links IoT growth with AI expansion.

FIGURE 1: THE AI-IOT FLYWHEEL

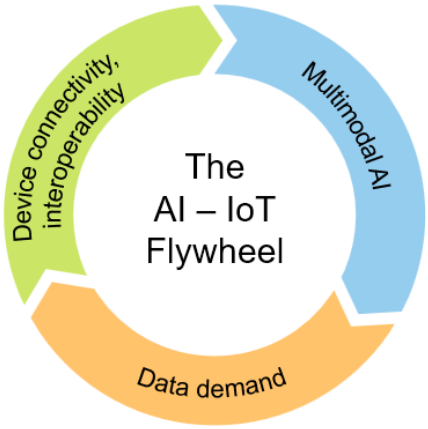

Source: Moor Insights & Strategy

Figure 1 illustrates the relationship between AI and IoT. Multimodal AI applications demand ever-increasing amounts of data from connected devices. Product companies satisfy this demand by maximizing connectivity and interoperability, and increased data availability accelerates AI growth. As the flywheel speeds up, economies of scale reduce device costs, adding more acceleration.

This is a perfect storm. Universal IP-based connectivity is AI's central nervous system — the only direct, unfiltered source of truth for training and inference. The combination of data-hungry AI-based applications, low-cost connected devices, and standards-based interoperability fundamentally changes the economics of embedded computing, delivering economic and social impact far exceeding the combined value of the individual connected devices.

[5] Multimodal AI is a machine learning model capable of using multiple types of inputs such as text, images, video, audio, and physical sensors. (Author's definition)

IMPORTANT INFORMATION ABOUT THIS PAPER

CONTRIBUTOR
Bill Curtis, Analyst In-Residence, Industrial IoT and IoT Technology

PUBLISHER
Patrick Moorhead, CEO, Founder and Chief Analyst at Moor Insights & Strategy

INQUIRIES
Contact us if you would like to discuss this report, and Moor Insights & Strategy will respond promptly.

CITATIONS
This paper can be cited by accredited press and analysts but must be cited in-context, displaying author's name, author's title, and "Moor Insights & Strategy". Non-press and non-analysts must receive prior written permission by Moor Insights & Strategy for any citations.

LICENSING

DISCLOSURES
NXP commissioned this paper . Moor Insights & Strategy provides research, analysis, advising, and consulting to many high-tech companies mentioned in this paper. No employees at the firm hold any equity positions with any companies cited in this document.

DISCLAIMER
The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions, and typographical errors. Moor Insights & Strategy disclaims all warranties as to the accuracy, completeness, or adequacy of such information and shall have no liability for errors, omissions, or inadequacies in such information. This document consists of the opinions of Moor Insights & Strategy and should not be construed as statements of fact. The opinions expressed herein are subject to change without notice.

Moor Insights & Strategy provides forecasts and forward-looking statements as directional indicators and not as precise predictions of future events. While our forecasts and forward-looking statements represent our current judgment on what the future holds, they are subject to risks and uncertainties that could cause actual results to differ materially. You are cautioned not to place undue reliance on these forecasts and forward-looking statements, which reflect our opinions only as of the date of publication for this document. Please keep in mind that we are not obligating ourselves to revise or publicly release the results of any revision to these forecasts and forward-looking statements in light of new information or future events.