Cooling of Data Center

Part III

王啟川, PhD

國立交通大學機械工程系教授

Fellow ASME, Fellow ASHRAE

Tel: 03-5712121 ext. 55105

E-mail:ccwang@mail.nctu.edu.tw

Outline

- Background
- Loading efficiency improvements
- CRAC & Its selection
- Data center environment
- Energy Monitoring Indices
- Economizers (free cooling)
- Cooling technology
- Liquid cooling and Piping short overview
- Physical design – air cooling
- Ventilation design – some guidelines
- Some Thermal Measurements – Guidelines
- Some calculations/suggestions from publications
- Short Summary

Data Center

- Highly energy-intensive and rapidly growing
- Consume 10 to 100 times more energy per area than a typical office building
- Large potential impact on electricity supply and distribution
- Used about 45 billion kWh in 2005 (USA)
- At current rates, power requirements could double in 5 years.

More Facts…

- Servers, including its infrastructure, account for 1.2% electricity consumption (US).
- Every Watt of electricity consumed by IT equipment, an extra 1.5 Watts is needed for infrastructure to support IT equipment.
- Most servers require 1 watt of cooling for every watt of power used in moderately dense server system.
- High dense servers requires 2 Watts of cooling for every watt used in the system.

Worldwide electricity used in data centers Environ. Res. Lett. 3 (2008)

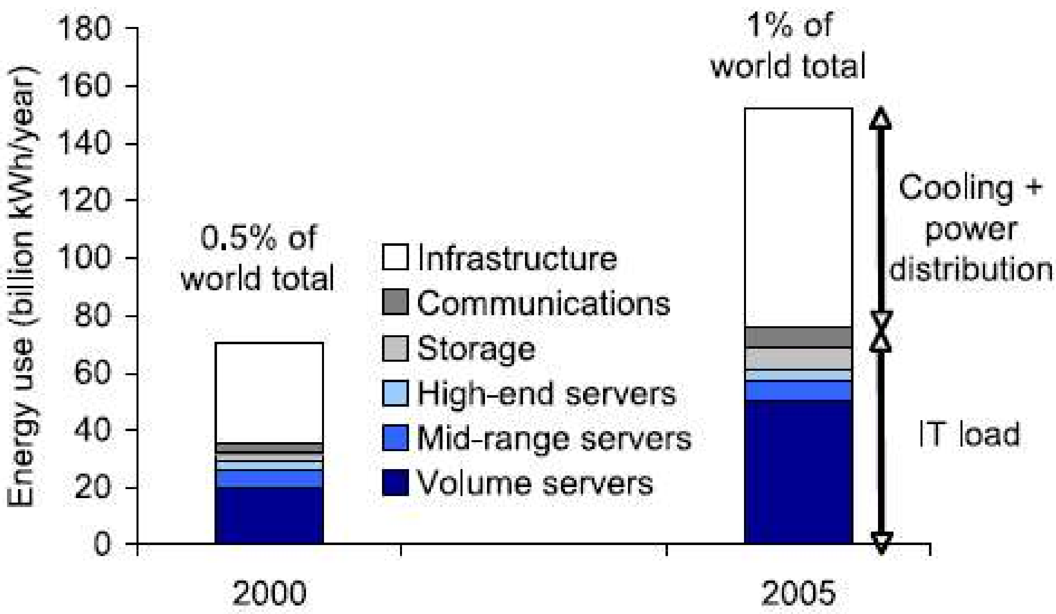

Figure 1. Total electricity use for data centers in the US and the world in 2000 and 2005, including cooling and auxiliary equipment. Total world electricity consumption was 13 238 billion kWh in 2000 and 15 747 billion kWh in 2005, according to the data in table 6.2 of the US Energy Information Administration's *International Energy Outlook*, downloadable at http://www.eia.doe.gov/iea/elec.html. Data center communications electricity use includes only that for networking equipment internal to data centers—it does not include the electricity use of the networks connecting data centers to the Internet as a whole or to the other parts of that broader network.

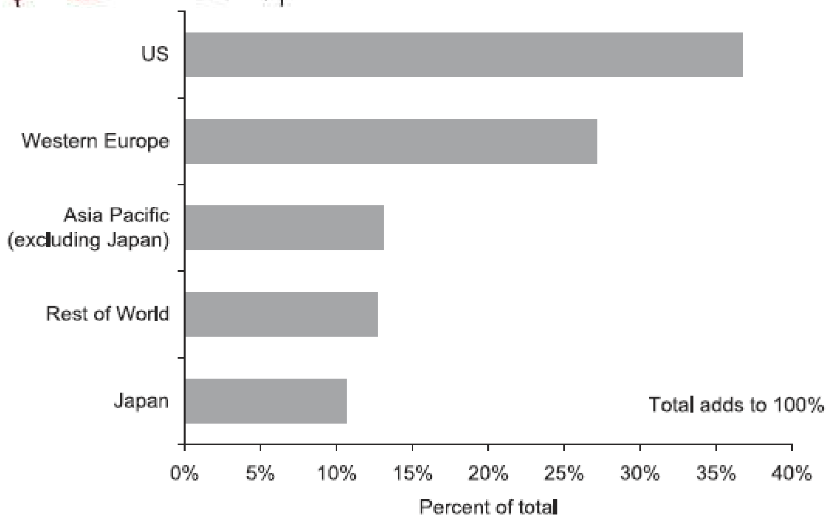

Figure 2. Regional distribution of electricity use for data centers in 2005.

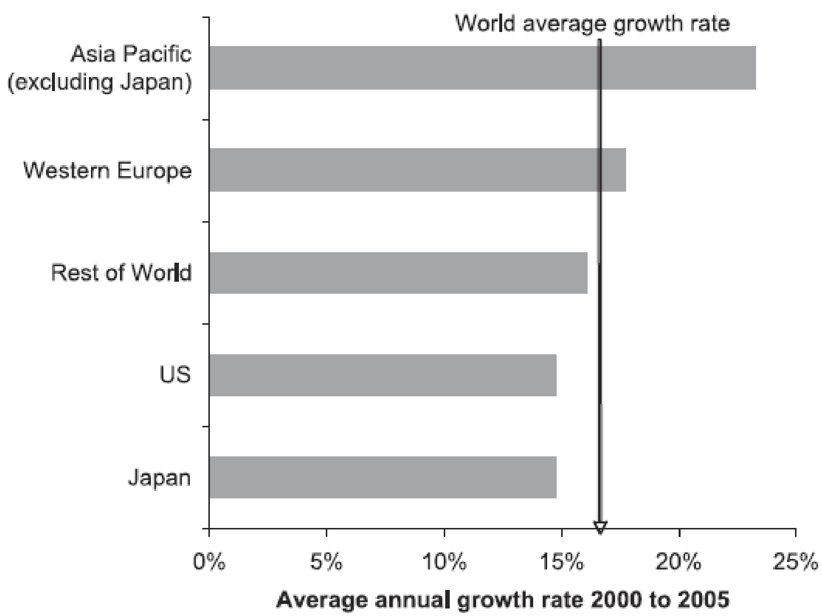

. Average annual percentage growth rates in data center electricity use by major world region, 2000–2005.

Table 1. Installed base and server power per unit in 2000 and 2005 by major world regions.

Installed base	Units	Volume	Mid-range	High-end	Total/avg
2000					
US	Thousands	4 927	663	23	5 613
Western Europe	Thousands	3 332	447	15	3 794
Japan	Thousands	1 140	250	15	1 405
Asia Pacific (ex. Japan)	Thousands	1 416	132	4	1 552
Rest of World	Thousands	1 425	317	8	1 750
Total	Thousands	12 240	1808	66	14 114
2005					
US	Thousands	9 897	387	22	10 306
Western Europe	Thousands	6 985	356	15	7 355
Japan	Thousands	2 361	185	12	2 558
Asia Pacific (ex. Japan)	Thousands	3 553	137	4	3 694
Rest of World	Thousands	3 162	199	7	3 368
Total	Thousands	25 959	1264	59	27 282
Average power used per server	**Units**	**Volume**	**Mid-range**	**High-end**	**Total/avg**
2000					
US	Watts/server	186	424	5534	236
Western Europe	Watts/server	181	422	4517	227
Japan	Watts/server	181	422	4517	271
Asia Pacific (ex. Japan)	Watts/server	181	422	4517	212
Rest of World	Watts/server	181	422	4517	246
Total	Watts/server	183	423	4874	236
2005					
US	Watts/server	219	625	7651	250
Western Europe	Watts/server	224	598	8378	258
Japan	Watts/server	224	598	8378	289
Asia Pacific (ex. Japan)	Watts/server	224	598	8378	247
Rest of World	Watts/server	224	598	8378	263
Total	Watts/server	222	607	8106	257

Note: (1) Installed base for US and World taken from Koomey (2007b). Non-US installed base by region was not available from IDC, so it was approximated using IDC shipments data by region and multipliers characterizing the relationship between installed base and shipments for all non-US regions in the aggregate (Koomey 2007a). This approach assumes that installed base for each non-US region grows in the same manner as does the sum of those regions. (2) Average power used per server for US and World taken from Koomey (2007b). Non-US average power per server calculated for non-US regions using the differences between US and World installed base and direct electricity consumption from Koomey (2007b).

