STULZ

CLIMATE. CUSTOMIZED.

Liquid Cooling

Exponentially increasing power density: At what pace should Liquid Cooling be implemented?



Release 01, October 2025

Contents

1. **Editorial** 3
2. **The AI-driven transformation in data centers** 4
3. **NVIDIA roadmap: The road to extreme power densities** 5
4. **The physical limits of Air Cooling** 8
5. **Global differences in markets and adoption tempo** 10
6. **The planning dilemma: Play the waiting game or make a start?** 12
7. **Planning strategies for AI data centers** 14
8. **Future-proof architectures** 16
9. **Conclusion** 17

© STULZ GmbH – All rights reserved

1. Editorial

The adaptation of AI workloads is pushing up power densities in AI ready data centers from 70 kW to levels currently at 125 – 130 kW per rack. By the end of 2027, we can expect this to increase exponentially to 600 kW and more. This trend marks the end of air cooling as the primary cooling medium in high-performance data centers. This is because air is physically unable to transfer enough heat to reliably cool high-density AI infrastructures. Liquid cooling transfers heat up to 1,000 times more efficiently and is therefore evolving into an indispensable component in modern data centers.

Where planning engineers are concerned, this development constitutes a paradigm shift in data center design. Today's data centers must be designed to withstand capacities of 125 to 140 kW, but also be modular to allow scaling to up to 600 kW per rack without having to rebuild the basic infrastructure from the ground up. Demand is growing for resilient, low-risk architectures with pre-dimensioned hydraulic paths, scalable CDUs (pay as you grow), and clearly defined interfaces that also support retrofit and greenfield projects.

Consequently, collaboration with data center cooling manufacturers like STULZ is shifting toward co-engineering based on KPIs: reference designs, validated test data and redundancy and control concepts are becoming an integral facet of the planning process. In addition, monitoring, group control of CDUs and the use of a digital twin early on in the planning phase ensure operational reliability, energy targets and compliance, bringing certainty despite ever faster hardware cycles.

Technological transformation
NVIDIA's roadmap shows exponential growth in capacity per rack: from the Blackwell generation with 125 kW to Rubin Ultra with 600 kW all the way to more than 1 MW for Feynman. This development follows a great leap in innovation that demands new planning strategies. As a strategic partner, STULZ develops validated infrastructure solutions for every generation.

Regional market dynamics
There are significant differences in global adoption: Asia leads the way with pragmatic learning-by-doing, followed by America with large capital investment, while Europe tends to be cautious due to regulatory complexity. With its 35 subsidiaries, STULZ responds to all approaches with a local presence and adapted solutions.

The central planning dilemma
Should data center operators start with the current chip generation, or wait for the next one? The answer to this question is fundamental to data center planning.

© STULZ GmbH – All rights reserved

2. The AI-driven transformation in data centers

Artificial intelligence marks an historic turning point for data centers. While traditional IT systems with 70 kW per rack were already considered high-density environments, we are now talking about 125–130 kW as standard, and already planning for infrastructures with a workload of 600 kW and more. These increases are impressive not just in terms of quantity, but also entail a qualitative paradigm shift.

Why AI places different demands on air conditioning

Multi-GPU modules with a shared memory and integrated onboard networks give rise to an exponential increase not just in computing power but also in thermal load – and in a way that pushes conventional cooling concepts to their limits. Here, we need to bear in mind that due to their thermal characteristics, AI applications differ fundamentally from conventional IT workloads.

Each software application has a different effect on heat generation in chips, which necessitates novel approaches in air conditioning. This variability is especially evident in different AI applications: chatbots work at a constant temperature at high load, Industry 4.0 applications follow a cyclical pattern in the medium to high load range, while machine learning training features extreme temperature fluctuations and requires ultra-fast reactions.

The power vs. sustainability dilemma

Another question with a growing impact on planning decisions is this: How can data center operators fulfill ever more demanding technological requirements without having to make sacrifices in terms of sustainability? Advanced AI platforms are defining new standards with high heat output per processor and rack densities of over 600 kW. At the same time, regulatory requirements are becoming more and more stringent. In Germany, the Energy Efficiency Act (EnEfG) dictates that new data centers commissioned on or after July 1, 2026 must achieve a PUE of at least 1.2 after no more than two years. Furthermore, a PUE threshold of 1.5 or less will be imposed on existing data centers as of July 1 , 2027, and 1.3 or less from July 1, 2030.

In this respect, the targeted integration of liquid cooling systems can help to gradually lower the PUE of existing data centers. Additionally, the EU's Energy Efficiency Directive requires the use of waste heat and corresponding cost/benefit analyses for new or significantly modernized plants with over 1 MW. Increasingly, these ESG (Environmental, Social, Governance) requirements are turning into important investment criteria and are therefore decisive factors in terms of competition.

