THE HEAT
IS ON

Why thermal management is crucial to NVIDIA in data center, fiber, and telecom performance

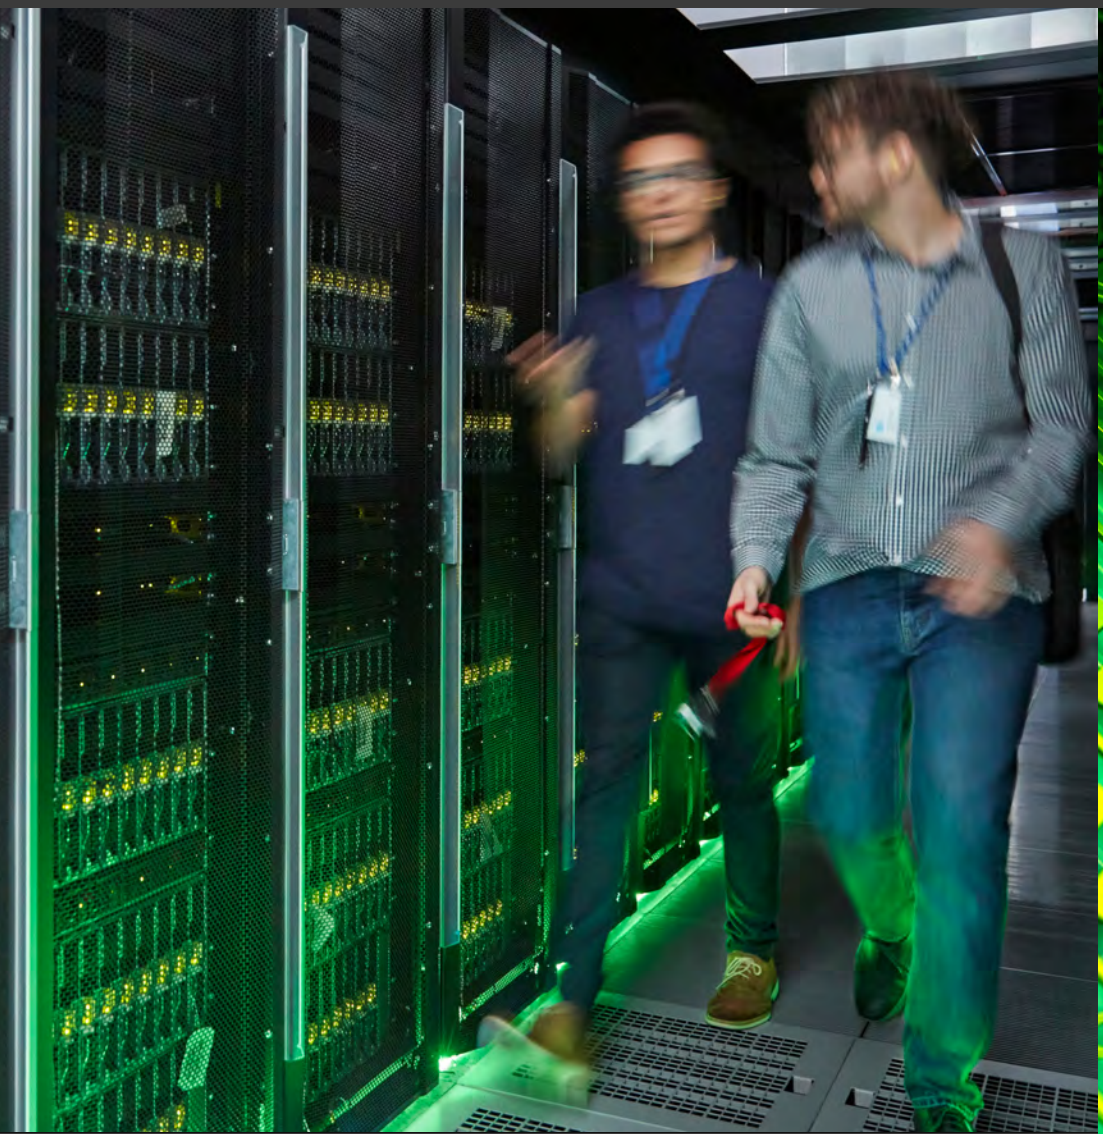

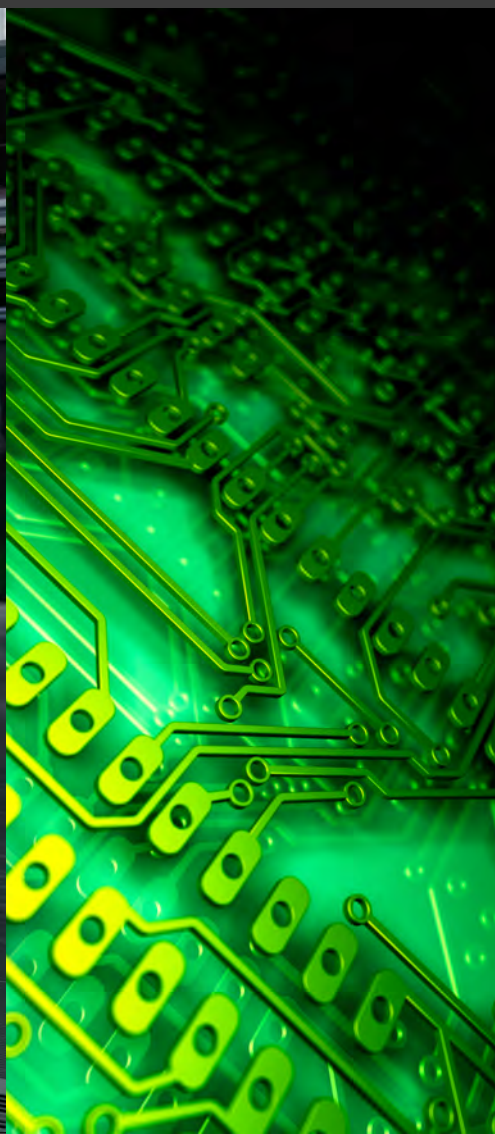

Contents

Executive summary

As engineers throughout NVIDIA know very well, thermal management—a small, silent, vital enabler to network infrastructure—significantly affects network operational performance. While thermal management may not be the focal point of today’s advancements, it’s intrinsic to the future success of areas like AI, VR, and 5G. That’s because electronic components get hot—a problem exacerbated by heightened network demands for improved reliability, higher power density, and speed.

Thermal management has been limited by conventional thermal interface materials (TIMs) and traditional active-cooling methods. NVIDIA has made great strides in advances in thermal management already, but additional results can be achieved by using effective new thermal management approaches in hardware components such as circuit boards in active antenna units (AAUs) and line cards for routers and switches.

Thermal management improvements can enable NVIDIA to help data centers and telecom operators to reach their target network performance metrics and to enable emerging technologies such as **AI, 5G, Wi-Fi 6, and supporting 400 GbE speeds.**

Perhaps it seems counterintuitive to start optimization at the integrated circuit level. However, data center and telecom infrastructures use components in mass quantities. One small uptick in thermal management at the integrated circuit level can create colossal, in-aggregate improvements, and NVIDIA can help to catalyze that.

Few materials master heat dissipation, but all networks need it. In router, switch, and circuit board design, effective thermal management is a notable competitive differentiator. Advanced thermal management materials include thermal gels, phase change materials, thermal GAP PAD® materials, and thin-film thermally conductive, dielectric coatings.

This paper explains how NVIDIA can help telecom operators and data center technologists elevate network outcomes by improving thermal management at the electronic component level—and why it matters.


Henkel - NVIDIA - “The heat is on”

A MASSIVE LEVER

The heat is on for data center and telecom operations. Global demand for data, internet access, and bandwidth is skyrocketing. Network infrastructure spending has escalated to respond to the worldwide need.

One vital linchpin holds the key to enabling expansion: thermal management. Networks need to operate integrated circuits at maximum processing power without overheating. As speed and component densities increase, however, so does the heat generated within those integrated circuits. Although active cooling assists with heat dissipation, it is expensive—and it is reaching its limit.