Background – General Trend

Datacom equipments power trends and cooling applications ASHRAE 2005

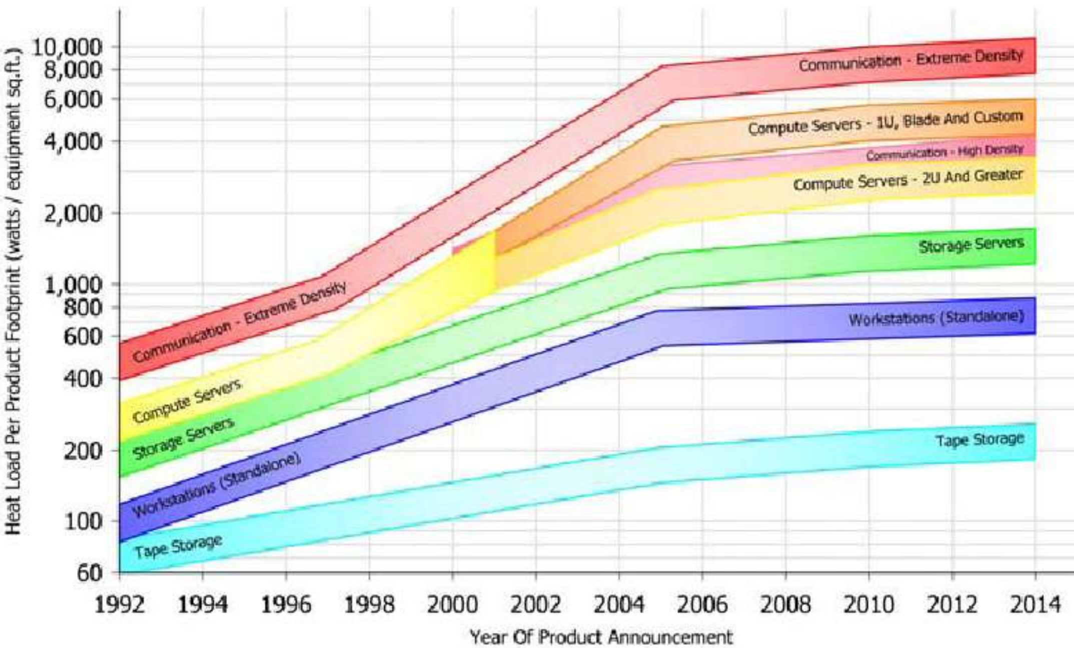

New ASHRAE updated and expanded power trend chart.

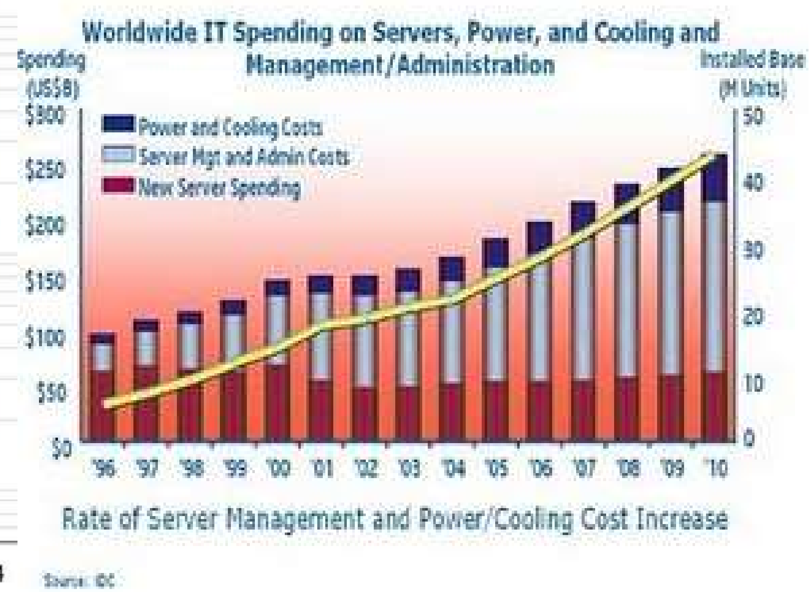

IDC Report – Cost Structure and Trends

2005 Predicted trend

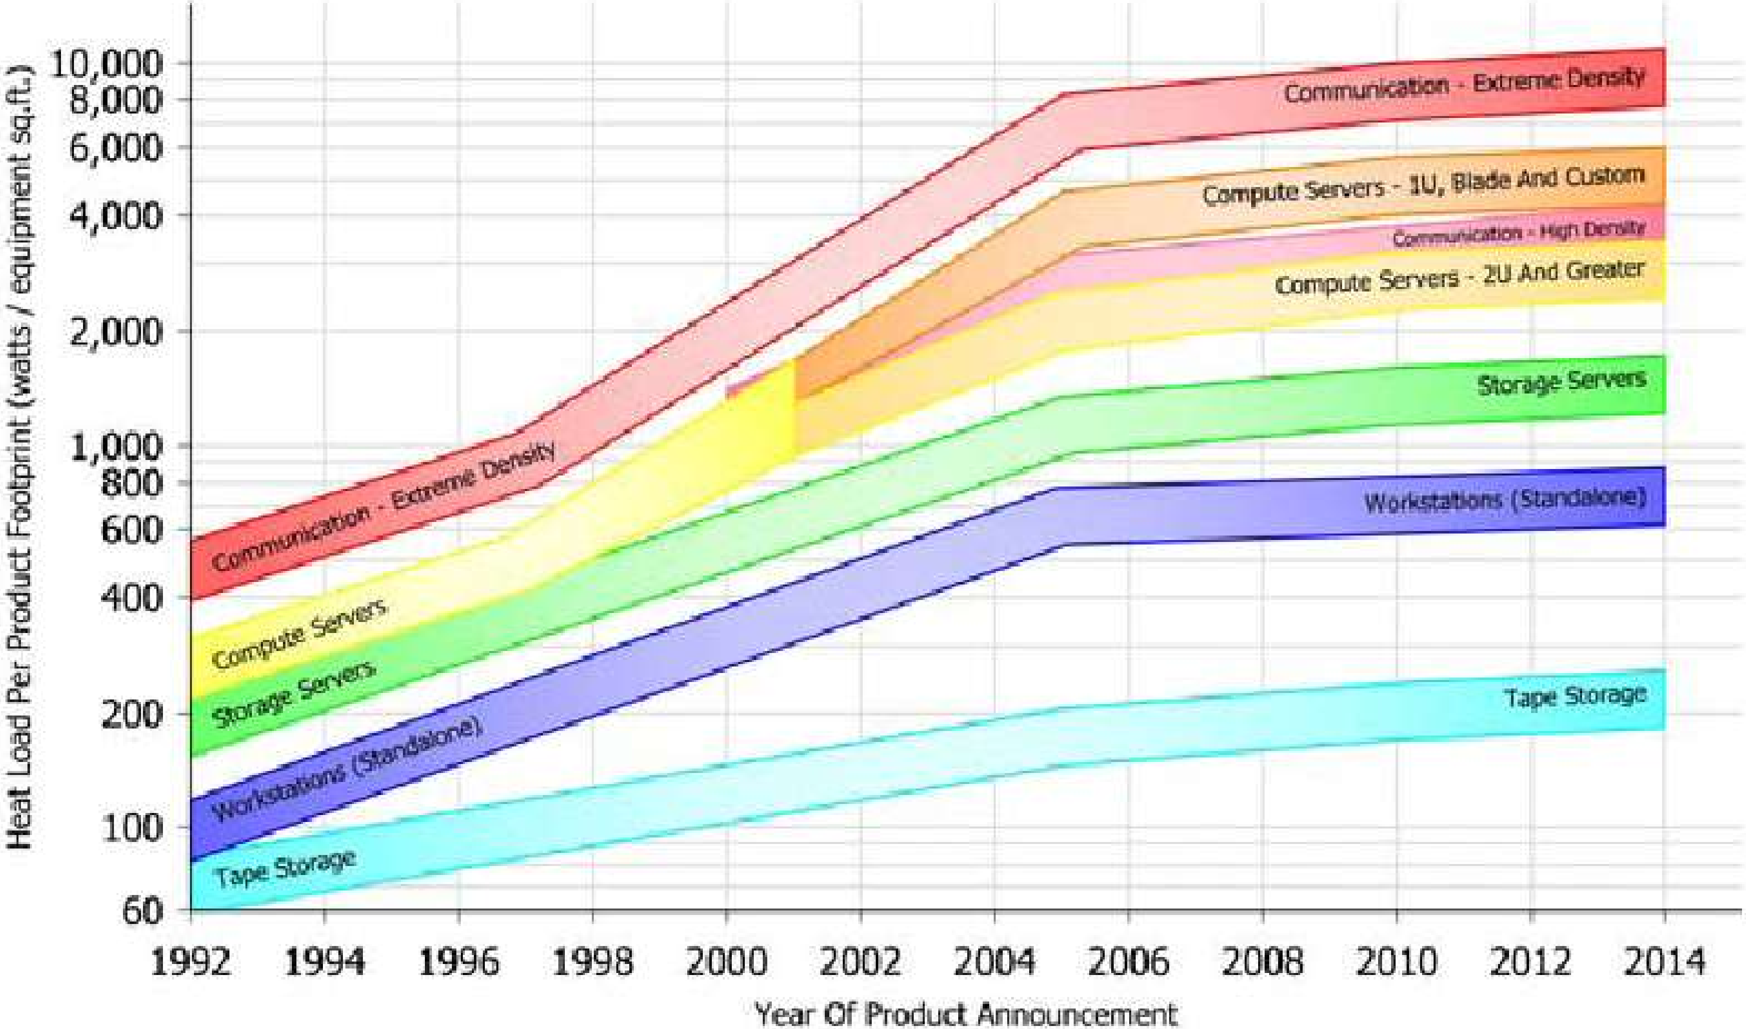

Figure 3-10 New ASHRAE updated and expanded power trend chart.

