Key Computing Technologies on the 2024–2025 Gartner Hype Cycles

Emerging computing hardware and architectures are evolving rapidly, and Gartner's Hype Cycles for 2024 offer a snapshot of where these innovations stand. This report examines major **hardware innovations** (like quantum, neuromorphic, GPU/accelerator chips, and cooling technologies) and **compute architecture trends** (edge computing, cloud-native platforms, composable infrastructure, etc.), highlighting their Hype Cycle positions and Gartner's perspective on their future benefits and risks. The analysis spans three Gartner Hype Cycles: **Emerging Technologies 2024**, **Compute (Infrastructure) 2024**, and **Data Center Infrastructure Technologies 2024** – which together outline the maturity of cutting-edge computing technologies through 2025. We provide an academic-oriented discussion for each technology, followed by a summary table of their Hype Cycle stages and maturity timelines.

Hardware Innovations in Computing

Quantum Computing (Quantum Processors)

Hype Cycle Position: *Peak of Inflated Expectations* (Compute 2024). Gartner places quantum computing near the peak of hype in 2024. In the *Hype Cycle for Compute 2024* chart, "Quantum Computing" appears at the top of the curve, indicating intense excitement but also a long road to maturity. Gartner analysts note that while prototype quantum processors exist, the technology remains **embryonic** with mass adoption still distant (Quantum Winter is Coming? Again and Again… | by Anastasiya Khromova, Dr. rer. nat. | Medium) (Gartner's Hype Cycle report suggests photonic computing is gaining traction | Electro Optics). Indeed, Gartner did *not* include quantum computing in its 2024 Top 10 Strategic Trends, observing that its near-term business impact is limited – "hardware is available, [but] our capabilities to work with it are not at the point where mass adoption can be expected… in the short term" (Quantum Winter is Coming? Again and Again… | by Anastasiya Khromova, Dr. rer. nat. | Medium).

Future Relevance & Benefits: In the long run, quantum computing holds **transformational potential**. It promises to solve classes of problems (e.g. cryptography, complex optimization, quantum chemistry simulations) that are intractable for classical computers. Gartner's emerging tech outlook suggests quantum computing is a "radical breakthrough" technology to watch (Gartner 2024 Hype Cycle for Emerging Technologies highlights ...). When (and if) fault-tolerant quantum processors become viable, they could provide exponential speed-ups for certain algorithms beyond the capability of any conventional supercomputer.

Risks & Challenges: The path to usefulness is fraught with challenges. Current quantum processors (NISQ devices) have limited qubit counts and require error correction breakthroughs. Gartner's commentary implies a risk of a "quantum winter" – a prolonged trough of disillusionment – if progress doesn't meet the hype. Organizations must temper expectations: Gartner advises that quantum computing's **commercial impact** is still 5–10+ years out (Gartner's Hype Cycle report suggests photonic computing is gaining traction | Electro Optics). In the interim,

businesses face risks investing too early: technical complexity, scarcity of quantum talent, and integration challenges with classical systems. The **upshot** is that quantum computing sits at an **early hype phase**, with tremendous future impact but significant scientific hurdles and unclear timelines for practical adoption (Quantum Winter is Coming? Again and Again… | by Anastasiya Khromova, Dr. rer. nat. | Medium).

Neuromorphic Computing (Brain-Inspired Chips)

Hype Cycle Position: *Trough of Disillusionment* (AI 2024). Neuromorphic chips – processors modeled after neural brain architectures – have advanced on Gartner's Hype Cycle but are still in an experimental stage. In the 2024 AI Hype Cycle, Gartner places **neuromorphic computing** as an innovation that has **passed the Peak** and is "soon to leave the peak or already in the Trough of Disillusionment" (Hype Cycle for Artificial Intelligence, 2024 | annotated by SHIH YEN). After some early hype, the technology has seen slower practical uptake, leading to reduced short-term expectations. However, Gartner notes that neuromorphic hardware **"advanced significantly in the past year"**, suggesting it could progress more rapidly through the remaining stages (Hype Cycle for Artificial Intelligence, 2024 | annotated by SHIH YEN).

Future Relevance & Benefits: Neuromorphic processors aim to achieve **orders-of-magnitude efficiency gains** for AI and sensor processing by mimicking how the human brain computes (spiking neural networks, event-driven processing). The potential benefits include ultra-low power consumption and real-time learning capabilities. Gartner views neuromorphic computing as a future enabler of sustainable AI hardware – balancing high computational power with minimal energy use (It's Here: The 2024 Gartner AI Hype Cycle™ | Gregory Renard). In an era where conventional chips hit energy and scaling limits, neuromorphic designs could unlock new applications in edge AI, robotics, and autonomous systems that require fast, local decision-making.