The rise of AI is driving the need for more and more speed and bandwidth. New technology innovations that NVIDIA is helping to move forward—among them, AI-focused data centers, 5G, Wi-Fi 6, and 400 GbE data transmission rates—are enabling dramatically increased processing speeds to meet this demand and further elevate the necessity of thermal management in each component.

Data centers rely on enormous scale and size to provide massive computing power—and the high computing power demand of AI will continue to push the demands. For example, most large data centers have around **100,000 servers**; hyperscale data centers can have many more.

The limiting factor to network size and operations is either power or cooling, both of which also represent a significant percentage of operating expense (OpEx).

It makes sense to minimize the aggregate heat generated by all that hardware. Thermal management is the gateway to future expansion, in part because of its domino effect on network performance. When it’s successful, thermal management helps to maximize network performance, boost reliability, and enhance component life. This is a key area of opportunity for NVIDIA.

As a result, NVIDIA and other electronic component manufacturers are motivated to improve thermal management at the source: the network infrastructure level within integrated circuits.

Akin to a powerful lever, the multiplier effect of a small heat dissipation improvement in NVIDIA circuit boards for routers, switches, servers, and AAUs adds up to big outcomes. Thus, thermal management materials help to minimize the aggregate heat generated at the network infrastructure level. Among them: circuit boards, line cards for routers and switches, and printed circuit boards (PCBs). The PCBs often are used in AAUs and baseband units (BBUs, sometimes called active band units or ABUs).

As electronic devices become more powerful, they generate more heat. Thermal management materials play a pivotal role in network performance and mitigate risks of system reliability drops, product failures, and degradation over time. The good news for NVIDIA is that it is possible to address these challenges with these exceptional thermal management materials—such as thermal gels, GAP PAD® materials, phase change materials, and microTIMs—that are used in the construction of the micro-sized electronic components.

Challenged by heightening demands for increasing reliability, density, processing power, and speed, thermal management is the pathway to data center and telecom performance, enabling 5G, Wi-Fi 6, 400 GbE data transmission rates, and other emerging innovations.

Why is heat a problem?

Heat accelerates a component’s performance degradation and reduces its lifespan.

Electronic components like those from NVIDIA are composed of semiconducting, conducting, and non-conducting materials. For any of them, the general rule of thumb is that the speed of chemical reactions doubles for each increase of 10° C. Electronic components need to be maintained at stable temperatures; otherwise, chemical reactions can break down or alter their materials.

Physical degradation can also occur. Thermal or power cycling can warp components when overheating occurs, such as when they run at maximum temperatures for extended periods of time.