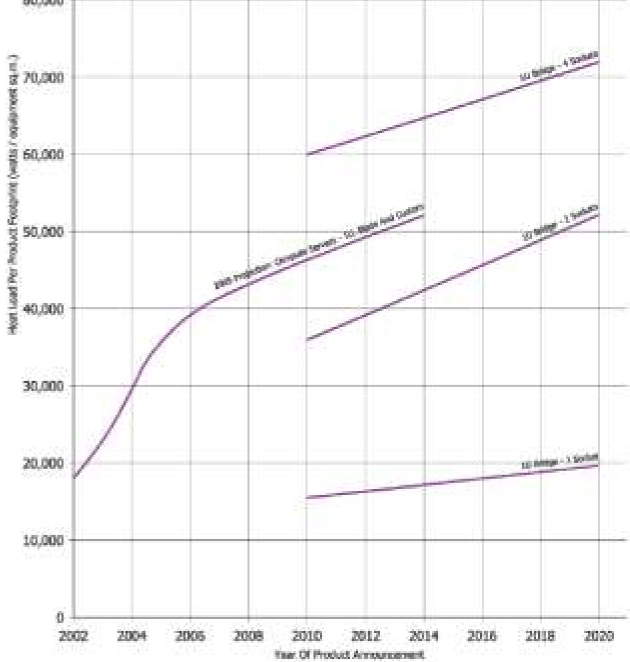

Figure B.9 1U Servers—2005 and 2012 trends (non-log scale, SI units).

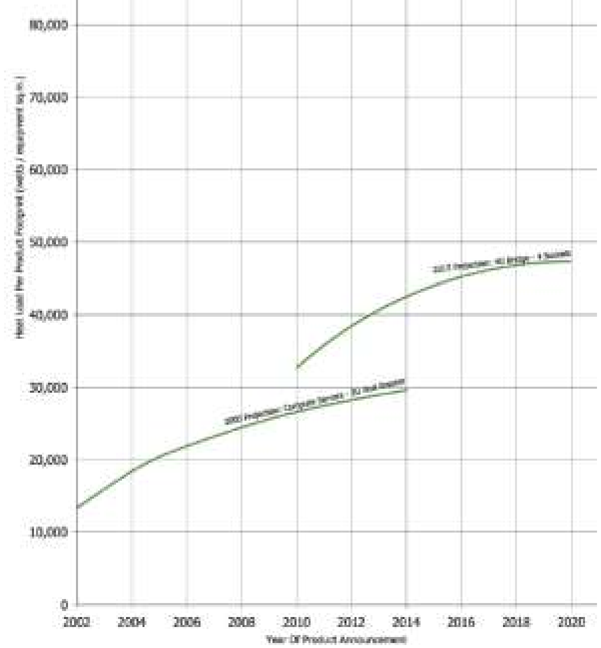

Figure B.11 4U Servers—2005 and 2012 trends (non-log scale, SI units).

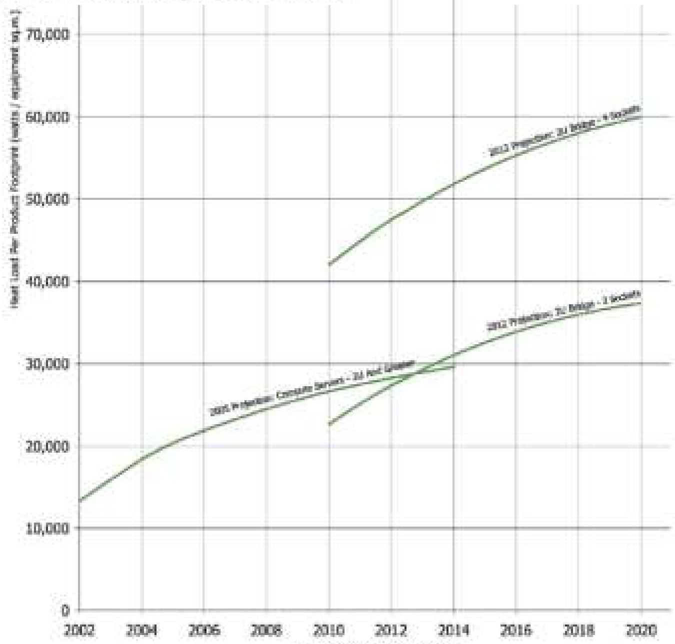

Figure B.10 2U Servers—2005 and 2012 trends (non-log scale, SI units).

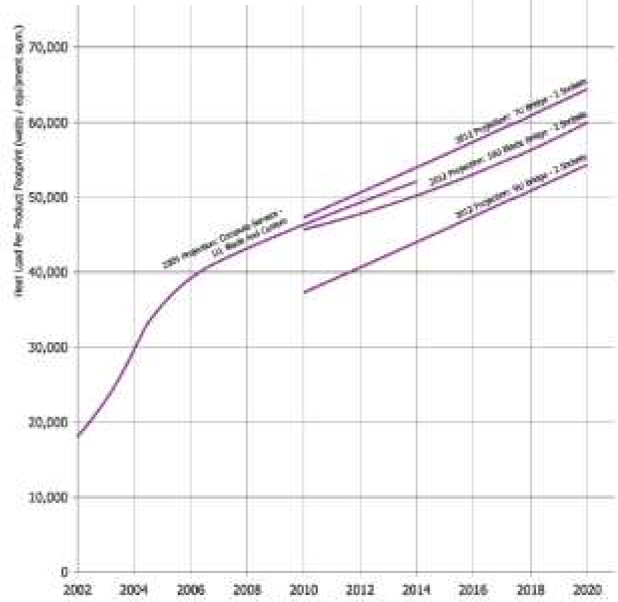

Figure B.12 Blade servers (7U, 9U, and 10U)—2005 and 2012 trends (non-log scale, SI units).

- This chart is based on ***maximum measured load***.
- The data shown in the power trend chart provide a general overview of the actual power consumed and the actual heat dissipated by data processing and telecommunications equipment.
- The data emulate the most probable level of power consumption assuming a fully configured, highly utilized system in the year the product was first shipped
- Finally, the intent of the trends is that they are to be used as a forecasting and planning tool by providing guidance for the future implementation of the different types of hardware.
- Not all products will fall within the trend lines on the chart at every point in time.

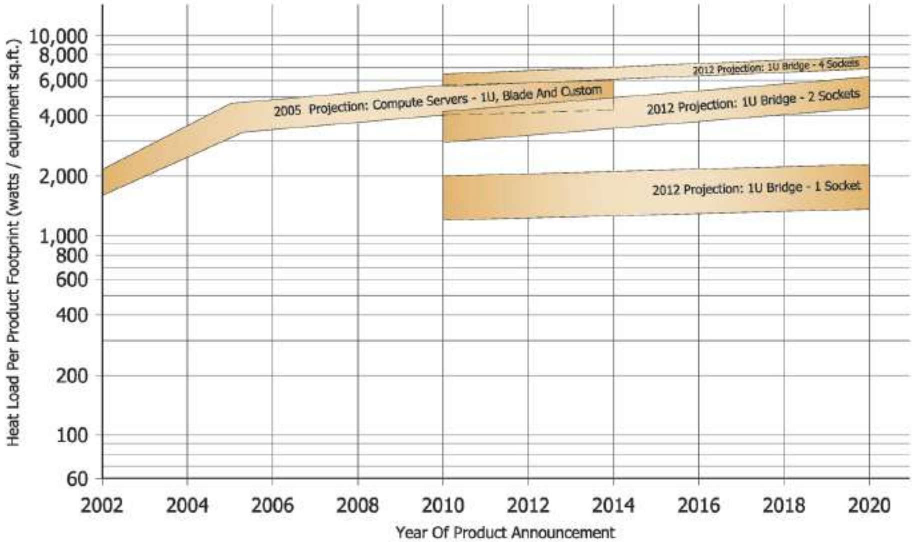

Figure 4.4 1U servers—2005 and 2012 trends.

