

Accelerating Hyperscale Cloud with 400G Network Solution

Presenter: Lucy Yu Faye Yang

Enhancing cloud performance with ultra-fast
networking technology



Introduction

Azure's Infrastructure Scaling and 400G Networking

Infrastructure Scaling

Azure is rapidly expanding infrastructure to support ever more demanding AI and high-performance computing workloads.

400G Networking Integration

Integration of 400G front-end networking enables ultra-low latency, high throughput, and energy efficiency.

SmartNIC Acceleration Evolution

SmartNIC technology advanced from 40Gbps to 400Gbps, achieving a tenfold bandwidth increase over a decade.

Qualification and Integration

Teams ensure secure, scalable deployment by qualifying hardware and firmware and validating platform integration.



Cloud Platform Complexity

Second Source & DDR5 Adoption

Second Source Vendor: Beyond Redundancy

It's about understanding the full support ecosystem of onboarding challenges.

Forces us to rethink how we evaluate platform fit, not just by specs, but by operational realities

Encourages deeper collaboration between engineering and supply chain teams.

DDR5 Adoption: Power, Thermal, and Integration Challenges

Redesigning Power Delivery Networks: DDR5 memory adoption required low-impedance decoupling and enhanced VRM (Voltage Regulator Module) tuning for system stability.

Advanced Thermal Management: Optimized PCB layouts and thermal interface materials were essential to manage hotspots and signal routing constraints.

Early Co-Design & Validation: Robust frameworks and cross-functional collaboration ensured reliable system integration of DDR5 memory.



Security and Qualification Challenges

Evolving Security Standards

Security standards must adapt to address hardware vulnerabilities such as Spectre and Meltdown in modern cloud environments.

Complex Qualification Process


Advancing SmartNIC generations increase the complexity of qualification processes requiring more robust validation.

Cross-Functional Collaboration and Co-engineering

Collaboration between Azure and silicon vendors ensures early silicon access and tooling readiness to align development.

Unified Roadmap and Transparency

Asynchronized roadmap and transparency in partner timelines streamline execution and reduce integration delays.



Firmware Integration Strategy & Flowchart

Firmware Integration Strategy

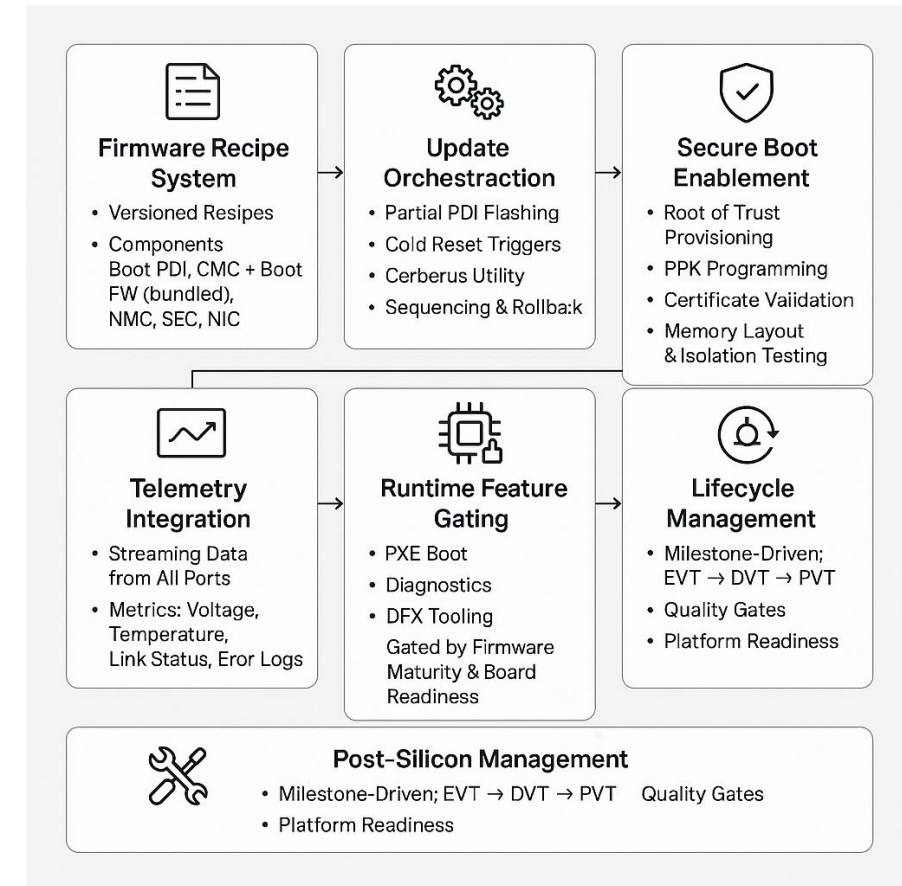
Modular Firmware Deployment: Scalable deployment using versioned firmware recipes with validated Boot PDI (Programmable Device Image), CMC (Cerberus Management Controller), NMC (Network Management Controller), and NIC (Network Interface Card) images.

Secure Boot & Telemetry: Root of Trust provisioning, PPK (Platform Provisioning Key) programming, and certificate validation ensure trusted firmware execution; integrated telemetry streams diagnostic data from all ports for performance monitoring and robust post-silicon support workflows.

Milestone-Driven Update Process

Update Orchestration: Firmware updates are sequenced through milestone-driven steps, using partial flashing and cold resets to ensure secure and reliable deployment.

Lifecycle & Post-Silicon Management: Progression through EVT, DVT, and PVT stages with quality gates ensures platform readiness, while coordinated agents and external services manage provisioning, revocation, and debugging after silicon deployment.



System Integration Variations

Multi-Host and Multi-NIC Configurations

Multi-Host Functionality

Multi-host setups allow a NIC to serve multiple CPUs through dedicated PCIe lanes and independent reset signals.

Multi-NIC Management

Multi-NIC configurations provide traffic isolation, redundancy, and bandwidth aggregation for enhanced network performance.

Thermal Strategy and Liquid Cooling

Limitations of Air Cooling

Traditional air-cooling methods struggle with increased data center workload heat dissipation demands. Next-gen servers use direct-to-chip liquid cooling with network accelerators for efficient thermal management.

Optimized Cold Plate Design

Carefully designed cold plates and thermal models ensure reliable, scalable, and energy-efficient deployments.



Deployment and Validation Challenges




Advanced System Designs

Modern 400G cards support different platform architect, increasing complexity in deployment and validation.



Validation Complexity

Firmware partitioning, PCIe coordination, and reset logic across configurations challenge qualification teams.



Simplification Benefits

Standardizing interfaces and simplifying hardware parameters enable efficient and predictable validation cycles.

Future Directions

Trends and Outlook for Azure Infrastructure

High-Performance Networking

Azure's 400Gb Boost offers top-tier networking and storage speeds, boosting cloud infrastructure performance.

AI-Driven Automation

AI enables real-time monitoring and self-healing, improving infrastructure reliability and operational efficiency.

Security and Compliance

Emphasis on encryption and zero-trust models ensures strong security and regulatory compliance.

Sustainability and Collaboration

Focus on energy efficiency and partnership with vendors promotes sustainable, interoperable cloud solutions.