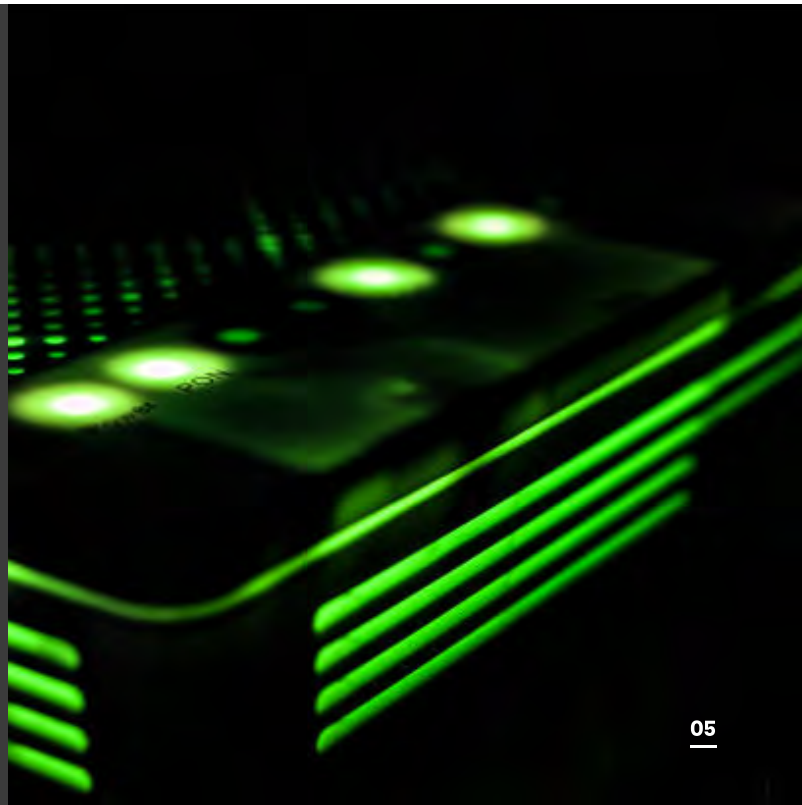

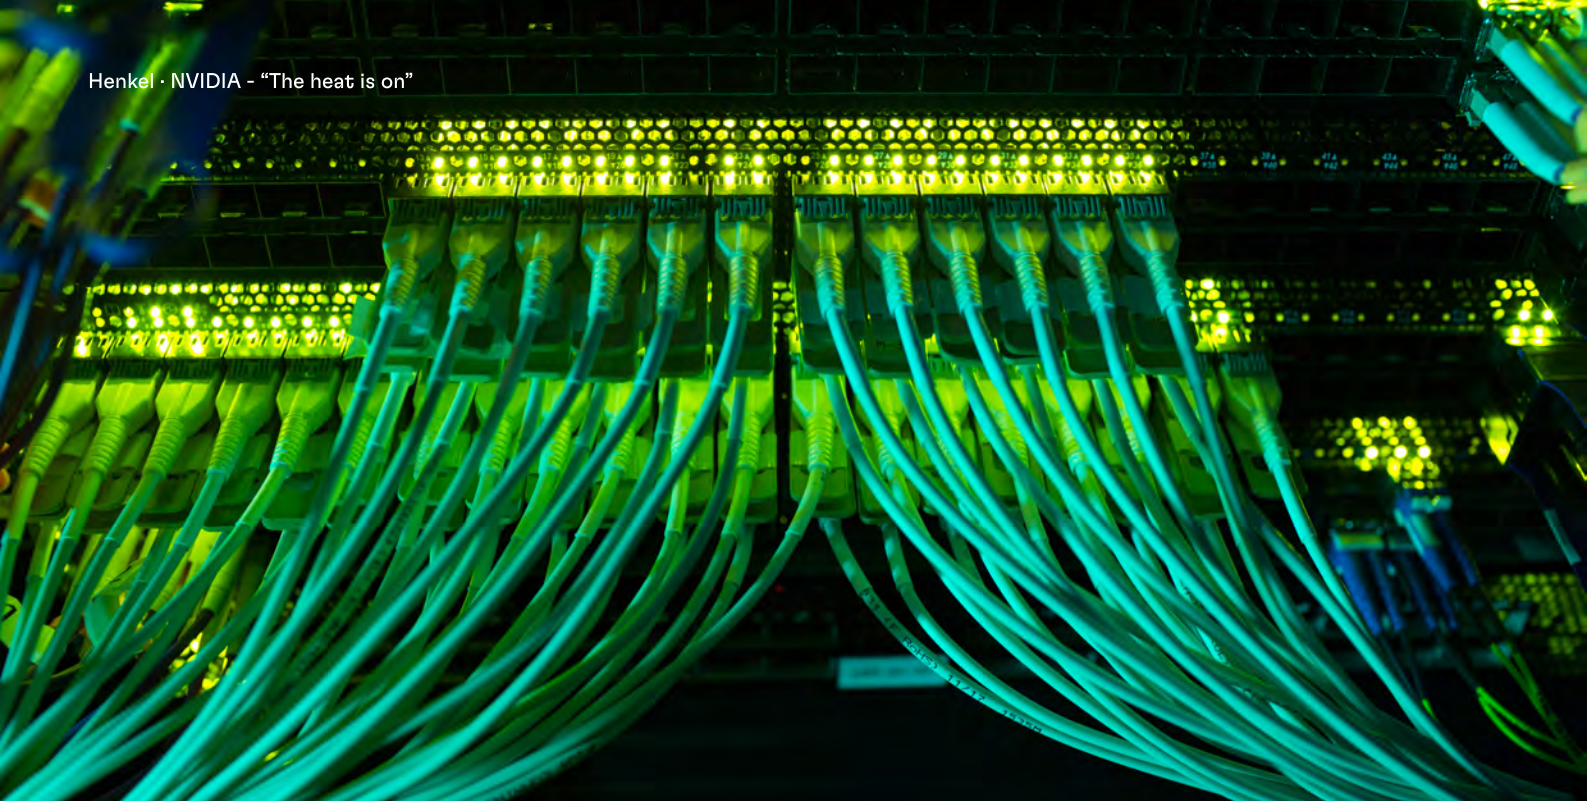
Henkel · NVIDIA - “The heat is on”