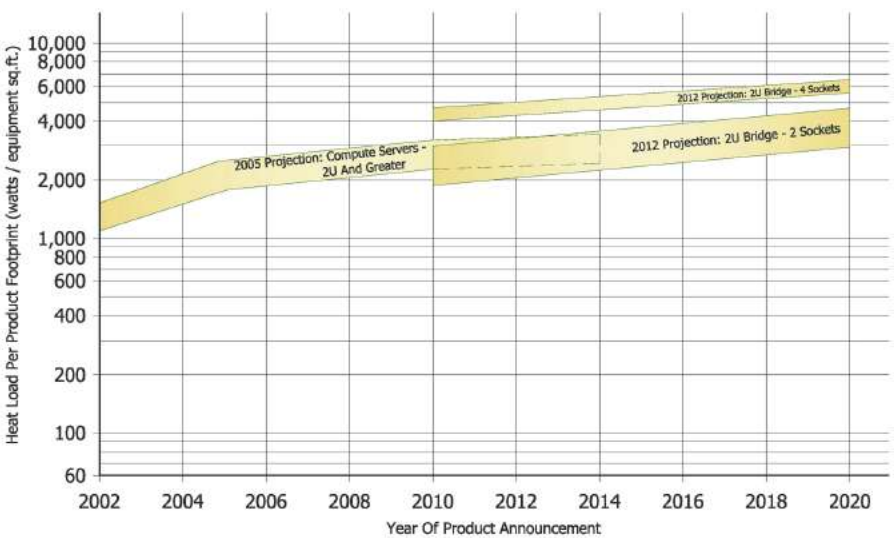

Figure 4.5 2U servers—2005 and 2012 trends.

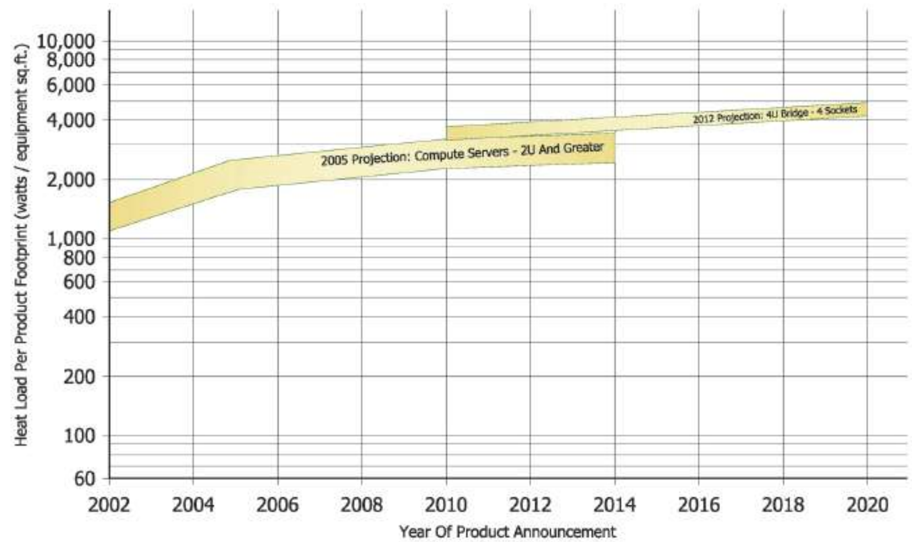

Figure 4.6 4U servers—2005 and 2012 trends.

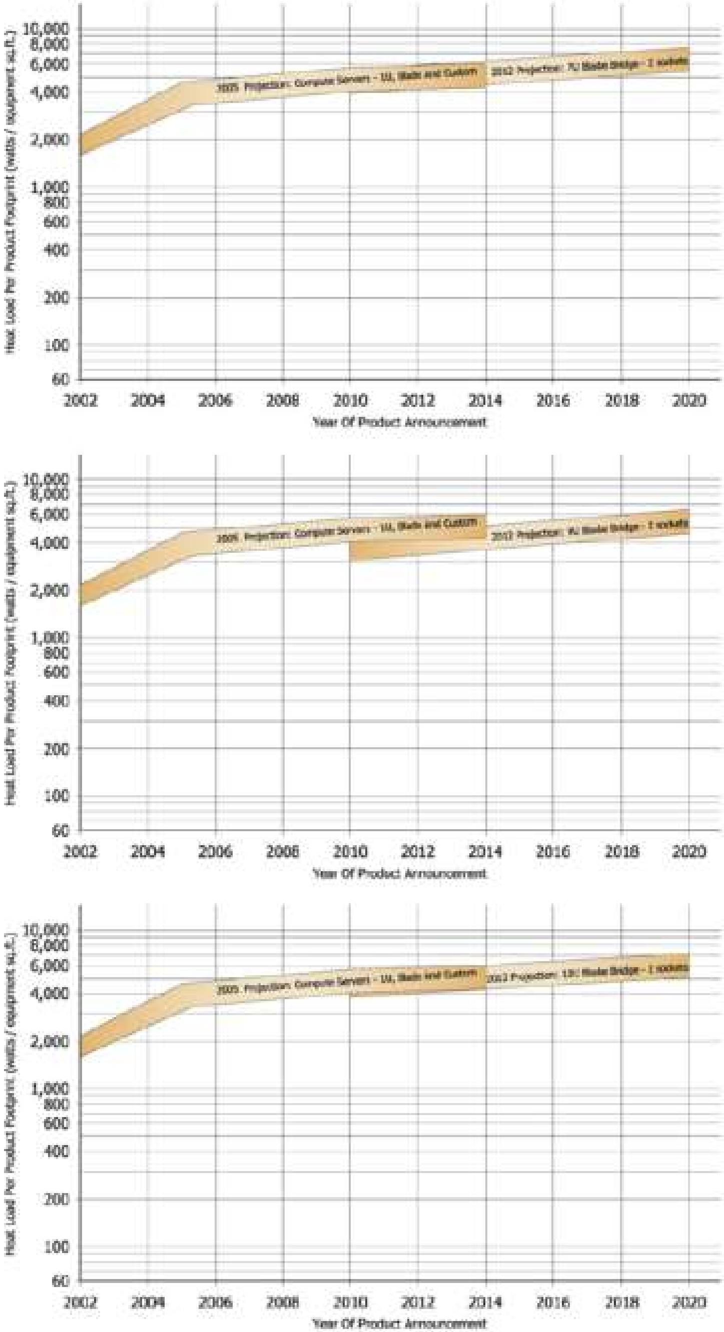

Figure 4.7 Blade servers (7U, 9U, and 10U)—2005 and 2012 trends.

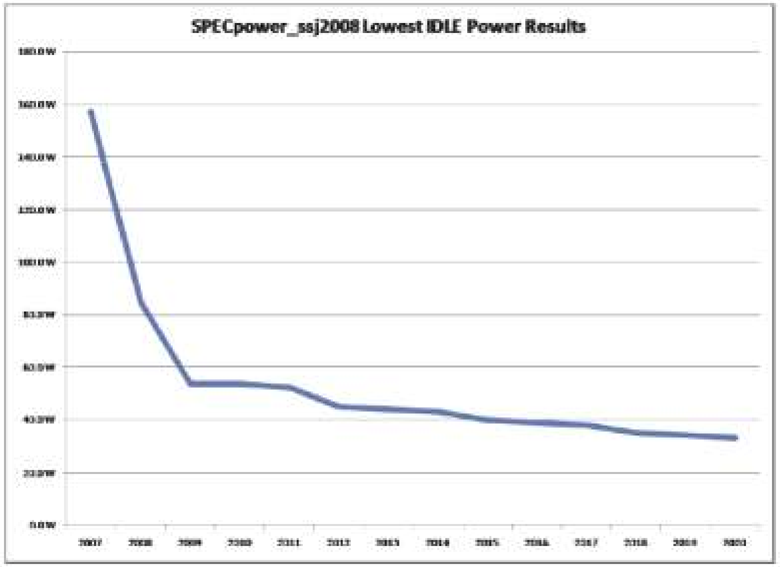

Figure 4.8 SpecPower trend in idle power.

Table 4.2 Power Trends of Nonstandard-Planform Equipment

Type	Range of Average Heat Loads	Range of Footprints, ft^2 (m^2)	Heat Load per Product Footprint, W/ft^2 (W/m^2) 2010	2015	2020
Storage Servers	±15%	6 to 13.5 (0.6 to 1.3)	700 (7500)	850 (9150)	1,100 (11,850)
Tape Storage	±30%	10 to 12 (0.9 to 1.1)	200 (2150)	200 (2150)	200 (2150)
Communications	±20%	6 to 12 (0.6 to 1.1)	2000 (21,500)	2550 (27,500)	3000 (32,300)

Cooling System Challenges

(APC white paper #5)

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

Adaptability / Scalability Challenges

Challenge	Underlying problems	Cooling System Requirements
Plan for a power density that is increasing and unpredictable	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.
Reduce the extensive engineering required for custom installations	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.
Adapt to ever-changing requirements	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
Allow for cooling capacity to be added to an existing operating space	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.

