IEEE PIMRC 2026

Keynote and Technical Session Highlights Report

Oscar Ashburn
Victoria University of Wellington

1st–4th September, 2026

Abstract

This report summarizes keynote and technical session content from IEEE PIMRC 2026. Keynote material is organized thematically around AI-native network design, distributed/privacy-preserving learning, capacity growth, and security. Two further sections cover an ISAC tutorial and a session on ceiling-mounted UWB indoor positioning.

Contents

1 Executive Summary

The overarching framing running through PIMRC 2026 is that the future of the internet is to become **AI-native**: the network itself is expected to host, train, and serve AI rather than simply carry traffic for AI applications running elsewhere. What follows are the open problems and current focus points driving this shift, as raised by speakers across academia, industry, and government.

A key takeaway across the conference: privacy, cybersecurity, and safety were spoken about positively, as a significant part of technology development rather than an afterthought. Security and trust were described as foundational design constraints, not add-ons (Wen; Ong Chen Hui), from privacy-preserving federated learning to starting quantum-safe migration now.

PIMRC 2026 framed 6G not as an incremental successor to 5G but as a platform-level convergence of **communications, computing, and AI**, with digital twins and edge intelligence as defining features (Tao, IMDA). A recurring architectural question was where distributed AI should sit (device/edge/cloud) and what network conditions that placement demands.

On the AI side, speakers argued for **communications-native foundation models** over task-specific ones, and for **goal-oriented/semantic communication** that optimizes for understanding rather than bit fidelity (Tao). Capacity growth remains a central driver: LLM text traffic under-delivered on predicted demand, but **video-generation AI** and **autonomous driving** are expected to drive the next major jump, motivating a $\sim$100$\times$ 5G capacity target via extreme MIMO (Fettweis).

Compute integration was flagged as one of the hardest open problems: GPU–memory bandwidth, power, and lagging semiconductor progress bottleneck AI supercomputing, with even 2-million-chip systems reportedly achieving only $\sim$40% effective throughput (Wen, Huawei). **Federated learning** was repeatedly proposed as a privacy-preserving alternative to centralized training, though it remains slower and raises open questions around poisoning defences, opt-out/unlearning, and contributor incentives.

Singapore's IMDA outlined the difficulty it expects when deploying robots in heterogeneous, multi-vendor, multi-owner environments, and raised the question of establishing **device/machine identity** before granting network access. **Quantum-safe migration** was presented as an urgent, multi-year undertaking that should begin now rather than once threats materialize. A consistent cross-cutting message: no single vendor, network, or country can solve 6G-era problems alone, requiring international, cross-vendor collaboration along an idea $\rightarrow$ prototype $\rightarrow$ deployment pipeline.

Among the technical talks, two stood out. The **ISAC tutorial** covered dual-functional radar-communications (DFRC), where a single waveform serves both sensing and communication roles, along with joint beamforming trade-offs and physical-layer security threats (notably sensing-privacy leakage to passive eavesdroppers). The **UWB radar session** addressed device-free indoor human positioning via ceiling-mounted sensors, where the central challenge is that a single person produces multiple distinct radar reflections that must be consolidated into one track rather than mistaken for several people. Objects positioned beneath the sensor were also noted as a significant source of noise.

2 Introduction

PIMRC 2026 featured keynotes on 6G's shift toward deep integration of communications, computing, sensing, and AI. Four keynotes are summarised: **Reshaping 6G** (Meixia Tao, Shanghai Jiao Tong

University); **Unlocking Frontiers of 6G Radio Access** (Gerhard Fettweis); **6G Integrated Communications and Computing** (Dr Tong Wen, CTO, Huawei); and **Networks for an AI-Native Future** (Dr Ong Chen Hui, IMDA Singapore). Two technical sessions are also covered: an **ISAC** tutorial and a session on **ceiling-mounted UWB indoor positioning**. This report is based on personal session notes, some of which were fragmentary, so details should be checked against the official proceedings.

3 AI-Native 6G: A New Architectural Paradigm

Tao; Ong Chen Hui. Unlike prior generations built principally for connectivity, 6G is framed as the age of **AI and digital twins**: deep integration of compute, sensing, communications, and intelligence, including on-device/edge training and inference. This raises a two-sided question — what AI needs from the network, and how the network must change to support AI, including where distributed intelligence should sit. IMDA's stated preference is to resolve these questions **early**, ahead of demand.

4 Foundation Models, Channel Modelling, and Semantic Communication

Tao. Purely statistical or purely ray-tracing channel models are no longer sufficient alone, motivating AI-based channel modelling combining data-driven and physics-based approaches. Task-specific AI models are costly to store at scale, motivating a **unified foundation model** over narrow single-purpose ones. A related theme was **goal-oriented/semantic communication**, optimising for whether the receiver understands meaning rather than bit-level fidelity. Models trained on non-communications data underperform on comms tasks, favouring training on **communications-only data from the outset**.

5 Radio Access, Capacity, and Spectrum for 6G

Fettweis; Tao. LLM-driven text traffic (post-2021) grew less than expected, being low-bandwidth; **video-generation AI** and **autonomous driving** are instead flagged as the next major demand drivers. A central goal is $\sim$**100× 5G capacity** via **extreme MIMO**. Tao added that 6G is expected to extend coverage via **non-terrestrial** and **low-altitude** platforms with new spectrum needs, still largely experimental.