TURNING UP THE **HEAT**

This isn’t a minor matter. Data center and telecom trends underscore thermal management’s crucial role.

Data center

Cloud hyperscale data centers and their colossal computing power capabilities are working to meet the rising demands for data volume and speed. In spite of economic turmoil, the world’s largest data center operators— Google, Amazon, Meta, and Microsoft—are leading the way in data center spending that’s expected to reach **$222 billion** in 2023. And the demand keeps rising. According to a 2023 Gitnux report, **“global internet traffic is expected to grow at a Compound Annual Growth Rate (CAGR) of 24% between 2021 and 2026,”** propelling investment and expenditures in data center networking to epic levels.

Data center OpEx, however, has reached a critical juncture, primarily due to cooling and power costs. According to Cushman & Wakefield’s 2023 Global Data Center Market Comparison, data center power demand increased from 4.9 GW to an estimated **7.4 GW** in one year. With an average of **40%** of data center power going to cooling costs, thermal management has to be top of mind. They are seeking solutions that NVIDIA can offer with the right thermal management approaches.

The takeaway: Data center OpEx threatens profitability. Given that thermal management can help maximize data center output while minimizing OpEx, it’s worth investigating.

Telecom and fiber-optics

These changes aren’t only in the data center. Concurrently, telecom infrastructure is growing rapidly. The telecom tower market was valued at **$56.94 billion** in 2021 and is projected to reach $189.05 billion by 2030. To fuel this growth, telecom operators are focused on controlling both capital expenditures (CapEx) and OpEx. According to a PwC survey, telecom operators indicate that **up to 20%, or $65 billion, per year in CapEx is wasted** due to under-optimized CapEx.

OpEx is also an ongoing challenge for telecom and fiber-optics infrastructure. Mobile base stations and cell towers are always-on equipment, and they generate rising levels of heat. Temperature and humidity fluctuations in outdoor environments are further burdens on remote telecom equipment that is equipped with limited active-cooling options.

Heat generated on the printed circuit boards used in base stations and cell towers can degrade equipment performance and reduce lifecycles. Because it costs so much to access and repair the technology, telecom and fiber operators need electronic components with maximum longevity and reliability. The operational target is near-zero failures and minimal maintenance. NVIDIA telecom and fiber components can help achieve this goal by relying heavily on effective thermal management, using advanced materials applied to the printed circuit boards that power AAUs and BBUs.

According to telecom operators, up to 20%—or $65 billion—is wasted per year in CapEx.