© STULZ GmbH – All rights reserved

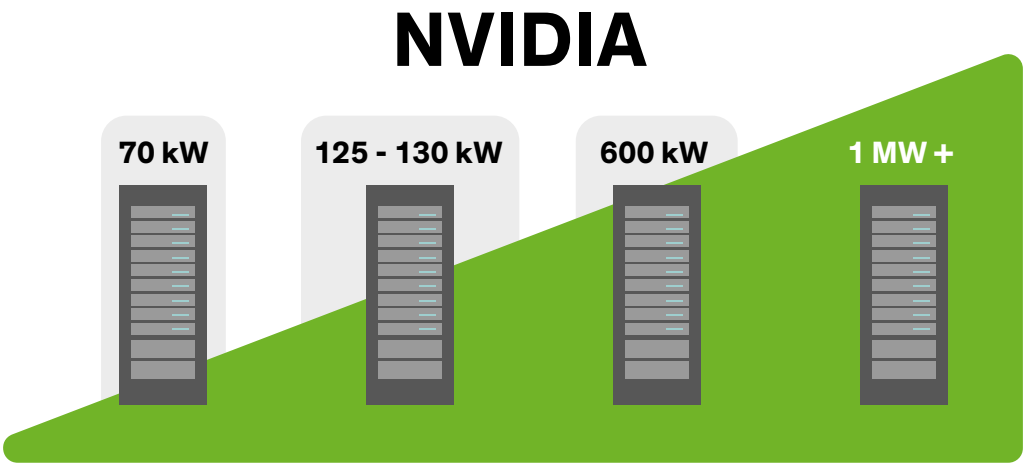
Tests under real conditions

At its headquarters in Hamburg, STULZ operates a state-of-the-art test center, which validates the critical aspects of liquid-to-liquid systems under real operating conditions. This collaboration with leading chip manufacturers ensures that STULZ cooling systems always satisfy the demanding requirements of each hardware manufacturer.

The test procedure encompasses complete system tests under real workload conditions, flow stability across multiple CDUs for high-density AI installations, filtration down to 25 µm for cold plate protection, and fluid monitoring to prevent corrosion and microbial growth. These exhaustive tests help data center operators to validate and gain certainty about their scaling requirements prior to rollout.

3. NVIDIA roadmap: The road to extreme power densities

Rapidly growing power densities and ever-shorter innovation cycles for AI accelerators are the main drivers behind the planning uncertainty currently besetting the entire industry – regardless of the individual manufacturer. Planned installations must therefore be designed to handle current capacities of 125 – 140 kW per rack in the near future, with modularity to allow scaling up to 600 kW per rack. Even 1 MW is on the horizon for this decade. The roadmaps of leading providers set out the timeframe and technical requirements for infrastructure design and cooling strategy and determine when hybrid solutions can be used, or when a transition to full liquid cooling will be necessary.



© STULZ GmbH – All rights reserved

The current generation: Blackwell as a turning point

The GB 200 Blackwell generation marks an historic turning point, with 125 – 130 kW per rack. For the first time, the critical threshold at which conventional air cooling is definitively not up to the job has been reached.

Blackwell Ultra (GB 300) raises the power density to 135 – 140 kW per rack. The main difference lies in the storage capacity, as Ultra chips can process larger volumes of data, which is particularly relevant for AI applications with high storage requirements.

2026: Rubin – Tried and tested system architecture for planning peace of mind

The Rubin generation will offer roughly the same memory as Blackwell Ultra, but its optimized design translates to greater speed and processing power. The result is more efficiency per watt – a significant step forward demonstrating that wins are in sight not just in terms of absolute performance, but also better energy efficiency.

Where cooling is concerned, the requirements with 135 – 200 kW per rack are similar to Blackwell's. This gives data center operators a great deal of certainty for their planning: infrastructures for today's high-density applications will also be sufficient for the Rubin generation.

2027: Rubin Ultra – The leap to 600 kW

In the second half of 2027, both memory and speed will double to power densities of up to 600 kW per rack. This generation marks a fundamental turning point in rack architecture generally.

NVIDIA has already developed a new rack design: the Kyber rack with 600 kW needs 100 % liquid cooling. There is no air cooling; servers are positioned vertically instead of horizontally and every component, including the power source, must be liquid-cooled. With power densities like this, all rack components will reach the hot spot range – even the mains supply and other components require liquid cooling.