6 Comms–Compute Integration and AI Supercomputing

Wen. Integrating wireless networking with computing infrastructure was framed as one of the hardest open problems, centred on **GPUs, memory, and power**. Wireless links aren't viable for GPU–memory interconnects, making bandwidth the bottleneck, compounded by lagging semiconductor progress. A guiding question: could a **500-billion-parameter LLM** train on a distributed network rather than one supercomputer? At-scale failures grow with system size; even $\sim$**2 million chips** reportedly achieved only $\sim$**40%** of theoretical throughput.

7 Privacy, Federated Learning, and Trust in AI Training

Wen; Ong Chen Hui. Centralized AI training requires trusting a data centre with one's data. **Federated learning** mitigates this by sharing only model coefficients, not raw data, though it remains much slower; **edge computing** and **federated distillation** were proposed to help close this gap. Poisoning attacks can be countered via periodic re-baselining, but post-hoc data removal (opt-out) and detecting malicious participants remain open problems. Incentivising contribution was raised as an unresolved audience question.

8 Anticipated Deployment Complexity: Robotics

Ong Chen Hui. Singapore builds network capacity **ahead of demand**, which is expected to enable heavier downstream usage. Looking at a mixed residential/business district, IMDA **expects** that robots reliable in one building may fail in others; that robots will need substantial **uplink** capacity and a way to communicate with building management systems; and that no single network, owner, or vendor will exist in these environments. Open questions include what connectivity quality robots actually need (task-dependent, not just throughput/latency), graceful degradation on outage, and whether **device/machine identity** should be verified before granting access.

9 Security and the Quantum-Safe Transition

Ong Chen Hui; Wen. Today's Public Key Infrastructure (PKI) could be compromised by future quantum computers; migration to **quantum-safe cryptography** is expected to take years, so the message was to start now. The question has shifted from designing quantum-safe systems to safely migrating existing operational ones. Singapore is investigating this nationally via a **Singtel–Ericsson–NCS–IMDA** collaboration on multi-vendor, end-to-end migration, echoing Wen's AI-security concerns: security and trust are foundational design constraints, not add-ons.

10 The Case for International, Cross-Vendor Collaboration

Wen; Ong Chen Hui. Both keynotes converged: no single vendor, network, or country can solve 6G-era problems alone. Validated, cross-vendor, cross-border development and common interoperability standards are needed, following an **idea → prototype → deployment** pipeline, with deployment itself expected to generate further research questions.

11 Technical Tutorial: ISAC Fundamentals and Prototyping

This hands-on tutorial covered **integrated sensing and communication (ISAC)**, also known as dual-functional radar-communications (DFRC): a single waveform serving both sensing and communication roles, turning every cell tower into a potential radar tower.

Applications (often UAV-paired): wide-area coverage, AI-driven drone delivery, joint air/ground infrastructure, smart homes, gesture recognition, and defence, all constrained by **spectral coexistence**.

Gains: joint DFRC design enables scalable spectrum trade-offs between sensing and communication, plus synergies (e.g. a channel estimate serving both functions).

Signal design: a comms-centric approach was presented, where the same signal used for sensing also carries embedded communication data, subject to bandwidth, frequency, polarisation, and throughput trade-offs. Coherent processing gain scales with integration time/pulses combined; range/Doppler ambiguity bounds decidable range and velocity; range accuracy depends on constellation geometry. Short-range profiles are harder to resolve due to clutter, requiring estimators focused on a defined range of interest.

Joint beamforming allows flexible time/frequency/spatial resource allocation between radar and comms, at the cost of high implementation complexity.

Security threats include data confidentiality, sensing-privacy leakage (unauthorized target localisation), and localisation uncertainty. A proof-of-concept eavesdropping demo showed that a passive eavesdropper (nicknamed "Eve") who only listens, and never transmits, cannot be spotted using the system's usual channel measurements (CSI). Even so, Eve could potentially pick up both the tracking information (such as where a person or object is) and the data being communicated, although the eavesdropper recovers meaningfully less than an intended receiver that knows how the signal was designed.

12 Technical Session: Indoor Human Positioning via Ceiling-Mounted UWB Radar

This session addressed **UWB radar** for **device-free** indoor human positioning, using a **ceiling-mounted**, largely top-down sensor configuration. A key challenge: a person is not a single point target — reflections from torso, arms, head, and legs can appear as multiple distinct detections, complicating person-level tracking. The notes suggest this ambiguity may be resolved by visual inspection of RF data (e.g. range-Doppler or point-cloud plots), distinguishing a single person's multi-part reflections from genuinely separate targets. Objects positioned beneath the sensor were also noted as a significant source of noise.

13 Conclusion

PIMRC 2026 points to a future internet that natively hosts, trains, and senses with intelligence, rather than merely carrying traffic for AI systems that live elsewhere. Across the keynotes, privacy and security were treated as enablers of this future rather than afterthoughts: federated learning is what makes distributed AI training trustworthy enough to scale, and early quantum-safe migration keeps tomorrow's network secure precisely because the work starts now, not once threats materialise. The recurring call for international, cross-vendor collaboration reflects the same instinct, that trust, like the technology itself, has to be built deliberately. The ISAC and UWB sessions offered a preview of this in practice: sensing-capable infrastructure that takes eavesdropping and privacy risks seriously from the start, and human-positioning systems solving the disambiguation problems any deployed sensing network will eventually face.