OCT 2025/ VOLUME 1

THE NERD'S SHOW MAGAZINE

HYPERSCALE COMPUTING ARCHITECTURE

SHESHA KRISHNAPURA

Intel Fellow and IT CTO

OCT 2025

THE FUTURE OF MOORE'S LAW

EXCLUSIVE INTERVIEW

FROM BANGALORE TO SILICON VALLEY

TOPICS COVERED:

- INNOVATIVE IDEAS CHALLENGING INDUSTRY NORMS
- EXCLUSIVE INSIGHTS FROM INTEL FELLOW
- SUSTAINABLE SERVER ARCHITECTURE
- BRIDGING TECHNOLOGY AND LEADERSHIP

Contents

1 Introduction: The Imperative for Change

The global data center industry consumes approximately 2%-3% of worldwide electricity today and growing significantly due to AI datacenter build outs, with server hardware representing the largest component of this energy demand. Yet beneath these aggregate statistics lies a more troubling reality: the systematic waste embedded in how we design, deploy, and replace computing infrastructure.

For over three decades, the enterprise computing industry has operated on a fundamental assumption that servers should be monolithic systems replaced every 4–6 years regardless of individual component health. This approach, rooted in the early days of computing when systems were less reliable and standardization was limited, continues despite dramatic improvements in component durability and manufacturing quality.

The Scale of the Problem:

- 13-14 million servers deployed globally each year
- $200 billion annual enterprise server market
- 62 million tons of electronic waste generated annually
- 35–40% of server capacity typically underutilized

Source: IDC Global Server Market Report 2024, EPA Electronic Waste Statistics, Gartner Infrastructure Analysis

2 The Visionary: Shesha Krishnapura's Journey

To understand the revolutionary potential of disaggregated computing, one must first understand the mind that conceived it. Shesha Krishnapura's journey from the digital voice/data switching Research and Development division of Indian Telephone Industries to the executive corridors of Intel represents more than career advancement—it embodies a philosophy of continuous learning, fundamental questioning, and practical innovation.

2.1 The Formative Years: Building Deep Foundations

Krishnapura's technical foundation was sculpted during his tenure at Indian Telephone Industries (ITI) from December 1986 to August 1989. This period, characterized by intensive work on real-time operating systems and assembly-level programming, established the systems-thinking approach that would later revolutionize server architecture.

> "The foundation of any successful career must be built on solid ethics and through that, credibility. One has to shape their direction through continuous learning and adaptation."

The intensity of those early years—working past 2 AM, sleeping on dot-matrix printouts, returning at 6 AM after brief rests—exemplifies the dedication required for breakthrough innovation. This experience with Intel microprocessor internals, particularly the 8085, real-time OS coding in assembly language, and digital switching systems scaling from 256 lines to 2,048 lines, provided the deep technical grounding that would prove essential for later innovations.

The ITI Foundation: Technical Mastery

Key Technical Competencies Developed (1986–1989):

- Intel microprocessor internals and architecture
- Real-time operating system development in assembly
- System-level design and component integration
- Analog-to-digital conversion systems
- Routing algorithms and network protocols
- Scalable digital switching systems
- Optimal coding within limited memory space (EPROM)—learning how to manage resource constraints

This foundational experience proved crucial not just for technical skills, but for developing the systematic thinking required to question fundamental assumptions about system design. The constraints of early computing—limited memory, processing power, and reliability—forced creative solutions that would later inform more efficient approaches to modern challenges.

2.2 Continous Learning Through Teaching

For the past eight years, Krishnapura has maintained a teaching role at Stanford University, sharing 40% of the ME344 curriculum (Introduction to High Performance Computing) with Steve Jones, the university's supercomputing director. This commitment to education reflects a deeper philosophy about knowledge synthesis and continuous learning.

> "When you share synthesized knowledge, you're not only teaching others—you're teaching yourself. You're constantly honing your own fundamentals and discovering new connections."

Krishnapura's teaching methodology emphasizes active engagement over passive absorption. Rather than simply presenting information, he consistently asks students: "What do you think this is? Why do you think it works this way?" This approach creates an interactive learning environment that challenges assumptions and builds deeper understanding.

Teaching Philosophy in Practice:

- Socratic questioning to build understanding
- Collaborative problem-solving approaches
- Bidirectional learning between instructor and students

This pedagogical approach extends beyond the classroom into his professional work. When evaluating technology solutions, he applies the same questioning methodology, challenging his team to articulate not just what they're building, but why they're building it that way and what alternatives might exist.