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.

Availability Challenges

Challenge	Underlying problems	Cooling System Requirements
Eliminate air mixing	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.
Assure redundancy when required	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.
Eliminate vertical temperature gradients at the face of the rack	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.
Minimize liquid sources in the mission critical installation	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.
Minimize human error	Uniquely engineered, poorly documented systems. Changing requirements require adjusting parameters of live systems.	Pre-engineered solutions that have comprehensive documentation and mistake-proofing features.

Lifecycle Cost Challenges		
Challenge	**Underlying problems**	**Cooling System Requirements**
Optimize capital investment and available space	System requirements are difficult to predict and systems are frequently oversized.	Modular systems that grow with the requirement.
Accelerate speed of deployment	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.
Lower the cost of service contracts	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.
Quantify the return on investment for cooling system improvements	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.

Serviceability Challenges

Challenge	Underlying problems	Cooling System Requirements
Decrease Mean-Time-To-Recover (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.
Simplify the complexity of the system	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.
Simpler service procedures	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.
Minimize vendor interfaces	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.
Learn from past problems and share learning across systems	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.

Manageability Challenges

Challenge	Underlying problems	Cooling System Requirements
Management system must give a clear description of any problems	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.
Provide predictive failure analysis	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.
Aggregate and summarize cooling performance data	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.

Loading efficiency Improvements

Loading of a data center
(APC white paper #126)

- In a typical data center, less than half of the electricity is go into IT

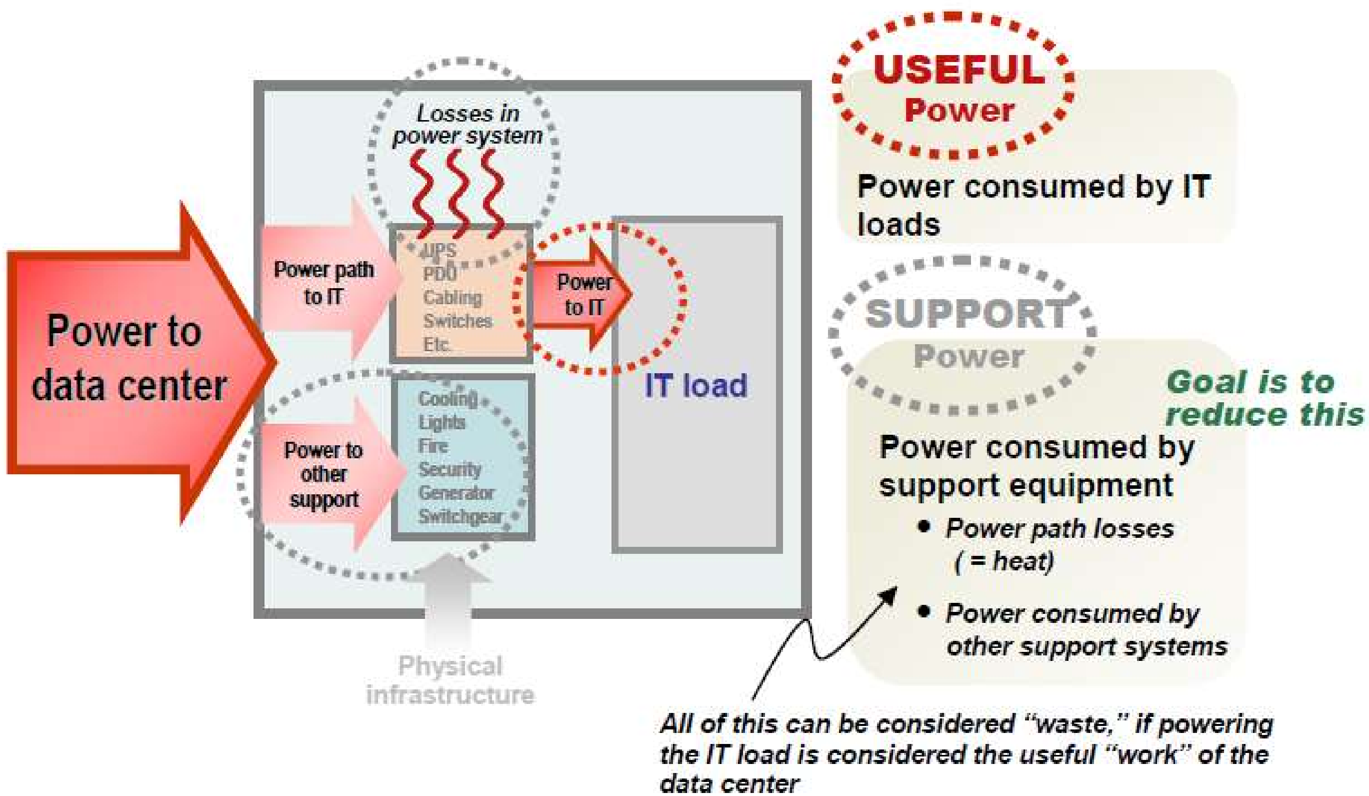

Figure 1

Power consumption in the data center

Five contributors to electrical inefficiency

*These all contribute to SUPPORT power in **Figure 1***

1. Inefficiencies of the power equipment
2. Inefficiencies of the cooling equipment
3. Power consumption of lighting
4. Over-sizing of the power and cooling systems
5. Inefficiencies due to configuration

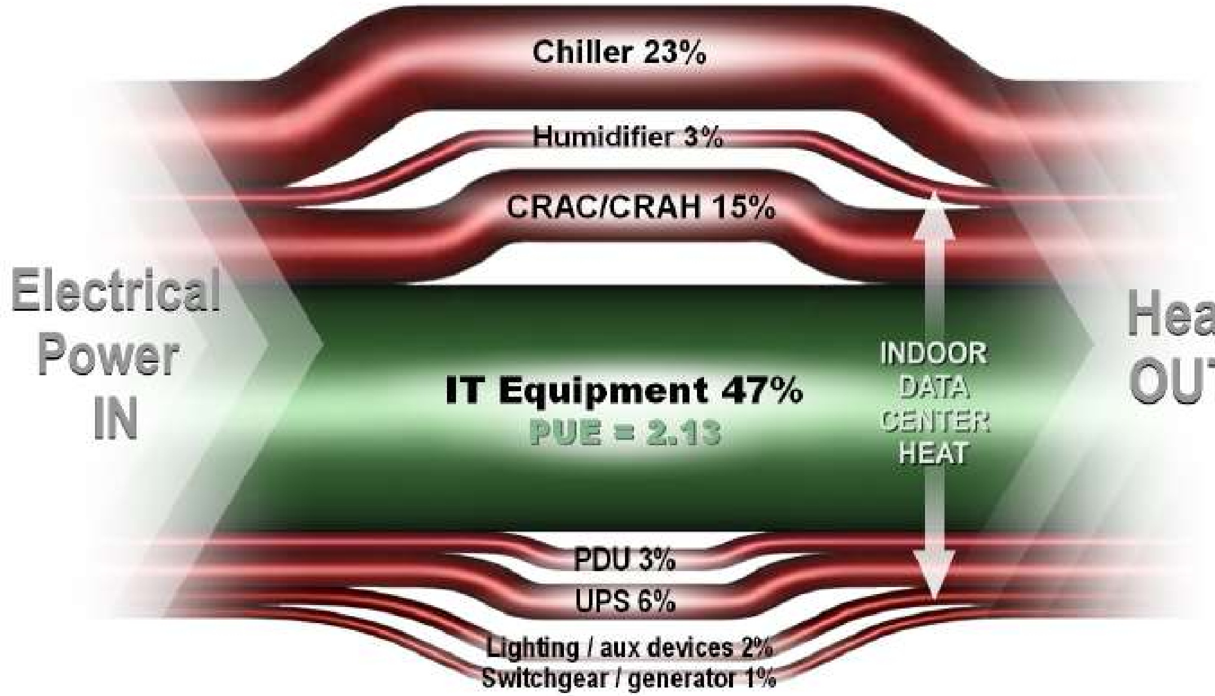

Figure 2

Power flow in a typical 2N data center

1. Inefficiencies of the power equipment

Equipment such as UPS, transformers, transfer switches, and wiring all consume some power (manifested as heat) while performing their function. While such equipment may have name-plate efficiency ratings that sound impressive – 90% or higher – these efficiency values are misleading and cannot be used to calculate the power wasted in real installations. When equipment is doubled for redundancy, or when the equipment is operated well below its rated power, efficiency falls dramatically. Furthermore, **the heat generated by this "wasted" energy in power equipment must be cooled by the cooling system**, which causes the air conditioning system to use even more electrical power.