Risks & Challenges: Thus far, neuromorphic hardware is **largely confined to research labs** and niche prototypes. Software ecosystems and programming paradigms for these chips are immature, making adoption difficult. Gartner's positioning in the trough reflects these challenges: early enthusiasm has given way to recognition that neuromorphic tech needs further R&D and specialized tooling before it's widely useful. There is also competition from more immediate AI chip innovations (GPUs, ASICs) that dominate current AI workloads. **In summary**, neuromorphic computing is promising for long-term energy-efficient AI, but Gartner assesses it as a nascent technology ~5–10 years from productivity, currently facing a reality check on the Hype Cycle (Hype Cycle for Artificial Intelligence, 2024 | annotated by SHIH YEN).

Photonic (Optical) Computing

Hype Cycle Position: *Innovation Trigger* (Compute 2024). Gartner identifies photonic computing as an **"embryonic" technology with transformational impact**, positioned at the very early stage of the Compute Hype Cycle 2024 (Gartner's Hype Cycle report suggests photonic computing is gaining traction | Electro Optics). In the Hype Cycle chart (as of July 2024), **Photonic Computing** sits in the Innovation Trigger phase (far left), indicating that only research prototypes exist and broader interest is just beginning to build. Gartner assigns it a **>10 year timeline to reach Plateau**,

reflecting the long-term horizon for optical processors to mature (Gartner's Hype Cycle report suggests photonic computing is gaining traction | Electro Optics).

Future Relevance & Benefits: Photonic computing uses photons (light) instead of electrons for computation and data transfer. According to Gartner, it could **"address performance and energy consumption challenges"** in AI and data centers as Moore's Law falters (Gartner Hype Cycle for Compute - Q.ANT) (Gartner Hype Cycle for Compute - Q.ANT). Light-based processors inherently offer massive parallelism and high bandwidth. This means photonic chips could perform certain computations (like matrix multiplications for AI inference) much faster and with far less heat dissipation than silicon chips (Gartner Hype Cycle for Compute - Q.ANT) (Gartner Hype Cycle for Compute - Q.ANT). Gartner highlights photonics as a key to **"sustainable, high-performance computing"**, potentially enabling future AI systems that current electronics cannot support (Gartner Hype Cycle for Compute - Q.ANT) (Gartner Hype Cycle for Compute - Q.ANT). If realized, optical computing would be a paradigm shift – delivering ultra-fast processors for HPC and AI with improved energy efficiency.

Risks & Challenges: As an emerging paradigm, photonic computing faces significant development hurdles. Current optical components (lasers, modulators, detectors) must be integrated into viable processors, and doing so at scale is non-trivial. Early photonic processors are specialized (often targeting optical communication or specific math operations) and not general-purpose yet. Gartner's placement of photonics in the Innovation Trigger implies it's still in labs, with **no guarantee of commercial success**. The risk is that technical barriers (like integrating optics with existing digital systems, or creating optical memory) could slow progress. Additionally, photonic computing will compete with improving electronic chips and alternative technologies (quantum, neuromorphic) for investment. Gartner nonetheless views it as a *"transformational innovation"* to watch, given the fundamental physical advantages of computing with light (Gartner's Hype Cycle report suggests photonic computing is gaining traction | Electro Optics).

AI Acceleration Hardware (GPUs, ASICs, Accelerators)

Hype Cycle Position: *Plateau of Productivity* (as of mid-2024 for GPUs); *Peak of Expectations* for emerging AI ASICs. General-purpose **GPU accelerators** have been a staple of AI and high-performance computing and are now considered a mature, productive technology. In fact, **GPU acceleration is cited as an example of a Plateau of Productivity technology** – widely adopted with proven benefits (Introduction to the Gartner Hype Cycle – BMC Software | Blogs). By 2024, GPU-based computing is mainstream in data centers, enabling deep learning, data analytics, and scientific simulation at scale. Gartner's 2019 analyses had already placed GPU accelerators on the plateau (Analysis of Gartner Hype Cycles for AI - LinkedIn), and the continued explosion of AI in 2023–2024 (driven by generative AI) has further cemented GPUs' pervasive role. *(For instance, NVIDIA's AI GPUs saw huge demand in 2024, though Gartner projects a potential leveling-off of growth after the initial boom (AI chip sales predicted to jump by a third – then cool off • The Register).)*

