

ID	Topic
1	CPU vs. GPU Architectures - Contrasting the Latency-Optimized design of CPUs (complex branching) with the Throughput-Optimized design of GPUs (massive parallelism)
2	Evolution of GPU Hardware - Tracing the path from early fixed-function graphics chips to the general-purpose GPUs used in modern AI
3	TPUs (Tensor Processing Units) - Google's custom ASICs that use Systolic Array designs to pass data through a grid of processors without constant memory access
4	Vector and Matrix Engines (AVX/SVE/AMX) - Specialized units within a CPU designed to perform vector and matrix operations to accelerate deep learning
5	Neuromorphic Computing - Computer chips that mimic how the brain uses "spiking neurons" to process information
6	Using the Cerebras AI Chip for Scientific Computing - Strategies for developing scientific computing applications on this non-traditional hardware platform
7	Julia for HPC - Utilizing the Julia programming language for high performance numerical and scientific computing
8	Domain-Specific Languages (DSLs) - Specialized languages like Halide that simplify writing high performance code for specific hardware targets
9	Lightweight Virtualization (Firecracker) - Using microVMs to provide the isolation of traditional virtual machines with the speed and efficiency of containers
10	Mixed-Precision Computing - Utilizing lower-bit formats (such as BF16) to accelerate scientific simulations while maintaining accuracy
11	In-Situ Analysis and Visualization - Processing and visualizing data in real-time while it is still in the memory of the supercomputer, avoiding the bottleneck of writing to disk
12	DNA Data Storage - The experimental use of synthetic DNA to store massive amounts of data in a biological format that can last for centuries
13	In-Network Computing with SmartNICs - Offloading computational tasks (like data reduction or encryption) directly to smart Network Interface Cards to improve HPC efficiency
14	Ultra Ethernet - An effort to evolve standard Ethernet into a high-speed, reliable fabric suitable for massive AI training clusters and HPC
15	Zero Trust Networking - A security philosophy where no device, user, or connection is trusted by default
16	Trusted Execution Environments (TEEs) - Hardware vaults such as Intel SGX that protect sensitive data and code even from the computer's own OS
17	Digital Sovereignty - The movement by nations to build independent hardware and cloud infrastructures to ensure data remains within their physical borders
18	Post-Quantum Cryptography Hardware - New chip designs built to run the complex math needed to stay safe from future quantum computer attacks
19	Processing-in-Memory (PIM) - While TPUs use systolic arrays to minimize memory access, PIM goes a step further by integrating logic directly into the memory chips (DRAM or SRAM)
20	Quantum Accelerators for HPC - Hybrid classical-quantum computing models and how emerging quantum processors may be integrated into traditional supercomputing workflows