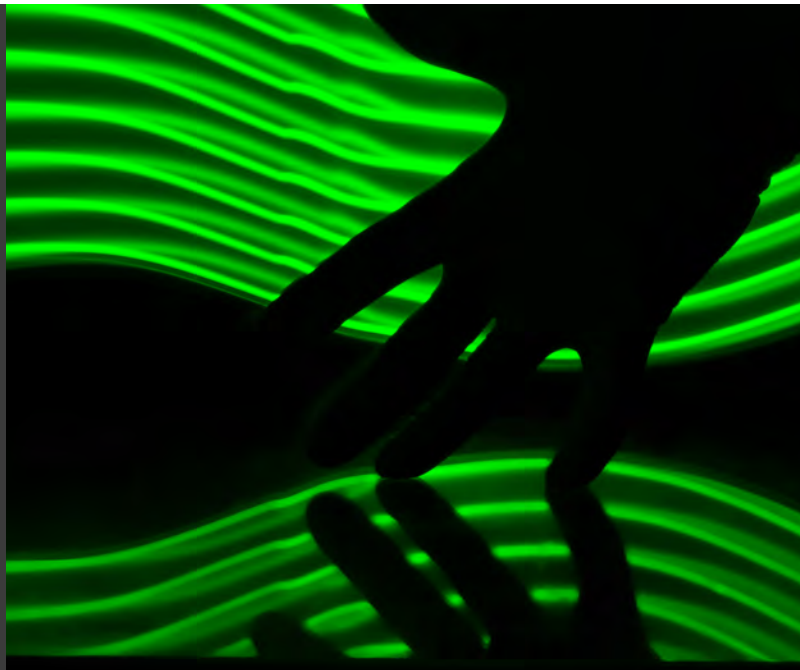

Emerging network enhancements

Everything is moving faster. 5G, Wi-Fi 6, and 400 GbE data transmission rates increase the need for effective thermal management in integrated circuits, and the bandwidth demands of AI and metaverse applications will amplify this need. Electronic components like those from NVIDIA must respond to speed and component requirements and all the other demands that enable our global always-on connectivity.

10x

5G has 10 times faster data processing speeds than 4G

9.6 Gbps

Wi-Fi 6 has the capacity for 9.6 Gbps, compared to 3.5 Gbps for Wi-Fi 5

4x

400 GbE data transmission rates are four times the speed of today’s 100 GbE

To keep up, thermal management capabilities must increase concurrently.

THERMAL MANAGEMENT **AT SCALE**

The higher processing density on integrated circuits generates more heat. Active-cooling options do help, but they are expensive, they are reaching their physical limitations, and they have adverse sustainability implications. Given the colossal scale of network infrastructure, thermal management is a matter of intense significance.

Data center

Hyperscale data centers enabled an unprecedented evolution of cloud computing. In 2022, an estimated **57% of businesses migrated their workloads to the cloud** across all industries. To keep pace with this demand and the heat dissipation it requires, the data center cooling market for active cooling is expected to grow at a compound annual growth rate (CAGR) of **17.1% from 2023 to 2030** as the public cloud computing market balloons to **$1 trillion by 2026**. Advanced materials help tackle the operational challenge of heat dissipation with small, incremental improvements in thermal management. For a participant as central to data centers as NVIDIA, this adds up to a substantial outcome.the operational challenge of heat dissipation with small, incremental improvements in thermal management. For a participant as central to data centers as NVIDIA, this adds up to a substantial outcome.

Another factor impacting thermal management is higher processing densities. Not long ago, racks with line cards in data centers handled around 100 GbE per pluggable optical module (POM). Since then, this number has quadrupled to 400 GbE speeds without any increase in rack size. And now it’s heading toward **800 GbE**. More processing is required per board, and that elevates the heat generated in the electronic components that NVIDIA produces.

Aggregate impact: What 5° C means in data centers

Can 5° C really make a difference? With the transition to 400 GbE-capable modules, the power level per POM—which number as many as 32 per line card—can reach as high as 15 watts. Alternatively, advanced, innovative microTIMs enable more heat to dissipate from the module, which reduces operational temperature at a rate of 0.33° C per watt. For a 15-watt module, temperature reduction is upwards of 5° C, which is significant in aggregate across the line card.

Telecom

Telecom operators face a growing energy challenge, too, which puts more emphasis on innovative thermal management. According to industry estimates, **as 5G replaces 4G, each telecom site will require two to three times more power**. On average, energy costs account for 5-7% of OpEx. These rising costs squeeze margins and limit the growth potential of on-site active-cooling options, making thermal management a vital determinant for effective telecom operations.

> **As 5G replaces 4G, each telecom site will require two to three times more power.**

Electronic components are susceptible to constant inoperation stress. They also must contend with outdoorrelated moisture and temperature fluctuations, which can lead to corrosion. Componentry is small, delicate, and temperature sensitive. It’s important to avoid overheating printed circuit boards within telecom’s AAUs and BBUs to sustain reliable operation and equipment life. 5G complicates this challenge; its higher switching and routing speeds increase heat generation at higher power densities.