At the same time, **new AI-specific accelerators (ASICs)** are at an earlier stage in the Hype Cycle. The Compute 2024 Hype Cycle shows **"Deep Neural Network ASICs"** at the Peak of Inflated

Expectations, reflecting the excitement around custom chips for AI (such as Google's TPUs, Graphcore IPUs, Habana Gaudi, etc.). These purpose-built chips promise even better performance-per-watt for AI workloads. Many startups and incumbents have touted novel AI accelerators, driving hype. However, like any peak-phase technology, not all will deliver, and some retrenchment is likely once real-world results come in.

Future Relevance & Benefits: Specialized AI hardware is critical for the **next era of computing performance**. Gartner's research emphasizes that as conventional CPUs plateau, acceleration via parallel architectures is essential (Gartner Hype Cycle for Compute - Q.ANT). **GPUs** (graphics processing units) pioneered this trend by offering massive parallelism for matrix and vector operations, enabling the deep learning revolution. Today, GPUs are indispensable for training large AI models and performing high-throughput scientific computations. They have a well-established ecosystem (CUDA, libraries, etc.) and deliver **high business value (high benefit)** on Gartner's scales (Introduction to the Gartner Hype Cycle – BMC Software | Blogs).

AI ASICs and accelerators take this further by customizing silicon to AI workloads (eliminating GPU features not needed for AI, and adding AI-specific optimizations). The benefit is higher speed and energy efficiency for neural networks. For example, tensor processing units can perform matrix multiplies extremely fast; neuromorphic chips (discussed above) aim for brain-like efficiency; and FPGAs offer reconfigurable logic for specialized tasks. Gartner sees such accelerators as crucial in building **"flexible, future-proof compute platforms for generative AI and other business needs" (Gartner Hype Cycle for Compute - Q.ANT)**. In essence, AI accelerators will continue to evolve so organizations can handle more advanced AI without prohibitive cost or power usage.

Risks & Challenges: While GPUs are firmly in productive use, the **cutting-edge AI chips** at the Peak face typical hype risks. Many AI accelerator startups are competing; some solutions may underperform or prove hard to program. Gartner cautions that "infrastructure leaders must balance hype with reality" – for example, adoption of FPGA or ASIC accelerators has been slower whenever software support is lacking or integration is complex. There is also the risk of **supply bottlenecks and costs**: 2024's surge in AI chip demand led to shortages, and companies investing heavily in AI hardware must ensure long-term ROI. Another concern is **power and cooling** – accelerators pack a lot of computation and can draw enormous power (GPUs in particular), straining data center energy and cooling capacities. This is why complementary innovations (like liquid cooling and photonics) are also on the Hype Cycle, aiming to mitigate these issues.

Overall, Gartner's view is that AI acceleration hardware is here to stay and will keep advancing. GPUs have reached a productive plateau (Introduction to the Gartner Hype Cycle – BMC Software | Blogs), and newer AI chips are climbing the curve. The key is to leverage their benefits (unprecedented AI capabilities) while minding the risks (cost, complexity, sustainability). Gartner's forecast even indicated that after an initial 2024 spike, AI chip demand might **"cool off"** to more sustainable growth (AI chip sales predicted to jump by a third – then cool off • The Register) – a sign that the hype is starting to give way to practical, optimized deployment strategies.

Advanced Cooling Technologies (Liquid & Immersion Cooling)

Hype Cycle Position: *Peak of Inflated Expectations* (Direct-to-Chip Liquid Cooling, 2024); *Trough of Disillusionment* (Immersion Cooling, 2024). Data center cooling is a critical hardware domain seeing innovation due to rising chip power densities. In the 2024 Data Center Infrastructure Hype Cycle, **Direct-to-Chip (D2C) liquid cooling** is at the Peak of hype (Bare metal can come back, says Gartner, citing VMware prices • The Register), meaning it's getting intense attention as a solution for high-heat processors. By contrast, **immersion cooling** – submerging servers in dielectric fluid – has slid into the Trough, as initial excitement gave way to practical challenges (Bare metal can come back, says Gartner, citing VMware prices • The Register).