3 The Innovation: Disaggregated Server Architecture

The most significant contribution of Krishnapura's career lies in his revolutionary work on disaggregated server architectures.

This innovation challenges the fundamental assumption that enterprise servers should be monolithic "pizza box" systems (Rack-Mount) replaced every 4–6 years, regardless of individual component health or it's longevity due to matured technology (Example: Power Supplies, Disk Drives, SAS Controller, Integrated Network Sub-System, Chassis Controller, etc.) or actual workload requirements/value of a subcomponent of a server.

3.1 The Efficiency Crisis

Traditional server replacement cycles create massive e-waste that most organizations accept as inevitable. Krishnapura's analysis reveals this significant e-waste through simple but profound observations about difference between typical consumer behaviour versus datacenter operator in infrastructure upgrade practices.

> "When did we last change the kitchen exhaust fan or bathroom exhaust fan? Only when it burned out. Then why do we discard perfectly working server fans every four to six years?"

This seemingly simple question reveals the profound e-waste embedded in datacenter operations.

The Waste Analysis: Shocking Statistics

- Storage Drives: 0.8% annual failure rate means 92% continue to function even after 10 years
- Network Components: Often remain unchanged for 10–20 years based on bandwidth requirements
- Memory Modules: Modern ECC memory has extremely low failure rates yet is replaced along with server as part of the refresh even if the existing memory speed is good enough

The economic implications are staggering. Organizations routinely discard millions of dollars in perfectly functional hardware simply because it's packaged in outdated server box.

Figure 1: Krishnapura celebrates the scale and flexibility enabled by disaggregated servers.

3.2 Component-Level Optimization

Krishnapura's solution involves disaggregating servers based on actual component technological value lifecycles and failure rates rather than arbitrary replacement schedules. This approach, implemented across 22 unique server designs as of June 2025 under his leadership, represents a fundamental shift toward sustainable computing infrastructure.

Benefits of disaggregated architecture

- Reduces technician time spent on refresh by 77% (lower operational cost).
- Decreases refresh materials' ship weight by 82% (reduction in e-waste).
- No need to replace good components (significant cost savings).
- No need to reinstall the OS (simplified faster operations).

3.3 Challenges and Solutions

The transition to disaggregated architecture requires overcoming significant organizational and technical through disruptive thinking. Traditional procurement processes, maintenance contracts, and operational procedures are all designed around monolithic server replacements.

Krishnapura's approach involves gradual implementation, starting with new deployments and expanding as organizations develop capability and confidence. This phased approach reduces risk while demonstrating clear benefits that justify broader adoption.

4 Understanding Moore's Law: Beyond the Headlines

Krishnapura's discussion of Moore's Law reveals the depth of understanding required to navigate the semiconductor industry's current challenges. Rather than accepting popular interpretations, he insists on fundamental comprehension of what Gordon Moore, Intel's co-founder, actually observed.

4.1 The True Definition

> "Many people don't understand what Moore's Law actually states. It's about transistor density—the number of transistors you can pack in a given space doubling every 18–24 months. But this is just one dimension of PPA: Power, Performance, and Area."

Moore's Law represents just one aspect of semiconductor advancement. The industry's focus has expanded to encompass the broader PPA framework, recognizing that raw transistor density alone doesn't determine system performance or value.

The conversation becomes particularly engaging when Krishnapura uses a COVID-19 analogy to illustrate the scale of modern semiconductor manufacturing. This comparison provides crucial context for understanding the extraordinary achievement of current chip manufacturing capabilities.

> "We are making transistors much smaller than a COVID virus. When COVID came, people didn't know how it spreads, what it is, how it moves. It took huge numbers of research people to figure it out. We're manufacturing billions of functional electronic components significantly smaller than that virus for decades."

Scale of Modern Manufacturing:

- COVID-19 virus diameter: 80-100 nanometers
- Intel's 14nm process: Achieved 12 years ago
- Current leading-edge processes: 3–5 nanometers, very soon 1.8 nanometer.
- Manufacturing precision: Creating functional structures smaller than viruses
- Complexity: Billions of transistors in space smaller than fingernail

4.2 The Evolution of Semiconductor

The progression of semiconductor manufacturing capabilities represents one of humanity's greatest technological achievements. Krishnapura traces this evolution through Intel's process generations, demonstrating the relentless pace of improvement.