2028+: Feynman and the megawatt challenge

The 3D stacking of the Feynman generation will reach power densities of over one megawatt per rack. With this technology, chips are stacked vertically, creating completely novel challenges for cooling systems:

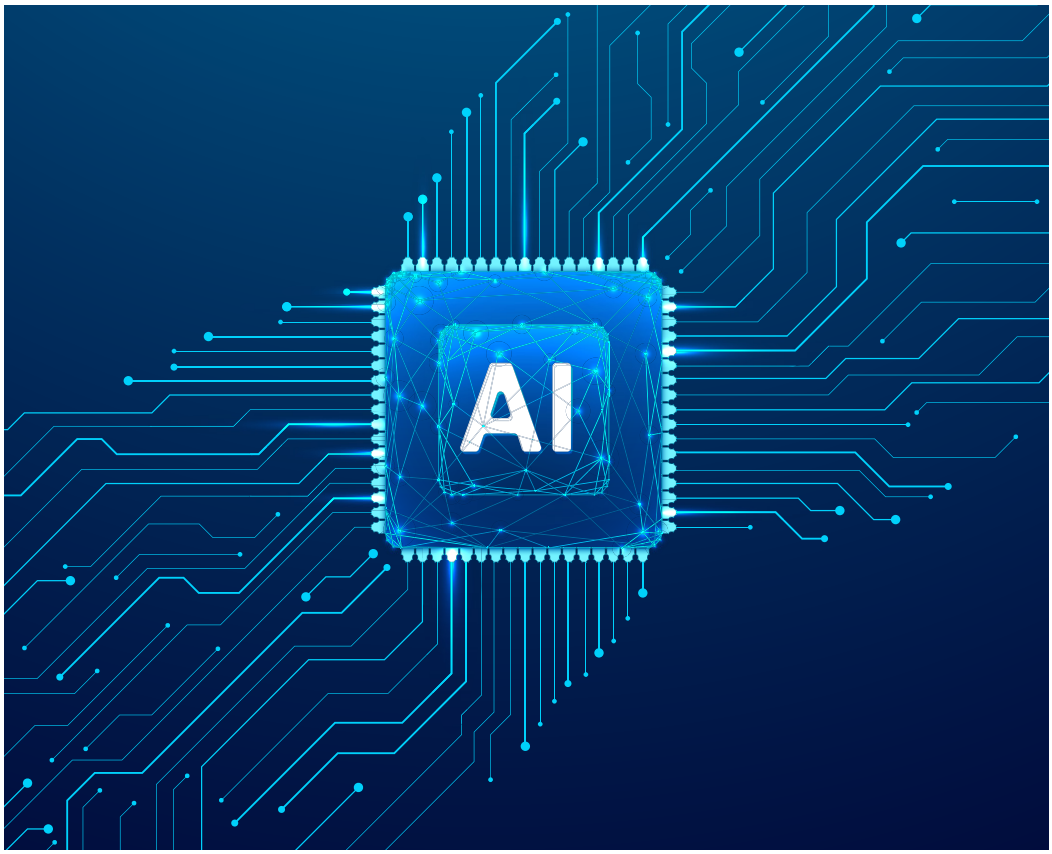
- Dissipation of heat from inside the chip stack
- Thermal hot spots in different layers
- Uneven heat distribution between the layers
- More complex cooling geometries

© STULZ GmbH – All rights reserved

Available production capacity and market dynamics

An often overlooked factor in the fulfillment of AI projects is limited production capacities for new generations of chips. Large investors already secure most of the available capacity for themselves on market launch. Smaller operators, who are waiting for the latest chip generation before entering into the AI sector, end up in a catch-22 situation due to lack of availability. This market dynamic reinforces the importance of rapidly building up expertise using currently available technologies. Here, existing relationships with manufacturers and technological partnerships with data center cooling specialists provide early access to reference designs, proven solutions for immediate implementation, and the possibility of step-by-step scalability to future generations.