Future Relevance & Benefits: Both liquid cooling approaches aim to more efficiently remove heat than traditional air cooling. Gartner considers them important for future data centers as CPU/GPU TDPs (thermal design powers) climb. **Direct-to-chip liquid cooling** uses cold plates or microfluidic channels to deliver coolant (water or refrigerant) directly to hot components (CPUs, GPUs). The benefit is highly targeted cooling, enabling higher rack densities and lower fan power. It can significantly improve PUE (power usage effectiveness) and support **sustainability goals** (part of the push for net-zero data centers) (Bare metal can come back, says Gartner, citing VMware prices • The Register). Many view D2C cooling as necessary for next-gen supercomputers and AI clusters, and Gartner rates it as a high-impact innovation (albeit amidst hype).

Immersion cooling involves placing entire servers in baths of non-conductive liquid. Single-phase immersion uses a pumped coolant bath, while two-phase uses fluid that boils off the heat. The allure is even greater cooling efficiency – potentially eliminating the need for any air cooling in the data center. It can dramatically reduce server hotspot issues and even reuse waste heat. Gartner did highlight immersion cooling in past cycles, and its presence (even in the trough) indicates it's a serious trend for specialized deployments (such as crypto mining farms, some HPC installations, etc.).

Risks & Challenges: At the peak of hype, direct liquid cooling may face overexpectations. Gartner notes that while many organizations are testing liquid cooling, broad **mainstream adoption may take 2–5+ years** as operational best practices mature. Challenges include the upfront retrofit costs, risk of leaks, and the need for plumbing infrastructure in data centers. For **immersion cooling**, the tumble to the trough reflects real hurdles: fluid costs, equipment maintenance (e.g. removing servers from fluid for service), and vendor ecosystem fragmentation. Gartner observes that immersion cooling, along with other hyped infrastructure tech like composable hardware, **has disappointed some early adopters and is "sliding" down in expectations (Bare metal can come back, says Gartner, citing VMware prices • The Register)**.

Another risk aspect is organizational comfort – enterprise data center teams may be hesitant to adopt liquid-based cooling without strong incentives, sticking to air cooling which is well understood. There are also **standards and interoperability questions** (each cooling vendor might have proprietary solutions). Gartner's inclusion of these cooling methods in the Hype Cycle, however, underscores that the industry must innovate in cooling to keep up with compute density. We can expect liquid cooling to gradually climb the "Slope of Enlightenment" as techniques improve and proof of value is demonstrated (for example, by showing energy cost savings or enabling otherwise impossible high-performance deployments). In summary, **Gartner's take**:

advanced cooling is hyped in 2024 – direct liquid cooling at a frenzy of interest (Bare metal can come back, says Gartner, citing VMware prices • The Register), immersion cooling in a bit of a backlash (Bare metal can come back, says Gartner, citing VMware prices • The Register) – but both are key to future data center designs as the heat output of hardware continues to soar.

Compute Architecture & Infrastructure Trends

Edge Computing

Hype Cycle Position: *Trough of Disillusionment* (Data Center Infrastructure 2024). After years of buzz, edge computing has encountered a reality check. Gartner's 2024 Hype Cycle for Data Center Tech assesses **edge computing as "in trouble," sliding into the trough of disillusionment** (Bare metal can come back, says Gartner, citing VMware prices • The Register). In practice, this means the initial hype – that placing compute power ubiquitously at the network edge would revolutionize latency and bandwidth for all applications – has outpaced real deployments. By 2024, many edge projects have struggled to scale or show return on investment, causing a dip in enthusiasm.

It's important to clarify what "edge computing" encompasses: it refers to decentralized compute infrastructure located closer to end-users or devices (e.g. micro data centers at cell towers, on-premises gateway servers, even on-device AI chips), rather than in central clouds. A few years ago, edge was at the Peak of Inflated Expectations with predictions of **massive IoT and 5G-driven demand**. Now, Gartner's data suggests only niche adoption so far, hence the disillusionment phase (Bare metal can come back, says Gartner, citing VMware prices • The Register).

Future Relevance & Benefits: Despite the current trough, Gartner does not dismiss edge computing's long-term importance. The core benefit of edge architecture is **ultra-low latency and local processing** of data, which is crucial for use cases like autonomous systems, industrial IoT control loops, AR/VR, smart cities, and real-time analytics. Offloading computation to the edge can also reduce backhaul costs and improve reliability when connectivity to cloud is limited. Gartner's own forecast (from other research) often cites that by mid-decade a significant portion of new data will be generated and processed outside traditional centralized data centers – a nod to the edge's rising role. Indeed, the presence of **"Edge servers"** on the 2024 Compute Hype Cycle (in the trough region) confirms that edge infrastructure is evolving, just slower than the hype predicted.