Intel Process Evolution

Historical Process Node Progression:

- 2006: 45 nanometer process introduction
- 2008: 32 nanometer process development
- 2010: 22 nanometer with FinFET technology
- 2012–2014: 14 nanometer mass production
- 2016–2018: 10 nanometer challenges and solutions
- 2019–2021: 7 nanometer and EUV lithography
- 2022–2024: 5 nanometer and 3 nanometer development

4.3 Beyond Moore's Law

As traditional Moore's Law scaling faces physical limits, the industry has evolved toward new approaches that continue performance improvements through architectural innovation rather than pure transistor scaling.

This shift toward architectural innovation aligns well with Krishnapura's work on disaggregated servers. Rather than relying solely on faster processors, optimal system design becomes the primary driver of performance improvement and significant economic value creation.

5 The Historical Context: Intel's Role in Creating Silicon Valley

Understanding Krishnapura's innovations requires appreciating Intel's foundational role in creating the collaborative ecosystem that defines Silicon Valley. This historical context explains not just technical achievements, but the cultural and organizational approaches that enable continued innovation through competitive collaborative mindset.

5.1 The Founding Vision

Intel's founding by Robert Noyce and Gordon Moore, with Andy Grove completing the leadership triumvirate, established principles that continue to influence the industry.

The company's focus on both technical excellence and systematic management created a model that other organizations emulated.

> "Silicon Valley is called Silicon Valley because of Gordon Moore and Robert Noyce. Steve Jobs used to visit Robert Noyce's home, discussing ideas in the basement. This collaborative spirit, this willingness to share ideas and challenge each other, defined the industry."

This collaborative approach, where competitors shared ideas and talent flowed freely between companies, created an innovation ecosystem that no single organization could have achieved alone. The culture of questioning, experimentation, and continuous learning that characterized early Intel continues to influence how technical leaders approach complex challenges.

Figure 2: Intel's Early Leadership Team

5.2 The Leadership Philosophy

The Intel leadership model emphasized systematic thinking, rigorous analysis, and constructive confrontation of ideas. Grove's famous "Only the Paranoid Survive" philosophy encouraged constant questioning of assumptions and preparation for disruptive changes.

Intel Leadership Principles:

- Constructive Confrontation: Challenge ideas, not people
- Systematic Analysis: Data-driven decision making
- Continuous Learning: Adapt to changing circumstances
- Technical Excellence: Never compromise on quality
- Collaborative Innovation: Share ideas to advance the industry

6 Leadership Philosophy: The Foundation of Innovation

Throughout our extensive conversation, several key intel leadership principles emerged from Krishnapura's approach to innovation and technical management.

These principles provide a framework for understanding how breakthrough innovations emerge from systematic yet disruptive thinking and disciplined quality execution.

6.1 Continuous Learning by Design

Krishnapura's journey from debugging digital switching systems to becoming an Intel Fellow illustrates the critical importance of continuous learning in technical leadership. His commitment to teaching at Stanford while maintaining his executive responsibilities demonstrates how sharing knowledge strengthens personal understanding and drives innovation.

> "Every interaction is a learning opportunity. When you teach, you're forced to synthesize your knowledge in new ways. When you listen to questions, you discover gaps in your understanding. When you collaborate with others, you see familiar problems from new perspectives."

6.2 Questioning Assumptions

The disaggregated server concept emerged from Krishnapura's keenness to question the fundamental assumptions about enterprise computing practices. This questioning methodology provides a framework for identifying innovation opportunities in established industries.

> "Most inefficiencies hide in plain sight, protected by assumptions that everyone accepts but nobody examines. The most powerful innovations often come from questioning the most basic assumptions about how things should work."