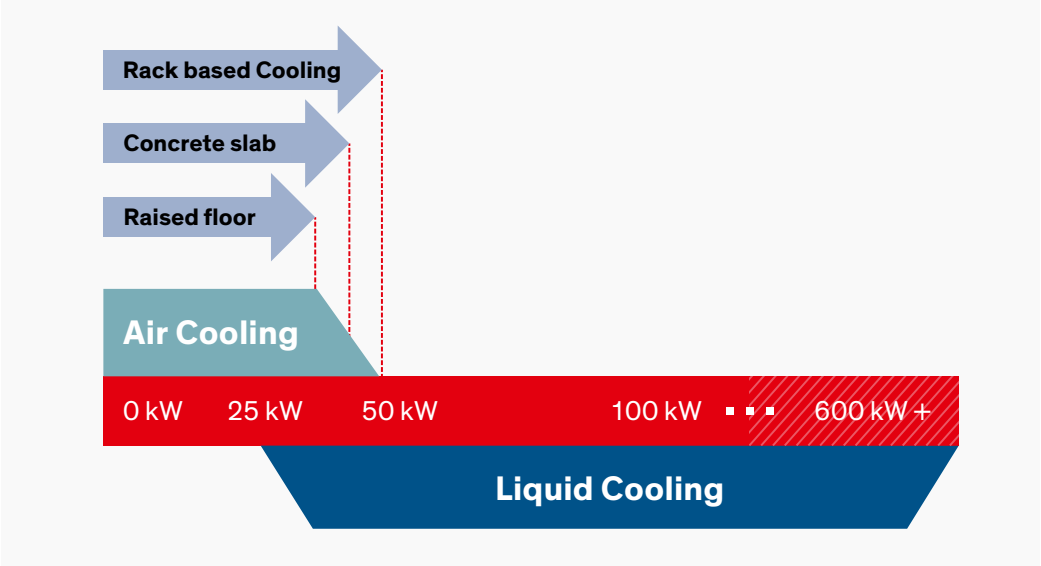
The roadmaps of many manufacturers clearly show that the development of GPU accelerators is exponential, not linear. For the next two to three years, liquid cooling systems will be a tried-and-tested and economical solution for all power densities up to 200 kW per rack. From 2027, modular, scalable concepts will enable a seamless transition to 600 kW systems with the same basic infrastructure. However, from 2028 at the latest, completely new, adaptive cooling strategies will be required. Their practical implementation will only be possible if it is based on decades of experience in air conditioning technology and high-density environments.



© STULZ GmbH – All rights reserved

4. The physical limits of Air Cooling

Air as a cooling medium has reached its physical limits and can no longer transfer sufficient heat as power densities increase. This is no purely theoretical limitation – it is already manifesting itself dramatically in the high-density applications of today.



The physical fundamentals of heat transfer

Air has a specific heat capacity of around 1 kJ / (kg × K), while water achieves 4.18 kJ / (kg × K) – already a factor of more than 4:1. The difference in density is even more telling: air at 20 °C has a density of roughly 1.2 kg / m³, whereas water's is 1,000 kg / m³. Consequently, liquid cooling systems are 30 to 46 % more efficient than air cooling systems.

These physical laws mean that as power densities climb, the required volumetric air flowrates also increase exponentially. At 600 + kW per rack, the volumetric air flowrates would be so great as to exceed the available space in the data center.

600 kW – the critical turning point

At 600 kW per rack, all rack components are in the hot spot range. This brings about a fundamental change: whereas only CPUs and GPUs require special cooling at low densities, at 600 + kW the mains supply and power source also need liquid cooling.

NVIDIA's 600 kW Kyber rack demands 100 % liquid cooling, as secondary components such as the power source and network infrastructure also generate too much heat for air cooling to cope with. Racks up to 140 kW are hybrid systems: 80 to 90 % of the heat is dissipated using liquid cooling, the remaining 10 to 20 % via air cooling. At 600 kW or more, this approach is no longer practicable.

© STULZ GmbH – All rights reserved

Thermal challenges for cooling AI workloads

AI applications come with specific thermal characteristics that extend far beyond simple heat dissipation. Chips produce heat differently depending on the software used, necessitating new approaches in cooling technology.

Load profiles of different AI applications:

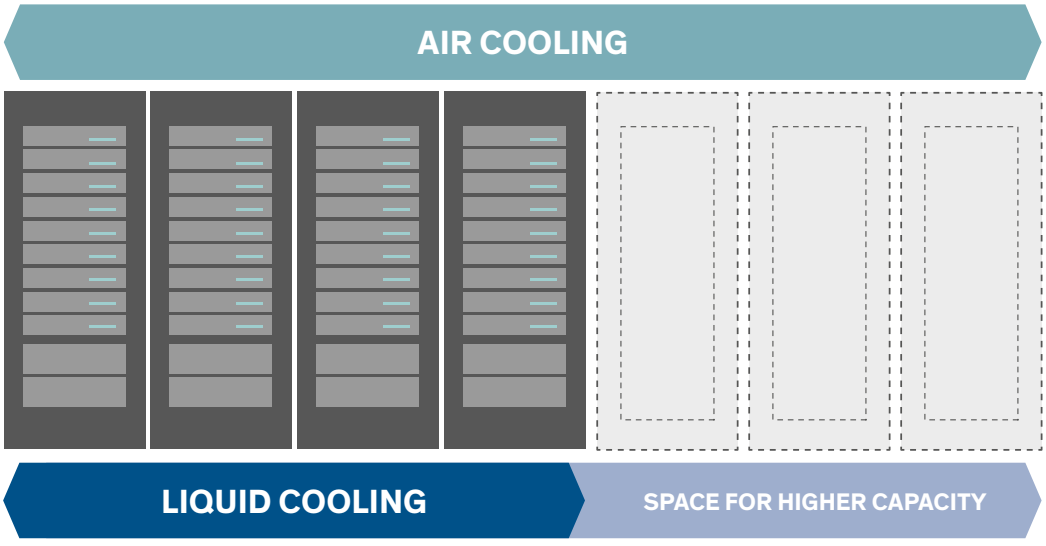
- Constant loads (chatbots): Uniform heat generation, can be perfectly managed with liquid cooling
- Medium fluctuations (Industry 4.0): Cyclic load patterns require adaptive control
- Extreme load fluctuations (ML training): Very fast reaction times needed, as major temperature variations occur