As the hype cycle implies, after the trough comes the **Slope of Enlightenment**. We can expect that through 2025 and beyond, successful patterns for edge computing will emerge. Likely, specific sectors (e.g. telecom with MEC - Multi-access Edge Computing, or manufacturing with on-prem AI) will lead the way. Gartner's commentary suggests that the concept is sound but **needed more time to find pragmatic applications** and scalable management models.

Risks & Challenges: The fall into the trough was driven by several challenges. One is the **complexity of managing distributed infrastructure** – deploying and orchestrating potentially

thousands of mini-sites is non-trivial (this fuels interest in orchestration, containerization, and automation tools). Another issue is that some anticipated killer apps for edge (like fully autonomous vehicles or ubiquitous AR glasses) are still in development, thus the demand pull hasn't fully materialized. Gartner specifically noted that edge computing has **"failed to deliver on its promise" for now (Bare metal can come back, says Gartner, citing VMware prices • The Register)**, which encapsulates the risk: early projects overestimated benefits or underestimated costs.

There is also a **scalability and standardization problem** – different use cases require different edge designs, and there isn't a one-size-fits-all edge stack yet. Additionally, ensuring security across many edge nodes can be difficult, and without clear ROI, business leaders may be hesitant to invest heavily in edge deployments.

Going forward, the **outlook** is that edge computing will climb out of the trough gradually. Gartner's Hype Cycle logic suggests that as real successes emerge (for example, content delivery networks and gaming have effectively used edge caching; retail companies using edge AI for inventory and analytics; telecom operators offering edge cloud services), confidence in the edge will grow again. In summary, **Gartner's perspective** is cautious optimism: edge computing is currently over-hyped and under-delivering, but it addresses genuine needs and is likely to achieve mainstream productivity in a 2–5+ year timeframe once technologies and use cases mature.

Cloud-Native Computing & Platform Engineering

Hype Cycle Position: *Peak of Inflated Expectations* (Infrastructure Platforms 2024; Emerging Tech 2024 theme). "Cloud-native" computing – characterized by containerization, microservices, DevOps, and dynamic orchestration (Kubernetes) – has become a dominant IT architecture. By 2024, many cloud-native techniques are actually quite mature (containers themselves are approaching the Plateau). However, Gartner's research highlights **a new wave of hype around internal developer platforms and cloud-native enablement**, often termed **Infrastructure Platform Engineering (IPE)** or platform engineering. In Gartner's *Infrastructure Platforms Hype Cycle 2024*, **platform engineering is at the Peak** and deemed a *"Transformational"* technology with 5–10 years to mainstream adoption (Our Key takeaways from the 2024 Gartner® Hype Cycle™ for Infrastructure Platforms, 2024 | Kubermatic). Similarly, the *Emerging Technologies Hype Cycle 2024* lists **cloud-native computing and related practices (GitOps, WebAssembly, internal developer portals)** under the "Developer Productivity" theme, indicating they are must-watch emerging tech for improving software delivery (Gartner 2024 Hype Cycle for Emerging Technologies highlights developer productivity, total experience, AI and security | EngineerIT).

What this means: Enterprises have widely adopted cloud (public and private) and are now seeking to optimize how developers utilize cloud-native infrastructure. Gartner notes a trend toward **self-service internal platforms** that abstract away complexity for developers (Our Key takeaways from the 2024 Gartner® Hype Cycle™ for Infrastructure Platforms, 2024 | Kubermatic). These include curated Kubernetes-based platforms, automated pipelines, and "golden path" development frameworks. By empowering developers and standardizing environments, organizations aim to ship software faster and more reliably. Gartner predicts that **over 50% of I&O organizations will form platform engineering teams by 2028 (up from <25% in 2024) (Our Key takeaways from**

the 2024 Gartner® Hype Cycle™ for Infrastructure Platforms, 2024 | Kubermatic), underscoring the strategic importance of this shift.

At the same time, core cloud-native technologies like containers and microservices are moving into the Slope of Enlightenment. The Compute 2024 Hype Cycle actually shows **"OS Containers" and "Micro OS for Containers" on the Slope** nearing the Plateau, which suggests containerization itself is no longer hype but accepted practice. What's emerging/hyped are improvements around it – e.g. **eBPF** (for better networking/observability inside clouds) is at peak hype, and **immutable infrastructure patterns** (treating servers as replaceable artifacts) are on the rise to mainstream (Bare metal can come back, says Gartner, citing VMware prices • The Register). All these tie into the cloud-native paradigm of treating infrastructure as code and automating everything.

