

# Essential Cooling System Requirements for Next Generation Data Centers

**White Paper #5**



## **Executive Summary**

**Effective mission critical installations must address the known problems and challenges relating to current and past data center designs. This paper presents a categorized and prioritized collection of cooling system challenges and requirements as obtained through systematic user interviews.**

# Introduction

Despite revolutionary changes in IT technology and products over the past decades, the design of cooling infrastructure for data centers has changed very little since 1965. Although IT equipment has always required cooling, the requirements of today's IT systems, combined with the way that those IT systems are deployed, has created new cooling-related problems which were not foreseen when the cooling principles for the modern data center were developed over 30 years ago. In this paper, a systematic approach of identifying and classifying user problems provides insight regarding the nature and characteristics of cooling systems in next generation mission critical installations.

This paper focuses on the problem of removing power in the form of heat from the mission critical installation. A related APC white paper #4: "Essential Power System Requirements for Next Generation Data Centers" addresses the related problems of providing power.

# Survey

A survey of management personnel relating to mission critical installations was conducted, interviewing corporate CIO's, Facility Managers, and IT Managers. Over 90 people were interviewed from over 50 different organizations including Fortune 1000 companies, Government and Education, and Service Providers. Approximately 50% of customers interviewed were from North America, 20% from Europe, and 30% from Japan, Pacific, Australia, and Asia (JPAA) region.

The nine-month survey utilized "Voice of the Customer" techniques, which relies on data collection of verbal and/or written responses to open-ended questions. This provides extremely unstructured responses, with the advantage that the responses are not limited or constricted by preconceptions within the question. During the course of the survey, some of the questions were expanded and/or changed in order to clarify ambiguous responses.

# Results: Cooling System Challenges in Mission Critical Installations

Survey responses were grouped according to common concepts, and for each group a solution requirement, corresponding to a challenge for mission critical installation design, was derived. This process identified 23 core challenges. These core challenges were then further grouped according to theme into the following 5 key theme areas:

- Adaptability / Scalability
- Availability
- Lifecycle Costs
- Maintenance / Serviceability
- Manageability

For each theme area, the challenge, underlying problem, and power system requirements are presented in tabular form. The highest priority problems are listed first under each theme, and were identified by priority as determined by number of mentions combined priority expressed by the respondents.

<b>Adaptability / Scalability Challenges</b>		
<b>Challenge</b>	<b>Underlying problems</b>	<b>Cooling System Requirements</b>
<b>Plan for a power density that is increasing and unpredictable</b>	Industry projections of power density requirements show great uncertainty but new data centers must meet requirements for 10 years. Must take into account IT refreshes that occur every 1.5 to 2.5 years.	System design that can be easily adapted, even retrofit, to cool high density racks which might be isolated cases or widespread in the future.
<b>Reduce the extensive engineering required for custom installations</b>	This engineering is time consuming, expensive, a key source of downstream quality problems, and it makes it very difficult to expand or modify the installation later.	Pre-engineered solutions that eliminate and/or simplify most planning and engineering.
<b>Adapt to ever-changing requirements</b>	Loads are frequently changed. It is difficult to know if the cooling system must be changed, and difficult to determine if the existing system can provide sufficient cooling.	A cooling system where it is possible to assure that a new load can be cooled, and where cooling can be easily and quickly directed to isolated high power loads without complicated construction and planning
<b>Allow for cooling capacity to be added to an existing operating space</b>	Many existing spaces were not designed for the power density that is currently being installed or planned. Adding cooling capacity to an existing operating data center or network room can be very difficult and expensive.	Retrofit options, which provide additional cooling capacity, possibly targeted at specific racks or equipment, which can be easily installed without complex planning or engineering, and without replacing or shutting down the existing systems.

The survey found adaptability challenges were the most important requirement. Particularly focused on problems involving the cooling of high density rack systems, and the uncertainty of the quantity, timing, and location of high density racks. This is complicated by IT refreshes in the data center or network room that typically occur every 1.5 to 2.5 years and is discussed in detail in APC White Paper #29: "Rack Powering Options for Data Centers and Network Rooms".

The survey showed that customers are often unable to predict if their cooling system will supply a future load, even when the characteristics of the load are known in advance.