With extremely fluctuating loads, conventional systems cannot provide sufficient volumetric flow and cold media fast enough to bring down the temperature. Liquid cooling systems respond much faster than air cooling due to their constant differential pressure control modes and intelligent control.

Space efficiency and energy efficiency

Liquid cooling systems fundamentally consume far less energy than air cooling systems, leading to lower operating costs and a smaller carbon footprint. They also enable higher heat loads per rack, so that less space is needed for the same cooling capacity.

At 600 + kW per rack, space efficiency becomes critical. Liquid cooling already achieves obvious space savings compared to air cooling, and can further improve this efficiency with more compact heat exchangers. High-temperature return circuits can be connected to local district heating networks, enabling waste heat recovery.



© STULZ GmbH – All rights reserved

5. Global differences in markets and adoption tempo

The approach to the introduction of liquid cooling systems differs hugely from one region to another. These different approaches directly influence technological development and accompanying investment strategies.



Asia: Pragmatic trailblazer

The Asian market leads the way through fast decision-making and learning-by-doing. This pragmatic approach is evident in several ways: Technological decisions are often made in weeks, not months, pilot projects with reduced capacity are systematically used as learning platforms, a higher tolerance of experimental approaches enables different system configurations, and continuous improvement is understood as a process, not a straightforward dash to perfection.

This learning-by-doing mentality has put Asia ahead of the game in terms of knowledge. Asian reference projects include proven systems for a variety of AI workloads, modular scaling strategies from pilot projects to fully developed systems, and local adaptations in line with different climates and infrastructures.

USA: The hyperscaler market with high capital investment

The American market combines high investment with dominance by established hyperscalers. American hyperscalers are technology leaders and set international standards for the entire industry. In contrast to Asia's rather experimental approach, American companies prefer systematic, well-documented implementation. Large investments enable solutions at different levels to be implemented in parallel – from standard systems all

© STULZ GmbH – All rights reserved

the way to fully customized infrastructures. In addition, the American market provides regulatory stability, which facilitates long-term planning and creates ideal conditions for pay-as-you-grow strategies.

Europe: Regulatory leadership with restrained adoption

Europe has a rather more hesitant approach to the adoption of AI, but leads the way in sustainability regulations. This apparent reticence has complex causes with strategic implications: Europe sets global standards for energy efficiency and environmental protection; however, stringent sustainability regulations result in greater planning complexity, while different national approaches hamper standardization. As a result, investors are already withdrawing from Europe because it is easier, in economical terms, to operate in other regions.

European companies preferred proven technologies, an approach which works in favor of established liquid cooling systems. Regulatory requirements become factors for competition: PUE targets of less than 1.1 mean that liquid cooling is virtually the minimum standard, the upcoming heat recovery requirement favors systems with high-temperature return circuits, and cooling technologies are evaluated in terms of ESG when making investment decisions.

Supply chain risks and geopolitical factors

The majority of chips are produced in Asia, resulting in global dependency. However, this dependency covers not just chips, but also GPU packaging, cold plates and flow control elements. Geopolitical tension can lead to trade conflicts and export restrictions, leading to delays all along the supply chain.

Moreover, critical components are often provided by only a few suppliers, increasing the risk. Therefore, modern testing capacities and second-source strategies for critical components are becoming ever more important. Nowadays, some data center operators are demanding supply chain verification documentation from cooling infrastructure manufacturers.

Convergence despite different approaches

In spite of these different approaches, all markets are moving in the direction of liquid cooling. The regional differences actually complement one another: Asia accumulates practical experience, America invests in scaling, Europe sets standards. The complementary nature of this situation is allowing liquid cooling solutions to become available more quickly, because the different markets are making advances in different aspects of the technology, which are then made available globally.