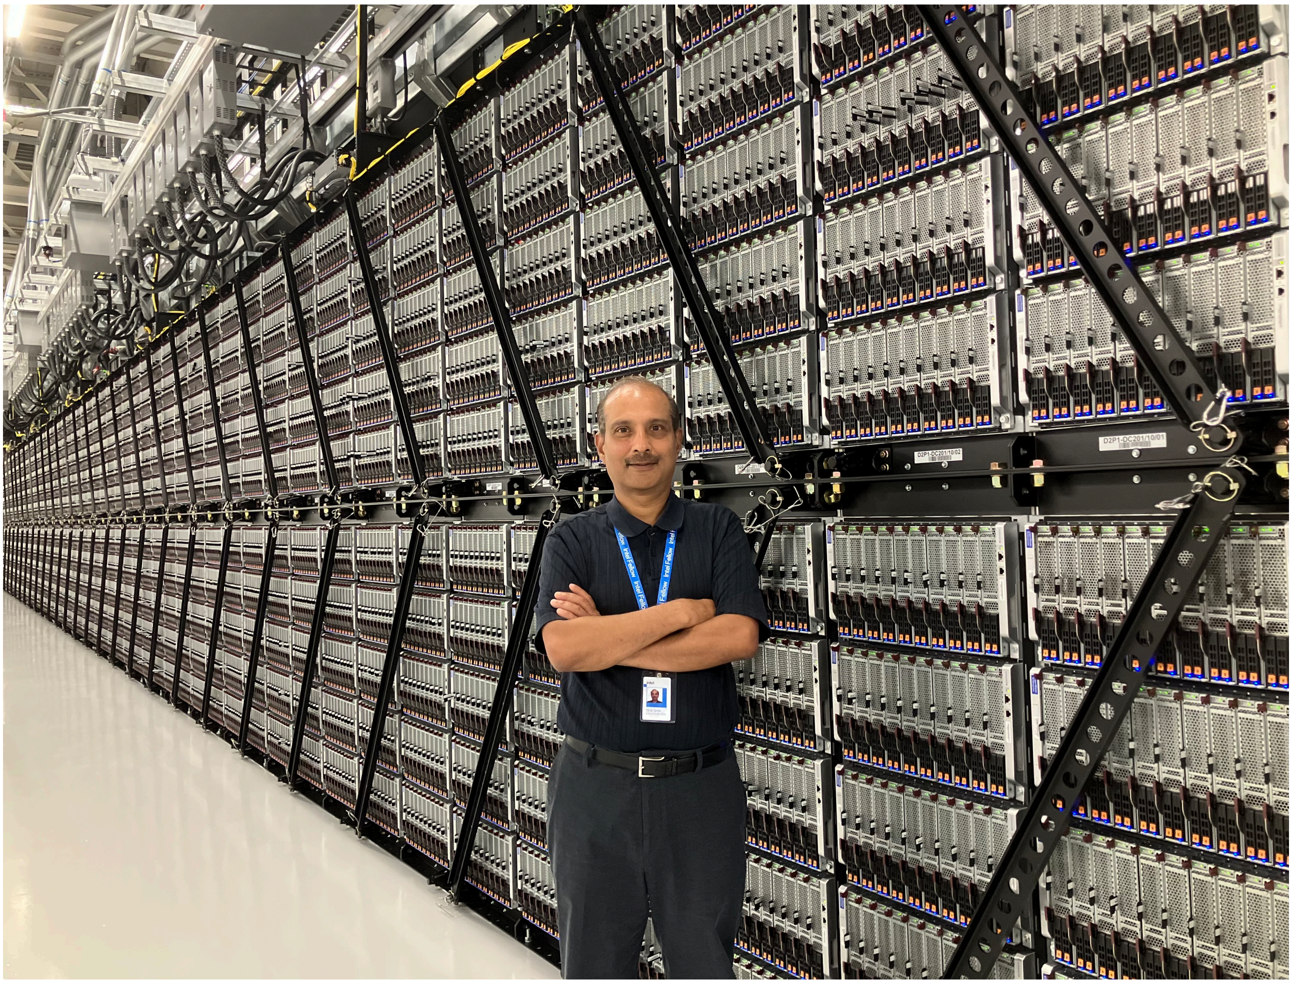

Figure 3: Continuous learning drives disciplined innovation across Intel's data centers.

6.3 Communication and Influence

Krishnapura's ability to explain complex technical concepts using everyday analogies demonstrates the importance of effective communication in technical leadership. His teaching methodology—using questions to engage audiences rather than simply broadcasting information—provides a model for influencing others and building consensus around new ideas.

> "Technical excellence alone is insufficient for creating change. You must be able to communicate complex ideas in ways that help others understand not just what you're proposing, but why it matters and how it connects to their experience."

7 Implementation Strategy: From Vision to Reality

Transforming revolutionary ideas into practical implementations requires systematic planning and execution. Krishnapura's approach to deploying disaggregated server architectures provides a model for implementing disruptive innovations in established organizations.