Henkel · NVIDIA - “The heat is on”

INNOVATIONS IN ADVANCED **MATERIALS**

Advanced materials—such as thermal gels, thermal GAP PAD® materials, phase change materials, and microTIMs—can make a huge difference when used in the micro-sized electronic components produced by NVIDIA.

It all starts with the circuit board, the primary source of heat generation. Circuit boards may use several thermal management materials, which in turn power a router, switch, server, AAU, or BBU. To better understand how to use thermal management, let’s look at each of these robust heat-dissipating materials.

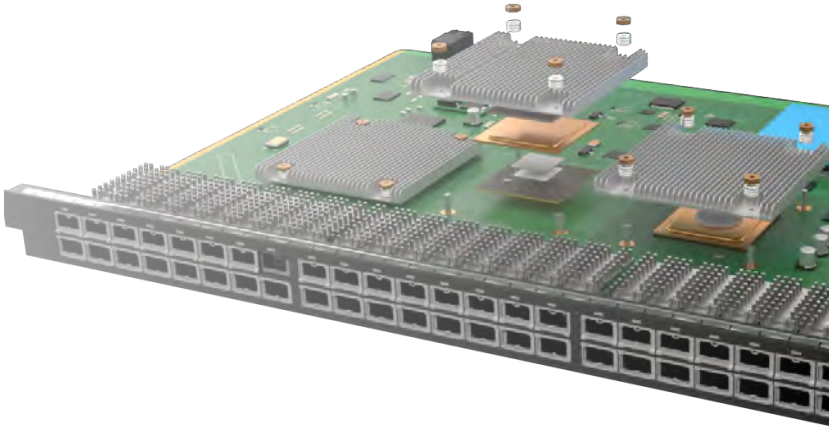

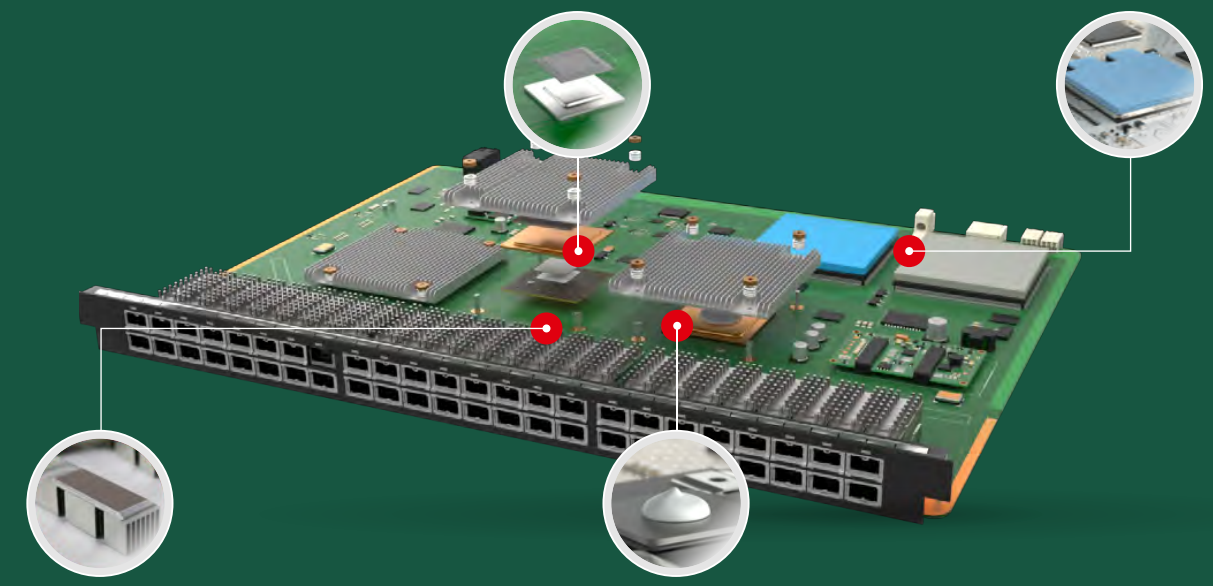

Caption: A typical circuit board using advanced materials.