This combination of regional strengths enables optimum solutions: Asian flexibility for rapid prototyping cycles, American scaling for hyperscale projects, and European standards for sustainable, compliant implementation.

© STULZ GmbH – All rights reserved

6. The planning dilemma: Play the waiting game or make a start?

The key dilemma in modern data center planning can be expressed by a seemingly simple question: should they start with the current chip generation, or wait for the next one? This question is preoccupying decision-makers all over the world, and leads to unnecessary delays in the planning of new zones specially designed for operating high-density AI systems.



The illusion of the waiting game

Many data center operators are faced with the question of whether they should start with the current generation of chips, or wait for the next one. This way of thinking seems rational – why not wait for the better, more efficient generation? However, the reality is more complex and demonstrates the weaknesses of the waiting approach.

This strategy leads to an availability catch-22: new generations are announced but available production capacity is limited, major investors claim a large portion of this output, and smaller operators go away empty-handed. If they then wait for the generation after this, the same process repeats itself.

© STULZ GmbH – All rights reserved

Building expertise as a competitive advantage

It is vitally important to gather experience and to install AI applications and high-performance chips now, in order to then be better prepared for the generations to come. Building up this expertise is valuable in several ways:

- Technical experience: understanding AI-chip reactions and heat fluctuations
- Operative experience: servicing and operating liquid cooling systems
- Strategical experience: supplier relationships and scalable strategies

Pragmatic solutions to planning uncertainty

For design engineers, uncertainty consists in wondering whether the design they are currently working on can be regarded as future-proof. This uncertainty is justified, but often leads to paralysis instead of adaptive solutions.

Key planning questions are: Will the infrastructure still be usable in 10 or 20 years' time? Can the system work with the next generation of chips? How do we plan for as yet unknown future requirements?

The answer lies in dimensioning without investment constraints. This can be achieved by means of pay-as-you-grow infrastructures with modular systems, oversizing basic infrastructures for future expansion, scalable, step-by-step concepts from kW to MW, and tests under real conditions for safe investment decisions.

© STULZ GmbH – All rights reserved

7. Planning strategies for AI data centers

The challenges of the AI revolution demand a fundamentally new approach to data center planning. In this age of extreme power densities, customized solutions rather than one-size-fits-all solutions are the winning strategy for success. In this regard, modern planning philosophy is based on three clearly defined levels for solutions, which can be scaled to suit the requirements and complexity.

Modular solution levels for different requirements

- **Level 1** offers standard systems with extensive equipment options without significant adaptation of mechanical components, which cover approximately 80 % of today's liquid cooling requirements. These systems are suitable for immediate entry and offer proven technology with a low implementation risk.

- **Level 2** enables standard equipment to be adapted to specific customers' needs, as different media mean different efficiencies. Here, systems are adapted to specific cooling media – from water to water/glycol mixtures to dielectric fluids.

- **Level 3** offers fully customized liquid cooling infrastructures as complete end-to-end developments for unusual requirements, which are produced in close partnership with customers.

System integration as a factor for success

Instead of a rigid separation between air and liquid cooling, STULZ's expertise and a product portfolio tailored to customer requirements enable the existing chilled water infrastructure to be supplemented in a targeted way by coolant distribution units (CDUs). This in turn enables liquid cooling systems to be incorporated precisely where high power densities arise. Existing chillers, pipework, pumps and heat exchanges remain usable in the shared facility water circuit. CDUs create an isolated loop separate from the IT side; they control temperature levels, volumetric flowrates and differential pressure, and provide the necessary level of redundancy.

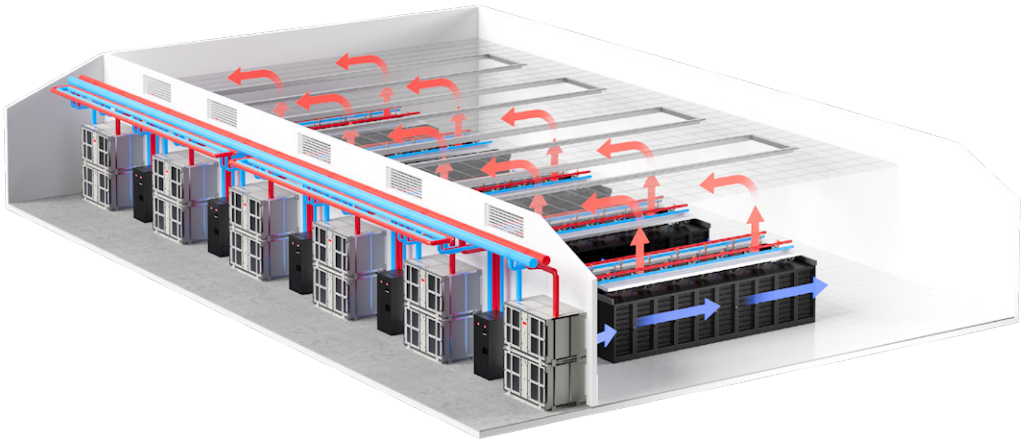
They open up the way to scalable integration, from individual high-density racks to inline segments to centrally supplied pods or zones, each with a demand-based N + 1 design and without the need to structurally separate the system. This approach shortens the time-to-value in brownfield environments, reduces conversion costs and therefore creates sufficient planning certainty for growing AI loads within the existing building envelope.