<b>Availability Challenges</b>		
<b>Challenge</b>	<b>Underlying problems</b>	<b>Cooling System Requirements</b>
<b>Eliminate air mixing</b>	Mixing of supply and exhaust air to IT equipment lowers return air temperature to the CRAC unit and raises the supply air temperature to the IT equipment. CRAC units must be set to deliver very cold air to overcome this, resulting in poor cooling performance.	Systems that minimize the mixing of the exhaust and supply air at the IT equipment.
<b>Assure redundancy when required</b>	The failure of a CRAC unit in a redundant system reduces cooling capacity but also affects the physical distribution of the airflow. It is very difficult to plan and verify redundancy.	Systems that, by design, assure airflow and supply temperature to all IT equipment during the failure of a CRAC unit or associated infrastructure.
<b>Eliminate vertical temperature gradients at the face of the rack</b>	The temperature up and down the front of a particular rack can vary 10 degrees C. This effect is unexpected and the reasons why this happens are unclear to the users. This places unexpected stress on individual pieces of IT equipment and results in premature failure of equipment above the temperature gradient.	Systems that prevent hot exhaust air from returning to areas on the front of the rack, and assure that cool supply air is distributed uniformly up and down racks.
<b>Minimize liquid sources in the mission critical installation</b>	Liquid spills can damage IT equipment and cause the need for data center shut-down. Clean up and damage assessment is very difficult.	Minimize the need for liquid in the data center. If needed, operate the liquid system at low or sub-atmospheric pressure to prevent leaks.
<b>Minimize human error</b>	Uniquely engineered, poorly documented systems. Changing requirements require adjusting parameters of live systems.	Pre-engineered solutions that have comprehensive documentation and mistake-proofing features.

Survey respondents universally identified frustration with the ability to assure required input temperature and airflow to all IT equipment in the data center or network room, even when the load is not changing.

Respondents showed very low confidence in the ability of any redundant cooling features in their data centers to perform.

<b>Lifecycle Cost Challenges</b>		
<b>Challenge</b>	<b>Underlying problems</b>	<b>Cooling System Requirements</b>
<b>Optimize capital investment and available space</b>	System requirements are difficult to predict and systems are frequently oversized.	Modular systems that grow with the requirement.
<b>Accelerate speed of deployment</b>	The planning and unique engineering involved takes 6-12 months, which is too long when compared with the planning horizon of the organization.	Pre-engineered solutions that eliminate and/or simplify most planning and engineering.
<b>Lower the cost of service contracts</b>	Service contracts on unused or underutilized equipment is wasted.	Rightsized systems that can be scaled rapidly with changing requirements would reduce oversizing and the wasted service contracts associated with underutilized equipment.
<b>Quantify the return on investment for cooling system improvements</b>	The options available in the design of a cooling system are very complex and vary widely in cost. It is very difficult to determine the value provided by the options. Particularly when the realized performance is typically much different than the design performance.	Standardized designs where the system performance can be predicted and quantified accurately.

The survey found the lifecycle cost challenges were of less concern than adaptability and availability requirements.

The cooling system requirements to meet the lifecycle cost challenges share many features in common with the solution requirements for adaptability. In particular, pre-engineered, standardized, and modular solutions are needed.

<b>Serviceability Challenges</b>		
<b>Challenge</b>	<b>Underlying problems</b>	<b>Cooling System Requirements</b>
<b>Decrease Mean-Time-To-Recover</b> (includes repair time plus technician arrival, diagnosis, and parts arrival times)	Spare parts are not readily available. Large systems that require complex disassembly process to diagnose and to repair.	Modular systems using standardized spare parts that are inventoried on-site or locally. Simple repair procedures that do not require complex disassembly. Accessibility to components which are designed for quick replacement.
<b>Simplify the complexity of the system</b>	Systems are so complex that service technicians and in-house maintenance staff make errors and cause malfunctions when operating and maintaining the system. Status of the system cannot be easily determined or communicated during a crisis. Third party control systems are complex and unique and are never thoroughly tested, resulting in unexpected behavior during fault conditions.	Standardized systems with standardized ancillary equipment and standardized nomenclature. Pre-engineered and pre-tested control systems that don't take a lot of time to set up. Advanced diagnostics that provide detailed information for troubleshooting.
<b>Simpler service procedures</b>	Routine service procedures require disassembly of unrelated subsystems. Some service items are not easy to access when the system is installed. Highly experienced personnel are required for many service procedures.	System should allow in-house staff to perform the most common service procedures. Modular subsystems with connectorized interfaces to mistake-proof service procedures.
<b>Minimize vendor interfaces</b>	Cooling systems often involve multiple vendors and contractors and it becomes difficult for in-house and even vendor personnel to determine who is responsible for a problem, leading to the wasting of time and money.	Pre-integrated, pre-engineered systems with minimal contractor-sourced components where it is clear who is responsible for a problem.
<b>Learn from past problems and share learning across systems</b>	Uniquely engineered systems where learning on one system cannot be transferred to another. No clear way that solutions for one customer's problem are communicated to other similar customers.	Pre-engineered standardized systems where learning is shared through manufacturer notifications and automatic upgrade procedures.

A common theme among the serviceability challenges is a belief by the respondents that cooling equipment could be designed to make it much easier to service.