Future Relevance & Benefits: Gartner's perspective is that cloud-native architectures are foundational for modern IT agility. Benefits include **scalability, resilience through microservice isolation, and faster innovation via CI/CD**. The current hype around platform engineering specifically is about boosting **developer productivity** – making cloud-native tech more accessible so developers can focus on code, not config. Gartner calls out technologies like **AI-augmented software engineering, GitOps, and Internal Developer Portals as enablers of developer "flow" state** (Gartner 2024 Hype Cycle for Emerging Technologies highlights developer productivity, total experience, AI and security | EngineerIT). By institutionalizing these, organizations can deliver applications quicker and with higher quality.

Cloud-native also underpins other trends (it's hard to do edge or AI at scale without cloud-native infrastructure behind the scenes). We see **complementary Gartner trends**: e.g., "Boosting developer productivity" was one of the four core themes for emerging tech 2024 (Gartner 2024 Hype Cycle for Emerging Technologies highlights developer productivity, total experience, AI and security | EngineerIT). Thus, cloud-native computing isn't a single technology but an ecosystem of tools and practices that together are reaching a high level of importance.

Risks & Challenges: At the peak of inflated expectations, there is a risk of **over-promising what platform engineering can do**. Gartner warns that simply creating a platform team doesn't automatically solve organizational silos or developer frustration. There are cultural and process challenges – e.g., balancing standardization with developer autonomy. If done poorly, internal platforms can become another source of friction (the very problem they intend to solve). **Governance and complexity** are concerns: running Kubernetes and a multitude of cloud-native tools at enterprise scale is complex, and requires new skills (hence the emphasis on platform engineers).

Another risk is hype divergence: some executives may think "cloud-native" is a solved problem since containers are everywhere, underestimating the effort required to actually optimize it for productivity (the reason IPE is hyped is because many organizations haven't gotten it right yet). Security in cloud-native environments (devsecops) is another ongoing challenge that can derail benefits if not addressed (this ties into Gartner's "Digital immune system" and AI TRiSM trends under security (Gartner 2024 Hype Cycle for Emerging Technologies highlights developer productivity, total experience, AI and security | EngineerIT) (Gartner 2024 Hype Cycle for

Emerging Technologies highlights developer productivity, total experience, AI and security | EngineerIT)).

In summary, **Gartner's stance** is that cloud-native computing is moving from an ad-hoc adoption phase into an engineered, productized phase within companies. It's at a hype peak in terms of tooling and methodologies to streamline it (Our Key takeaways from the 2024 Gartner® Hype Cycle™ for Infrastructure Platforms, 2024 | Kubermatic). Done well, the result will be faster software cycles, more resilient systems, and happier developers. The academic takeaway for students is that cloud-native principles (containerization, orchestration, serverless, etc.) are now essential knowledge – and the frontier is how to effectively operationalize these at scale (platform engineering) while avoiding the common pitfalls (overcomplexity, lack of developer buy-in). Gartner expects significant progress in the next 5 years, moving this domain to full **Plateau of Productivity** once best practices solidify.

Composable & Disaggregated Infrastructure

Hype Cycle Position: *Trough of Disillusionment* (Data Center 2024). Composable infrastructure refers to the ability to flexibly assemble and reassemble compute, storage, and network resources on the fly (often via a software control plane), rather than having fixed physical servers. It promises cloud-like agility with on-premises hardware by disaggregating resources into pools. In Gartner's 2024 Hype Cycle for Data Center Infrastructure, **Composable Infrastructure has sunk into the Trough of Disillusionment** (Bare metal can come back, says Gartner, citing VMware prices • The Register). This indicates that while the idea has been around for a few years, real-world adoption has been slower than initially hoped, and some early efforts haven't met expectations.

A related concept trending in 2024 is **"Devirtualization"** – essentially the opposite of hyper-converging, where certain workloads are moved *off* virtual machines to bare metal for performance or cost reasons. Gartner introduced *devirtualization* as a new item in the 2024 Hype Cycle, driven by factors like hypervisor licensing costs (Bare metal can come back, says Gartner, citing VMware prices • The Register) (Bare metal can come back, says Gartner, citing VMware prices • The Register). One can see this as part of a broader rethinking of infrastructure composition: organizations are re-evaluating the trade-offs of traditional virtualization versus newer approaches (bare metal plus containers, or disaggregated racks with pooled components).