© STULZ GmbH – All rights reserved

Fully-fledged partnership instead of simple supplier relationships

A common problem in current planning practice is that cooling specialists are often the last to be involved in the planning process. The result is insufficient space for optimum cooling solutions, compromises in energy efficiency and performance, higher costs due to subsequent adaptations, and suboptimal implementation overall. In contrast, successful AI data centers demand an integral technological partnership right from the first draft.

Modern partnerships between manufacturers, specialist design engineers and data center operators therefore rely on KPI-based consulting that enables optimum system design. At the same time, specialized companies and planning agencies act as an important bridge between IT requirements and building services management systems. Instead of the traditional supplier relationship, this creates a strategic partnership that is geared towards joint concept development with the manufacturer and maximum future-proofness. Theoretical calculations are replaced by tests under real conditions, giving rise to a solid foundation for planning cost-effective and technically optimized solutions.



Pilot projects as a vital learning platform

Pilot projects with reduced capacity have proven to be a strategically important entry point into the operation of AI workloads. They allow technical tests with functional verification of new cooling concepts and enable the build-up of knowledge as strategic preparation for larger rollouts. At the same time, they enable supplier assessment with quality assurance and step-by-step migration to fully developed systems. This approach keeps risks to a minimum and ensures continuous learning before major investments are made.

The transition from experimental pilot projects to fully developed, production-ready systems takes place in clearly defined stages. Each stage builds on experience from the previous stage, and also provides the basis for the next level of scaling. This allows the creation of robust, scalable infrastructures that can satisfy both current and future requirements.

© STULZ GmbH – All rights reserved

8. Future-proof architectures

Dynamic roadmaps in AI accelerator production are shortening innovation cycles and causing a major shift in the thermal profiles of modern IT applications. Therefore, today's data center infrastructures should be designed to be as modular and adaptive as possible, without committing to one particular chip generation. The aim is to use the existing building services management system as a stable foundation and to only expand those segments in which high power densities will be produced by AI accelerators or can be expected in the foreseeable future.

Strategic oversizing at a few key points can considerably lower the cost of conversion in future expansion stages. Reserves in pipe sizes, heat transfer surfaces and pump assemblies allow future adjustments to volumetric flowrates and temperatures without the need for a fundamental rebuild. Coolant distribution units (CDUs) as an interface are a decisive aid here. They isolate the liquid cooling circuits of the AI accelerators from the existing facility water circuit and maintain defined temperature and differential pressure setpoints in the secondary LC circuit. This separation guarantees the necessary volumetric flow to the cold plates, while the chilled water circuit as a whole can remain unchanged.



+ CyberCool CMU

Coolant Distribution Unit for seamless heat transfer in liquid cooling solutions

- Up to 1,380 kW rated heat exchange
- Optimized for
  - high external pressure
  - low internal pressure drop
  - low Delta T
- Improved heat exchanger surfaces
- Maximum energy efficiency
- Ensures high liquid quality

© STULZ GmbH – All rights reserved

Uniform monitoring via air and liquid paths forms the basis for predictive operation, as real-time temperature, volumetric flowrate and pressure values enable control quality and energy consumption to be measured. The database accumulated in this way can be used for the step-by-step development of a digital twin for optimizing control parameters, calculating maintenance windows, and the transparent representation of available capacity reserves. At the same time, technological partnerships with tested reference designs reduce implementation risks and dramatically shorten the time-to-value.

Where regulatory trends are concerned, the integration of interfaces for heat recovery and validation in the planning process is absolutely essential. From October 2025, the EU's Energy Efficiency Directive will require a waste heat recovery cost/benefit analysis for new or significantly modernized systems over 1 MW. Data centers therefore remain connected to local heating networks, while also being audit-compliant as regards future testing and reporting obligations.