<b>Manageability Challenges</b>		
<b>Challenge</b>	<b>Underlying problems</b>	<b>Cooling System Requirements</b>
<b>Management system must give a clear description of any problems</b>	Cooling management systems report data, which often bears little relation to the actual problem symptoms. Cooling management systems rarely provide information that helps with a diagnosis to the component level when a fault occurs.	Provide data reports, which better match problem symptoms. Eliminate arcane terminology. Provide information, which assists in diagnosing faults to the component level. Provide detailed snapshot of system performance during problems for troubleshooting.
<b>Provide predictive failure analysis</b>	Many cooling components fail or trip unexpectedly, or degrade without being noticed. No advance warning is provided that could allow corrective actions that might prevent cooling loss.	Instrument the cooling system in a way that provides advance warning of component failures. In the case of consumable or finite-life items, automatically notify regarding remaining expected life and replacement intervals. Adjust system performance to accommodate degrading consumables where applicable.
<b>Aggregate and summarize cooling performance data</b>	Cooling performance data is often not summed from separate CRAC units, providing poor insight into overall system performance. Operation of separate CRAC units is often not coordinated.	Graphical user interfaces and automatic notification, which report, manage, and notify based on parameters at the consolidated system level and at the individual CRAC level. Communication between systems to prevent demand fighting.

The manageability solution requirements are expensive to design, install, and test in uniquely engineered systems. These challenges clearly suggest the need for pre-engineered, pre-tested, and standardized management tools.

Survey respondents showed a lack of awareness regarding the time-varying power consumption of the newest generation of IT equipment, a characteristic which will give rise to time-varying heat outputs. Therefore managing this issue did not emerge as a challenge. Nevertheless, this issue is expected to emerge as a key manageability challenge in the near future and is discussed in detail in APC White Paper #43: "Dynamic Power Variations in the Data Center and Network Room".



## Contrast Between Power and Cooling Challenges

A related study of power challenges in mission critical installations reveals that 13 of the cooling challenges are in common with powering challenges, 4 are closely related, and 6 of the cooling challenges are completely distinct from the powering challenges.

The common power and cooling challenges share the theme of adaptability, particularly to the unpredictable and changing requirements of IT equipment that is added or swapped out during the life of the data center. Modular, scalable, pre-engineered, and standardized systems are the common solution for these problems. A number of concerns regarding service were also common between power and cooling needs.

The needs relating to system management were similar.

The biggest differences between power and cooling challenges relate to performance and cost concerns. The respondents clearly indicated that they were very concerned with the performance and availability of cooling systems. There was a clear pattern that customers did not believe that their existing cooling systems were operating as intended, and most users were uncertain as to whether their planned cooling redundancy, if any, would actually function during a fault condition. Performance issues were a greater concern than life cycle cost for cooling users. This is a striking contrast to the power survey, in which life cycle cost issues were prioritized higher suggesting a greater satisfaction with the performance of power systems.

## Cooling Systems for Mission Critical Installations

To satisfy the mission critical installation cooling challenges identified in this survey, there are a number of changes required from current design practice. Many of these changes will require changes in the technology and design of cooling equipment, and how it is specified. Integration of the components of the cooling subsystem, particularly the air distribution and return systems, must move away from the current practice of unique system designs, and toward pre-engineered and even pre-manufactured solutions. Such solutions would ideally be modular and standardized, expandable at will, and would ship complete but in parts that would rapidly plug together on site. Standardization will facilitate the learning process. By spreading the cost of developing high performance management systems across large numbers of standardized installations, coordinated management would be affordable to all customers.

One suggested solution to the problem of cooling for high density racks is direct water cooling to the rack. IT equipment itself may be water cooled, or heat exchangers in the rack may be used. Advances in air distribution systems will affect the power level at which this approach is preferred. For isolated racks above 10kW, or large groups of racks above 6kW, water-cooling currently appears to have benefits. However, this represents a very small fraction of racks today. Therefore, direct water-cooling to the rack is not expected to become a mainstream solution for the foreseeable future, and few references to it were found in the survey data.

## Conclusions

A systematic analysis of customer problems relating to data center and network room cooling systems provides a clear statement of direction for mission critical installations. The most pressing problems that are not solved by current design practices and equipment have the common theme of the inability of the data center to adapt to change. Data center cooling systems must be more adaptable to changing requirements, in order to improve both availability and cost effectiveness.

Cooling users are not confident that their current and planned systems will be able to cool high density racks. This is primarily a problem relating to air distribution and mixing. Data center cooling systems must provide the capability for greater control of airflow at the rack.

In many industries, a maturity level is reached where new advances in reliability, cycle time, and cost require standardization, pre-engineering, and modularization. Designers of mission critical installations, designers of the cooling equipment used in them, and owners should consider whether this point has been reached. The results of the survey in this paper suggest the need for a new generation of adaptable cooling systems for mission critical installations.

## References

- 1) FIPS PUB 94 "Guideline for Computer Power for ADP Installations"; National Technical Information Service