Phase change materials

Phase change materials replace grease as the interface material between power devices and heat sinks. Thermal phase change compounds can form a thin bondline with low thermal resistance. Application is also simple, so the material can be integrated into fully automated processes, which supports mass production.

Micro thermal interface coatings

MicroTIM is a durable, thermally conductive, thin film coating applied to networking line card heat sinks that come in contact with POMs. MicroTIMs enhance the thermal performance and durability of heat-generating devices—including POMs—and aluminum heat sink assemblies. MicroTIMs also provide additional resilience to POMs’ repeated plug-and-pull action, and they help lower networking line card heat.

Thermal GAP PAD® materials

Thermal GAP PAD® materials eliminate air space between hot components and heat sinks to manage the challenging thermal loads inherent with higher power densities. These are soft, high-compliance materials that conform to irregular surface shapes. They fill small gaps to enable interface wet-out and optimize thermal transfer. Thermal GAP PAD® materials minimize the adverse impacts of heat on device operation and lifetime across data center and telecom applications. Newer GAP PAD® materials, which help to address heat from 5G infrastructures, are rated at 12.0 W/m-K.

Thermal gels

Thermally conductive gels are conformable. They fill gaps across many different types of assemblies, where they provide low component stress, simple rework, process flexibility, in-application stability, and efficient thermal conductivity. Thermal gels often are used for stationary applications that require stability, a common requirement in the telecom sector. They are available in one-component liquid and curable formulations.

NETWORK INFRASTRUCTURE OUTCOMES

Advanced materials help NVIDIA with thermal management at the electronic component level to drive improved outcomes at the data center, telecom, and fiber network levels.

The results	The root cause adressed
Efficiency	Better thermal management enables maximum operating power and maximum processing.
Less data latency	More heat dissipation lets facilities use higher power densities. That means more processing power and therefore lower latency.
Reliability	Components don't degrade as fast due to greater thermal efficiency.
Less downtime	Better component life expectancy and minimal repairs reduce maintenance costs.
Reduced hardware costs	With better hardware life expectancy, component replacement costs decrease.
Reduced cooling costs	Less need for additional cooling hardware and fans means less money spent on active or system cooling.

Hot stuff: The takeaway

Rising demand for data, internet access, and bandwidth has increased the need for thermal management. This presents both challenges and opportunities for NVIDIA. NVIDIA GPU and DPU offerings will help meet the demands for AI throughout the world of data centers and telecom networks, and to fulfill the potential, NVIDIA will have to overcome obstacles of heat mitigation and environmental variations. Advanced, innovative component materials provide next-level thermal management that can help you reach your goals of improved reliability, performance, and cost management. They are the gateway to AI-focused data centers, 5G, Wi-Fi 6, 400 GbE data transmission rates, and beyond.

Selecting the right materials for routers, switches, and circuit boards is essential for data center, telecom, and fiber network performance. We invite you to align with an innovative, advanced materials partner to source and, in some cases, co-innovate material development. Henkel stands ready to work with you toward that goal.

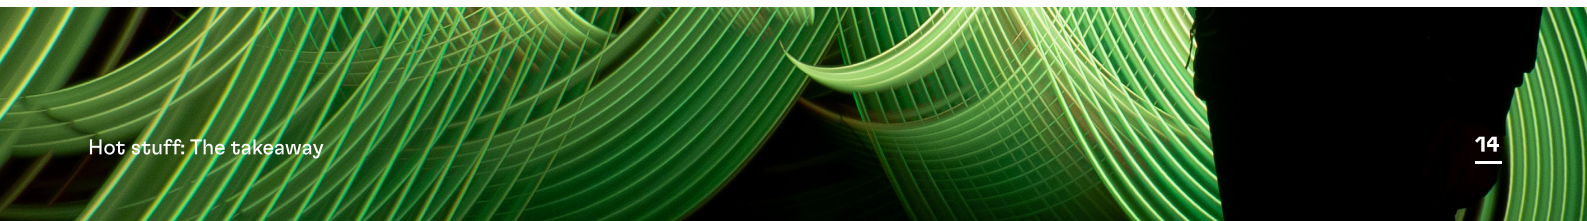

Henkel Adhesive Technologies