9. Conclusion

NVIDIA's latest GPU roadmap impressively illustrates the innovation cycles that data centers for AI workloads currently have to keep pace with. While today's Blackwell generation still has power densities of 125 – 130 kW per rack, in 2027 the Rubin Ultra generation will already reach 600 kW. Just one year later, Feynman will hit the 1 MW threshold. Even today's Blackwell generation marks a significant turning point: for the first time, power densities are so high that the resulting heat can only be dissipated by liquid cooling.

This development necessitates a completely novel approach to data center planning. To achieve this, the key lies in well thought-out entry based on solid, scalable standard components in the form of reference designs, combined with an expansion strategy based on the available data. Data center concepts that are currently designed to handle capacities of 125 to 140 kW must allow scaling to up to 600 kW per rack without the need to rebuild the basic infrastructure from the ground up. This can be achieved using architectures with pre-dimensioned hydraulic paths and scalable CDUs, plus clearly defined interfaces that can be used for both new and existing data centers.

Regions sometimes differ greatly in their approach to developing and implementing data center infrastructures for AI workloads: The Asian market takes a pragmatic learning-by-doing approach, entailing fast-track planning and the implementation of many projects with a high adoption rate, which are gradually optimized when operation is ongoing. In the American hyperscaler market, the prevailing approach is systematic with thorough documentation and high capital investment, which sets international standards.

© STULZ GmbH – All rights reserved

The European market has the strictest regulations, especially when it comes to sustainability. This results in far more complex planning procedures and a somewhat more hesitant pace in the adoption of new technologies.

Regardless of regional influences, a good starting point for project planning is the definition of a distinct zone or pods specially intended for the integration of AI workloads. In this zone, reliable reference values for different operating temperatures, energy efficiency and availability can be established. These values serve as the basis for all further expansion stages. This kind of empirical approach avoids costly planning miscalculations and creates trust in the chosen architecture among all involved parties – from the management to the IT team.

The nature of project management is also changing. Traditional tendering processes with rigid specifications are increasingly giving way to KPI-based co-engineering processes with technology partners and other data center operators. Joint pilot projects, proof of concepts, reference designs and intelligent control systems not only reduce technical risks – they also speed up the time-to-market for new services. By working together in this way, everyone can benefit from the experience of other data center operators and technology partners, and successfully adapt proven solutions for their own infrastructure.

Irrespective of the latest GPU generation, regional preferences regarding implementation and changes in tendering and planning processes, the important thing is to start planning and introducing new tech at an early stage. All the findings gained by doing this will help future generations to be implemented more quickly, ensuring long-term competitiveness.

© STULZ GmbH – All rights reserved



ABOUT THE AUTHOR

Jonathan Aaron Jürgens is a Junior Product Manager at STULZ GmbH, where he has been shaping the future of data center cooling since 2020. Holding a bachelor's degree in industrial engineering, Jonathan plays a key role in the development and market implementation of the CyberCool CMU. With hands-on experience from over 100 global liquid cooling projects since the CMU's launch, Jonathan has established himself as the in-house expert on liquid cooling infrastructures. To further share this expertise, he conducts liquid cooling seminars worldwide, equipping data center stakeholders with the knowledge needed to optimize cooling performance while enhancing sustainability.

ABOUT STULZ GMBH

STULZ is a global provider of innovative, mission-critical cooling solutions. Founded in 1947 and with its headquarters in Hamburg, Germany, it is represented by 35 subsidiaries, a network of more than 150 partners and ten production sites worldwide. Its solutions provide precise temperature and humidity control, as well as reliable cooling capacity, to safeguard the operation and longevity for data centers, telecom infrastructure and industry. STULZ prioritises energy efficiency and sustainability in its product development and manufacturing processes to deliver future proof cooling solutions.
For further information, please visit our website at **www.stulz.com**

© STULZ GmbH – All rights reserved

CLOSE TO YOU AROUND THE WORLD



STULZ Company Headquarters

STULZ GmbH
Holsteiner Chaussee 283
22457 Hamburg
Tel. +49 40 5585-0
products@stulz.de

For further information, please visit our website at **www.stulz.com** or download our **STULZ Products and Services** app.

AUSTRALIA
AUSTRIA
BELGIUM
BRAZIL
CHINA
FRANCE
GERMANY
INDIA
INDONESIA
IRELAND
ITALY
MALAYSIA
MEXICO
NETHERLANDS
NEW ZEALAND
POLAND
SINGAPORE
SOUTH AFRICA
SPAIN
SWEDEN
THAILAND
UNITED ARAB EMIRATES
UNITED KINGDOM
USA

+ 150 PARTNERS WORLDWIDE

V1.0 1025
© STULZ GmbH – All rights reserved