2. Inefficiencies of the cooling equipment

Equipment such as air handlers, chillers, cooling towers, condensers, pumps, and dry coolers consume some power while performing their cooling function (that is, some of their input power is dispersed as heat instead of contributing to the mechanical work of cooling). In fact, the inefficiency (waste heat) of cooling equipment typically greatly exceeds the inefficiency (waste heat) of power equipment. When cooling equipment is doubled for redundancy or when the equipment is operated well below its rated power, efficiency falls dramatically. Therefore, **an increase in the efficiency of the cooling equipment directly benefits overall system efficiency**.

3. Power consumption of lighting

Lighting consumes power and generates heat. The heat generated by lighting must be cooled by the cooling system, which causes the air conditioning system to consume correspondingly more electrical power, even if the outdoor temperature is cold. When lighting remains on when there are no personnel in the data center, or when unutilized areas of the data center are lit, useless electrical consumption results. Therefore, increases in the efficiency of the lighting, or controlling lighting to be present only when and where needed, materially benefits overall system efficiency.

4. Over-sizing

Over-sizing is one of the largest drivers of electrical waste, but is the most difficult for users to understand or assess. Over-sizing of power and cooling equipment occurs whenever the design value of the power and cooling system exceeds the IT load. This condition can occur from any combination of the following factors:

- The IT load was overestimated and the power and cooling systems were sized for too large a load
- The IT load is being deployed over time, but the power and cooling systems are sized for a future larger load
- The cooling system design is poor, requiring over-sizing of the cooling equipment in order to successfully cool the IT load

5. Inefficiencies due to configuration

The physical configuration of the IT equipment can have a dramatic effect on the energy consumption of the cooling system. A poor configuration forces the cooling system to move much more air than the IT equipment actually requires. A poor configuration also causes the cooling system to generate cooler air than the IT equipment actually requires. Furthermore, physical configuration may force various cooling units into a conflict where one is dehumidifying while another is humidifying, a typically undiagnosed condition that dramatically reduces efficiency. The current trend of increasing power density in new and existing data centers greatly amplifies these inefficiencies. These configuration problems are present in virtually all operating data centers today and cause needless energy waste. Therefore, an architecture that systematically optimizes the physical configuration can dramatically reduce energy consumption.

An ideal and optimized data center

- Power and cooling equipment that is not currently needed should not be energized
- Over-sizing should be reduced wherever possible, so equipment can operate within the optimum region of its efficiency curve[1]
- Power, cooling, and lighting equipment should take advantage of the latest technologies to minimize power consumption
- Subsystems that must be used below their rated capacity (to support redundancy) should be optimized for that fractional-load efficiency, not for their full-load efficiency
- Capacity management tools should be used to minimize "stranded capacity" within the data center, allowing the maximum amount of IT equipment to be installed within the gross power and cooling envelope, pushing the system to the highest point on its efficiency curve
- Optimized, integrated physical configuration should be inherent *within* the system, and not tied to the characteristics of the room where it resides — for example, row-based cooling should be integrated with the IT racks, independent of room-based cooling
- The system should be instrumented to identify and warn about conditions that generate sub-optimal electrical consumption, so that they can be quickly corrected
- The system should include installation and operation tools and rules that maximize operating efficiency and minimize or eliminate the possibility of sub-optimal configuration or installation

Stranded capacity

- Stranded capacity is capacity that cannot be utilized by IT loads due to the design or configuration of the system. It is very costly that data centers can't meet the operational and capacity requirements of their initial design and it is a hurdle to get your data center green.
- The presence of stranded capacity indicates an imbalance between two or more of the following capacities:
  - Floor and rack space , Power, Power distribution , Cooling, Cooling distribution
- Most data centers are overcooled & inefficient waste of energy waste of scientific capacity.

Simple Cooling Load Estimation APC white paper #25

- Total IT load power in Watts - The sum of the power inputs of all the IT equipment.
- Power system rated power - The power rating of the UPS system. If a redundant system is used, do not include the capacity of the redundant UPS.

Table 2

Data center or network room heat output calculation worksheet

Item	Data required	Heat output calculation	Heat output subtotal
IT equipment	Total IT load power in Watts	Same as total IT load power in watts	____________ Watts
UPS with battery	Power system rated power in Watts	(0.04 x Power system rating) + (0.05 x Total IT load power)	____________ Watts
Power distribution	Power system rated power in Watts	(0.01 x Power system rating) + (0.02 x Total IT load power)	____________ Watts
Lighting	Floor area in square feet, or Floor area in square meters	2.0 x floor area (sq ft), or 21.53 x floor area (sq m)	____________ Watts
People	Max # of personnel in data center	100 x Max # of personnel	____________ Watts
Total	Subtotals from above	Sum of heat output subtotals	____________ Watts

Capacity Supply and Demand

APC white paper # 150

- While having power and cooling supply and demand information at the room or facility level helps, it does not provide sufficiently detailed information to answer the questions about specific IT equipment deployments. On the other hand, providing power and cooling supply and demand information at the IT device level is unnecessarily detailed and intractable. An effective and practical level at which to measure and budget power and cooling capacity is **at the rack level**.

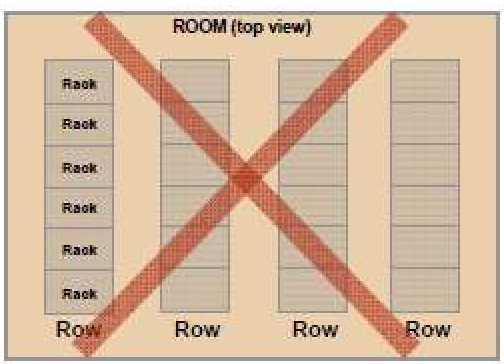

ROOM level
Too broad

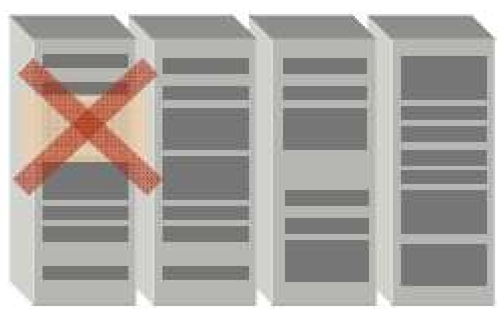

DEVICE level
Too specific

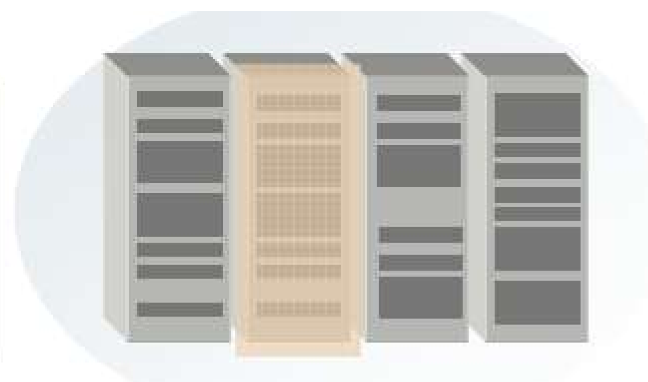

RACK level
Best for capacity management

Once the cooling requirements are determined, the following factors, must be considered:

- The size of the cooling load of the equipment (including power equipment)
- The size of the cooling load of the building
- Oversizing to account for humidification effects
- Oversizing to create redundancy
- Oversizing for future requirements
- The required oversizing for a CRAC unit therefore ranges from 0% for a small system with ducted exhaust air return, to 30% for a system with high levels of mixing within the room.

Avoiding Costs from Oversizing Data Center and Network Room Infrastructure, APC white paper #137

Table 1

Definitions related to oversizing

Term	Definition
Design lifetime	The overall planned life of the data center. Typically 10-15 years.
Design capacity	The maximum IT load the data center is ultimately capable of supporting. All or part of the power and cooling equipment needed to provide this capacity may be installed at start-up.
Installed capacity	The load capability of the power and cooling equipment installed. Equal to or less then the design capacity.
Expected load	The estimated IT load at the commissioning of the system and over its lifetime. This typically changes with time and increases from time of commissioning.
Actual load	The actual IT load at the commissioning of the system and over its lifetime. This typically changes with time and increases from time of commissioning.