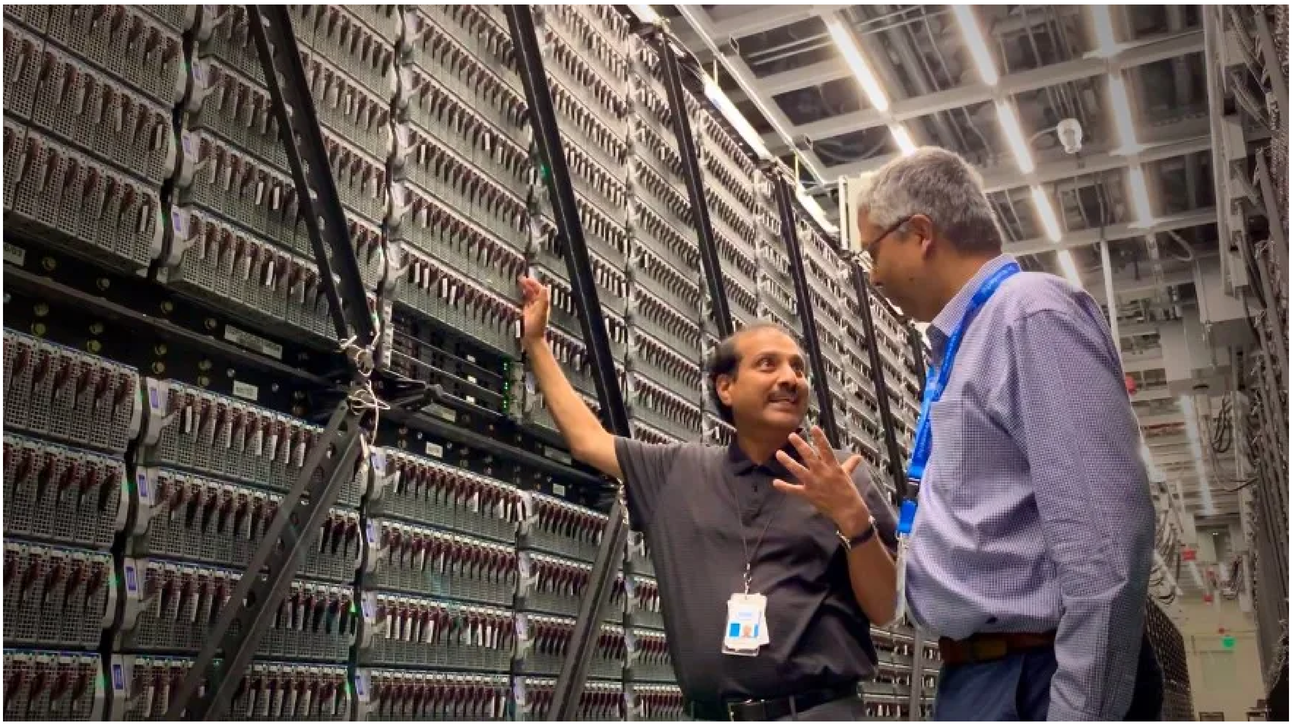

Figure 4: Shesha Krishnapura Explaining Disaggregated Server Infrastructure in a Data Center

7.1 The Phased Deployment

Rather than attempting wholesale change, Krishnapura advocates for phased deployment that reduces risk while demonstrating clear benefits. This approach enables organizations to build capability and confidence while validating assumptions about new technologies. By starting small and iteratively expanding, teams can learn from early stages, fix mistakes, refine processes, and secure organizational buy-in.

Each phase serves as both a test bed and a foundation for the next, ensuring that technical, operational, and cultural challenges are addressed before broader rollout. This proven technique not only avoids any negative effects of disruption but also accelerates adoption through tangible, incremental wins.

Implementation Phases

Phase 1: Proof of Concept (Months 1–3)

- Identify and resolve initial challenges

Phase 2: Pilot Deployment (Months 3–6)

- Expand to production workloads with limited risk

Phase 3: Scaled Implementation (Months 6–12)

- Deploy across multiple data centers and workloads

Phase 4: Full Transformation (Months 12+)

- Replace legacy systems with disaggregated architecture

7.2 Change Management

Implementing disaggregated architectures requires very simple changes in organizational processes, skills, and culture. Success depends on continous innovation and systematic change management that addresses both technical and human factors.

> "Technology innovation is only successful when it's accompanied by operational improvements. You must change not just what you're building, but how you're building it, who's building it, and why they're motivated to make it successful."

The change management approach includes changing mindset, comprehensive training, and new performance metrics that align with disaggregated architecture benefits. Organizations must also develop new supplier relationships and support processes that accommodate systematic continuous component-level innovation and optimization.

Critical Success Factors:

- Leading with purpose and clear vision communication
- Comprehensive training and skill development programs
- Revised metrics for value realization and new skill development
- Strong supplier partnerships and support relationships
- Continuous improvement culture and feedback mechanisms

About The Nerd Show Podcast

The Nerd Show is a monthly podcast that features in-depth conversations with leading technologists, scientists, and innovators who are shaping the future of technology. Each episode explores the intersection of technical excellence, leadership, and vision.

Credits

Figure 5: Editor-in-Chief: Mayank Sharma

THE NERD SHOW MONTHLY

Oct 2025 | Volume 1, Issue 1