Future Relevance & Benefits: The vision of composable infrastructure is quite powerful. By decoupling CPU/Memory, storage, and accelerators into separate sleds or chassis connected by high-speed fabrics (like Compute Express Link, which Gartner shows on the rise), data centers could allocate resources on-demand to applications. For example, if an AI workload suddenly needs more GPU or memory, a composable system could attach more of those resources dynamically. This would lead to **much higher utilization efficiency** (no more stranded resources trapped in one physical server) and **flexibility** in deploying varied workloads. Gartner classifies composable/disaggregated infrastructure as having a potentially **High transformational benefit** if realized, as it essentially brings cloud-model elasticity to on-prem hardware.

Already, we see enabling technologies: **CXL (Compute Express Link)** is a new interconnect standard that Gartner listed in the Innovation Trigger, aimed at enabling memory pooling and

resource disaggregation. Also, **hyperconverged infrastructure (HCI)** – while different (it tightly couples compute+storage) – reached the Plateau in the Hype Cycle, so the market is used to abstracting resources to some degree. Composable infrastructure can be seen as an evolutionary next step beyond HCI, promising even more fluid resource management.

Risks & Challenges: The fact that composable infrastructure is in the trough means early experiments have exposed difficulties. Gartner's analysis suggests that the technology **"has not yet lived up to its hype, disappointing buyers" (Bare metal can come back, says Gartner, citing VMware prices • The Register)**. Challenges include: a lack of standards (early solutions are often vendor-specific and not interoperable), high costs to implement new fabric hardware, and limited support in existing software stacks. For instance, truly disaggregating memory and accelerators across a rack requires ultra-low latency networks and changes to how operating systems handle remote resources. The industry is only partway there (CXL is one step, but widespread adoption and OS support will take a few years).

Moreover, Gartner mentions **"Revirtualization"** – moving to new virtualization or container platforms – at the Peak of Expectations in 2024 as a response to vendor licensing issues (Bare metal can come back, says Gartner, citing VMware prices • The Register). This churn shows that enterprises are currently grappling with how best to configure their infrastructure: some swing from virtualized to bare metal (devirtualize) for cost/performance, others invest in new orchestration to recombine workloads (revirtualize on different tech). In this turbulent period, composable infrastructure might be a hard sell until it proves clear value over these simpler approaches.

Gartner's guidance emphasizes careful consideration of maturity. The trough placement indicates composable hardware is **5–10 years from mainstream**. In the interim, organizations might achieve some similar goals with existing tech (e.g., using Kubernetes to treat a cluster of nodes somewhat fungibly, or employing HCI for simplicity). There is also a human/process element: adopting composability might require rethinking capacity planning and operational processes, which can encounter resistance or misalignment.

In academia, the notion of disaggregated infrastructure ties to research in computer architecture and systems (think of projects on rack-scale computing, or memory disaggregation). The Hype Cycle teaches us that while the *concept* is very promising, *implementation maturity* lags. **Bottom line:** Gartner still considers composable infrastructure strategically important, but due to early over-hype, it's in a phase of reassessment. As underlying technologies (like faster interconnects and better orchestration software) improve, composability could climb into enlightenment. Students should watch this space – e.g., developments in CXL, smart NICs (DPUs), and software-defined infrastructures – as they will likely drive the resurgence of composable infrastructure toward the end of the decade.

Summary of Technologies, Hype Stages, and Maturity (2024–2025)

The table below summarizes each technology discussed, the Hype Cycle(s) in which it appears, and its current stage (as of 2024), along with Gartner's estimated timeline to maturity:

Technology	Hype Cycle (Year)	Position (Stage)	Time to Plateau
Quantum Computing	Emerging Tech 2024; Compute 2024	Peak of Inflated Expectations	10+ years (long-term) ([Gartner's Hype Cycle report suggests photonic computing is gaining traction
Neuromorphic Computing	AI 2024; Compute 2024	In/near Trough of Disillusionment ([Hype Cycle for Artificial Intelligence, 2024	annotated by SHIH YEN](https://readwise.io/reader/shared/01j346cwkhtw8xqbazpgy6mx99/#:~:text=Soon%20to%20leave%20the%20peak,rest%20of%20the%20Hype%20Cycle))
Photonic Computing	Compute 2024	Innovation Trigger (embryonic) ([Gartner's Hype Cycle report suggests photonic computing is gaining traction	Electro Optics](https://www.electrooptics.com/article/gartners-hype-cycle-report-suggests-photonic-computing-gaining-traction#:~:text=,Hype%20Cycle%E2%84%A2%20for%20Compute%2C%202024))
AI Accelerators (ASICs)	Compute 2024	Peak of Inflated Expectations	2–5 years (for ASICs); GPUs at Plateau ([Introduction to the Gartner Hype Cycle – BMC Software
GPU Acceleration	(Mature technology)	Plateau of Productivity ([Introduction to the Gartner Hype Cycle – BMC Software	Blogs](https://www.bmc.com/blogs/gartner-hype-cycle/#:~:text=Examples%20of%20the%20Plateau%20of,Different%20technology%2C%20different))
Edge Computing	DC Infra 2024	Trough of Disillusionment (Bare metal can come back, says Gartner, citing VMware prices • The Register)	2–5 years (to climb slope)
Cloud-Native / Platform Eng.	Infra Platforms 2024; Emerging Tech 2024	Peak of Inflated Expectations ([Our Key takeaways from the 2024 Gartner® Hype Cycle™ for Infrastructure Platforms, 2024	Kubermatic](https://www.kubermatic.com/blog/our-key-takeaways-from-the-2024-gartner-hype-cycle-for-infrastructure-platforms/#:~:text=Infrastructure%20Platform%20Engineering%20,the%20%E2%80%9CPeak%20of%20Inflated%20Expectations%E2%80%9D))
Composable Infrastructure	DC Infra 2024; Compute 2024	Trough of Disillusionment (Bare metal can come back, says Gartner, citing VMware prices • The Register)	5–10 years (needs maturation)
Direct Chip Liquid Cooling	DC Infra 2024	Peak of Inflated Expectations (Bare metal can come back, says Gartner, citing VMware prices • The Register)	2–5 years (early adoption)
Immersion Cooling	DC Infra 2024	Trough of Disillusionment (Bare metal can come back, says Gartner, citing VMware prices • The Register)	~5+ years (niche use now)

(Sources: Gartner Hype Cycle reports 2024; positions interpreted from text and charts (Bare metal can come back, says Gartner, citing VMware prices • The Register) (Bare metal can come back, says Gartner, citing VMware prices • The Register) (Hype Cycle for Artificial Intelligence, 2024 |

annotated by SHIH YEN) (Gartner's Hype Cycle report suggests photonic computing is gaining traction | Electro Optics) (Introduction to the Gartner Hype Cycle – BMC Software | Blogs).)

Conclusion: The Gartner Hype Cycles for 2024 reveal that many cutting-edge computing technologies are in the early or intermediate stages of maturity. **Hardware innovations** like quantum, neuromorphic, and photonic computing are mostly embryonic or peaking in hype – their transformative potential is recognized, but significant R&D and time are needed before they become practical. Current workhorse technologies such as GPUs are already at plateau, delivering real value, while new AI accelerators ride a hype wave that must eventually stabilize into tangible improvements. In the realm of **compute architectures**, trends like edge computing and composable infrastructure have faced setbacks after initial excitement, showing the importance of realistic goals and incremental progress. Conversely, cloud-native architectures have largely proven their worth, and the focus has shifted to optimizing their use (through platform engineering and automation) at scale.

From Gartner's perspective, each technology carries unique benefits (whether it's unprecedented speed, efficiency, flexibility, or productivity gains) as well as risks (technical hurdles, integration complexity, skill gaps, or simply being ahead of its time). For a university-level AI and Future Computing course, the key takeaway is to understand *where* each technology lies on its adoption curve and *why*. This helps predict the trajectory: which innovations might become dominant in a few years versus which require longer-term breakthroughs or ecosystem development. Gartner's Hype Cycle provides a framework to critically assess emerging tech – to separate the hype from the genuine drivers of progress (Itential Recognized in Six Gartner® Hype Cycle™ Reports for its Innovative Infrastructure Orchestration & Automation Technology for the Second Year in a Row) – and to strategize accordingly. In summary, the coming years (2025 and beyond) will likely see the climbing of some of these technologies out of the trough (e.g. more edge computing successes, initial quantum applications, etc.), the plateauing of others (containers and cloud-native methods becoming standard IT plumbing), and the entry of entirely new innovations. Keeping an eye on the Hype Cycle stages helps academics and practitioners alike to navigate this evolving landscape of future computing.