Figure 1

Design capacity and expected load requirement over the lifetime of a data

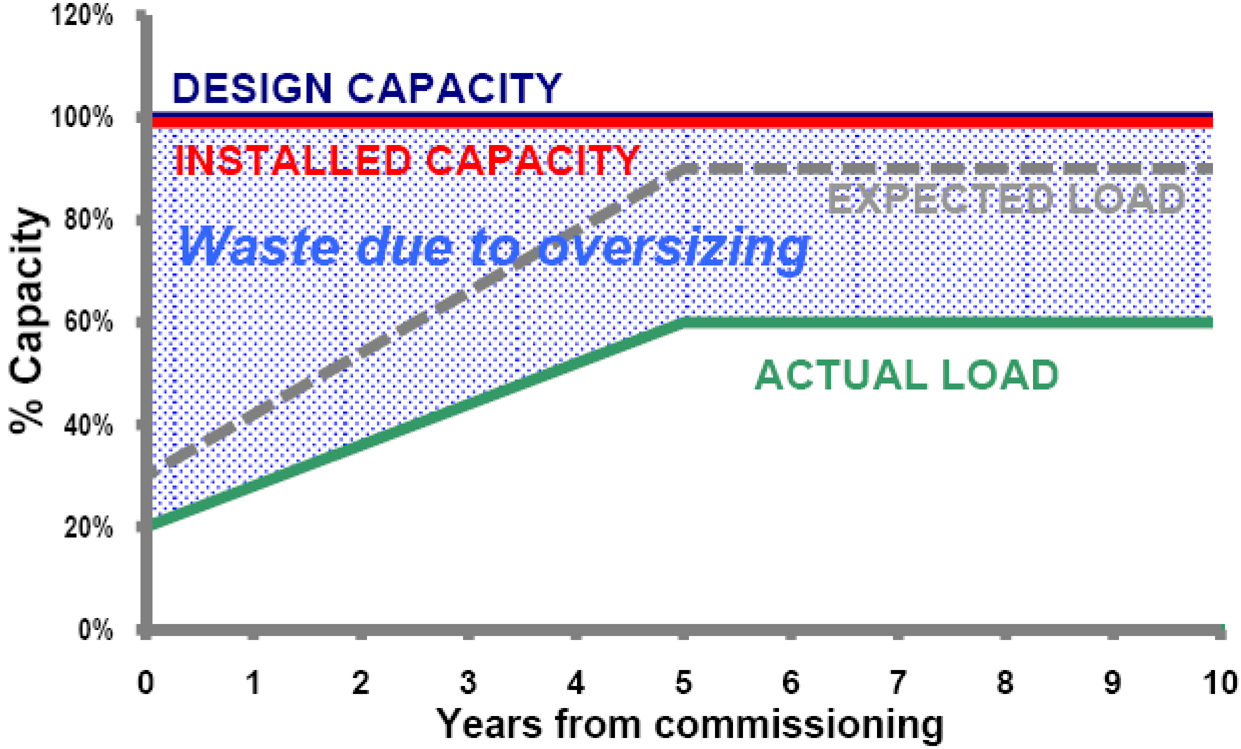

It was found that the start-up expected load is typically 30% of the ultimate design capacity and that the ultimate expected load is typically 80-90% of the ultimate design capacity allowing for a safety margin). It was further found that the start-up actual load is typically 20% of the ultimate design capacity, and that the ultimate actual load is typically about 60% of the design capacity.

国立交通大學
National Chiao Tung University

Figure 3

Design power capacity and requirement over the lifetime of a data center

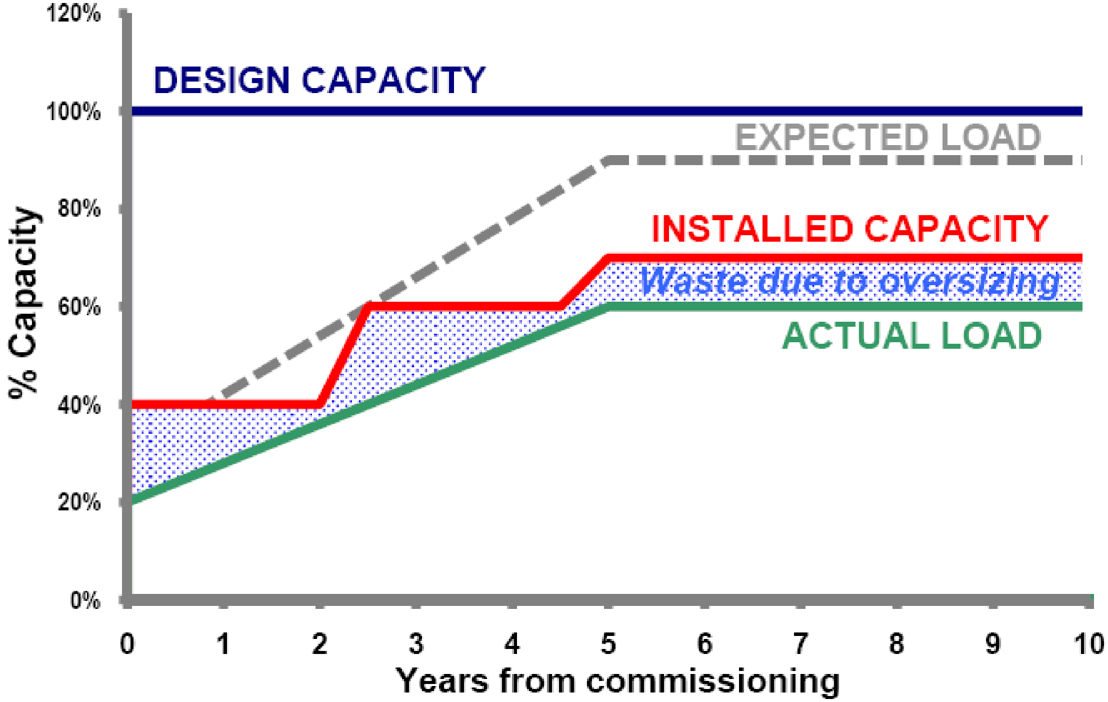

APC white paper #143

Table 2

Parameters of the growth model

	Growth model parameter	Meaning
IT load profile	1 MAXIMUM final load	Maximum anticipated IT load
	2 MINIMUM final load	Minimum anticipated IT load
	3 INITIAL load	IT load of initial installation
	4 Ramp-up time	The time it takes to go from initial load to final load
System capacity plan	5 Step size	Incremental step size of the physical infrastructure system, if full buildout is deferred
	6 Margin	Extra capacity to cover the unexpected – either an unexpected addition to the IT load or an unexpected drain on system capacity

Figure 4

Complete growth model

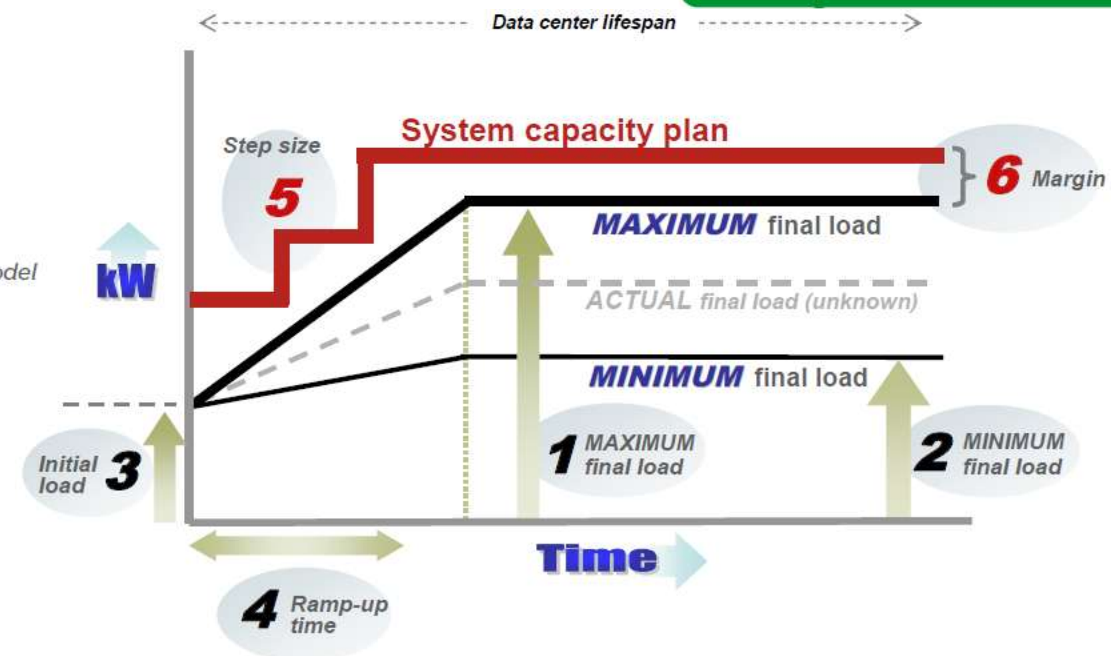

For a particular project, the growth model is developed in two parts

- ***First, develop the IT load profile***. The **IT load profile**, consisting of parameters 1-4 of the model, is created early in the planning process, **based on an understanding of the organization's business needs**. In some cases, this may require consulting expertise from someone familiar with the organization's business and general IT issues, or reference to standard profiles describing the IT growth parameters of similar organizations. The key at this step is for the participants of the planning process to develop a shared view of the projected IT load.

- ***Second, develop the system capacity plan to support the IT load profile***. The **system capacity plan** is represented by parameters 5 and 6 of the growth model (**step size** and **margin**). Development of the system capacity plan is begun early in the planning sequence, with a rough estimate of step size that will guide the choice of reference design. The system capacity plan is finalized later in the planning sequence, after the basic system architecture and the floor plan (row layout of the room) have been determined. The user will typically not have expertise in this area, and so will rely upon the equipment vendor or other qualified consulting services. Incremental phase-in steps provide the option to delay, adjust, or cancel full build out based on actual conditions as they develop during the ramp-up time. The benefits of a stepped phase-in are discussed later in this paper in the section **“The value of stepped phase-in”**.

Technology to implement the design principles *(APC white paper #126)*

- **Scalable power and cooling**, to avoid over-sizing
- **Row-based cooling**, to improve cooling efficiency
- **High-efficiency UPS**, to improve power efficiency
- **415/240 V AC power distribution**, to improve power efficiency
- **Variable-speed drives on pumps and chillers**, to improve efficiency at partial load and on cool days
- **Capacity management tools**, to improve utilization of power, cooling, and rack capacity
- **Room layout tools**, to optimize room layout for cooling efficiency

APC White Paper #114

Practical strategies for reducing electrical power consumption for data centers, indicating range of achievable electrical savings

	Savings	Guidance	Limitations
Right-size DCPI	10 – 30%	• Using a modular, scalable power and cooling architecture • Savings are greater for redundant systems	• For new designs and some expansions
Virtualize servers	10 – 40%	• Not technically a physical infrastructure solution but has radical impact • Involves consolidation of applications onto fewer servers, typically blade servers • Also frees up power and cooling capacity for expansion	• Requires major IT process changes • To achieve savings in an existing facility some power and cooling devices may need to be turned off
More efficient air conditioner architecture	7 – 15%	• Row-oriented cooling has higher efficiency for high density (White Paper 130) • Shorter air paths require less fan power • CRAC supply and return temperatures are higher, increasing efficiency, capacity, and preventing dehumidification thereby greatly reducing humidification costs	• For new designs • Benefits are limited to high density designs
Economizer modes of air conditioners	4 – 15%	• Many air conditioners offer economizer options • This can offer substantial energy savings, depending on geographic location • Some data centers have air conditioners with economizer modes, but economizer operation is disabled	• For new designs • Difficult to retrofit
More efficient floor layout	5 – 12%	• Floor layout has a large effect on the efficiency of the air conditioning system • Involves hot-aisle / cold-aisle arrangement with suitable air conditioner locations (White Paper 122)	• For new designs • Difficult to retrofit

Practical strategies for reducing electrical power consumption for data centers, indicating range of achievable electrical savings

Coordinate air conditioners	0 – 10%	•Many data centers have multiple air conditioners that actually fight each other •One may actually heat while another cools •One may dehumidify while another humidifies •The result is gross waste •May require a professional assessment to diagnose	•For any data center with multiple air conditioners
Locate vented floor tiles correctly	1 -6%	•Many vented tiles are located incorrectly in the average data center or the wrong number are installed •Correct locations are NOT intuitively obvious •A professional assessment can ensure an optimal result •Side benefit – reduced hot spots	•Only for data centers using a raised floor •Easy, but requires expert guidance to achieve best result
Install energy efficient lighting	1 – 3%	•Turn off some or all lights based on time of day or motion •Use more efficient lighting technology •Don't forget that lighting power also must be cooled, doubling the cost •Benefit is larger on low density or partly filled data centers	•Most data centers can benefit
Install blanking panels	1 – 2%	•Decrease server inlet temperature •Also saves on energy by increasing the CRAC return air temperature •Cheap and easy with new snap-in blanking panels such as those by Schneider Electric	•For any data center, old or new

Use high efficiency UPS

APC White Paper #126

Figure 8

UPS efficiency as a function of load comparing latest generation UPS to historic published data

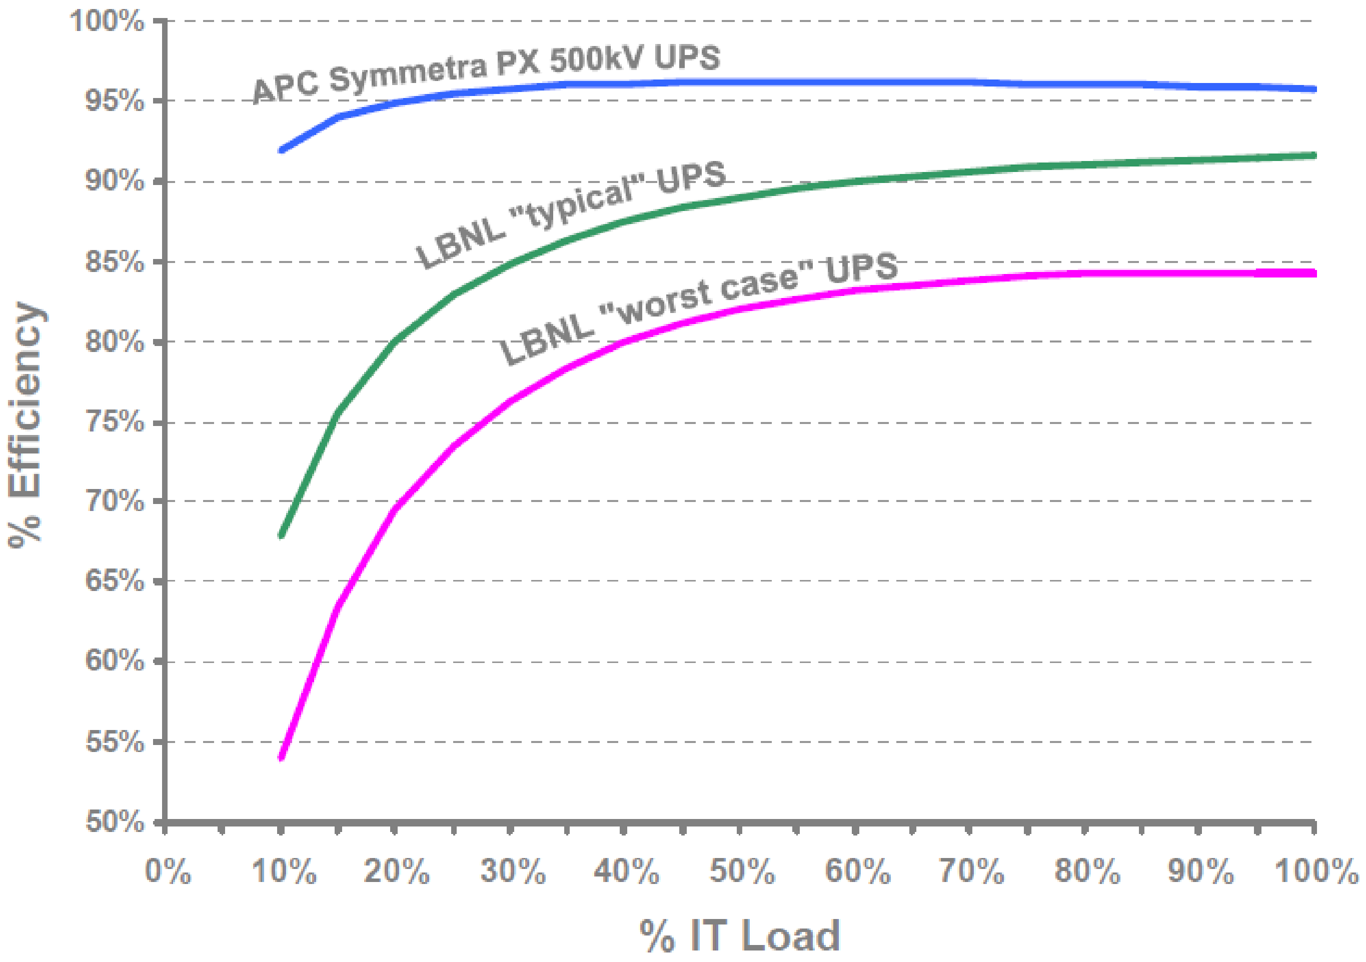

Use variable speed Chiller & Pump

Pumps and chillers equipped with variable-speed drives (VFDs) and appropriate controls can reduce their speed and energy consumption to match the current IT load and the current outdoor conditions. The energy improvement varies depending on conditions, but can be as large as 10% or more, especially for data centers that are not operating at full rated IT load, or for data centers with chiller or pump redundancy. Variable-speed drives on pumps and chillers can be considered a form of "automatic rightsizing."

Capacity Management

- Shorter airflow paths, resulting in the need for less fan horsepower
- Less air mixing, resulting in higher heat rejection temperatures
- Higher heat rejection temperatures, resulting in improved chiller efficiency
- Higher heat rejection temperatures, resulting in increased air conditioner capacity
- Shorter wiring lengths, resulting in lower wiring and PDU losses
- More IT load can be powered by the same power and cooling infrastructure

APC White Paper #127

The power path is divided into three segments:

- UPS
- Distribution wiring
- IT device power supplies (PSUs)

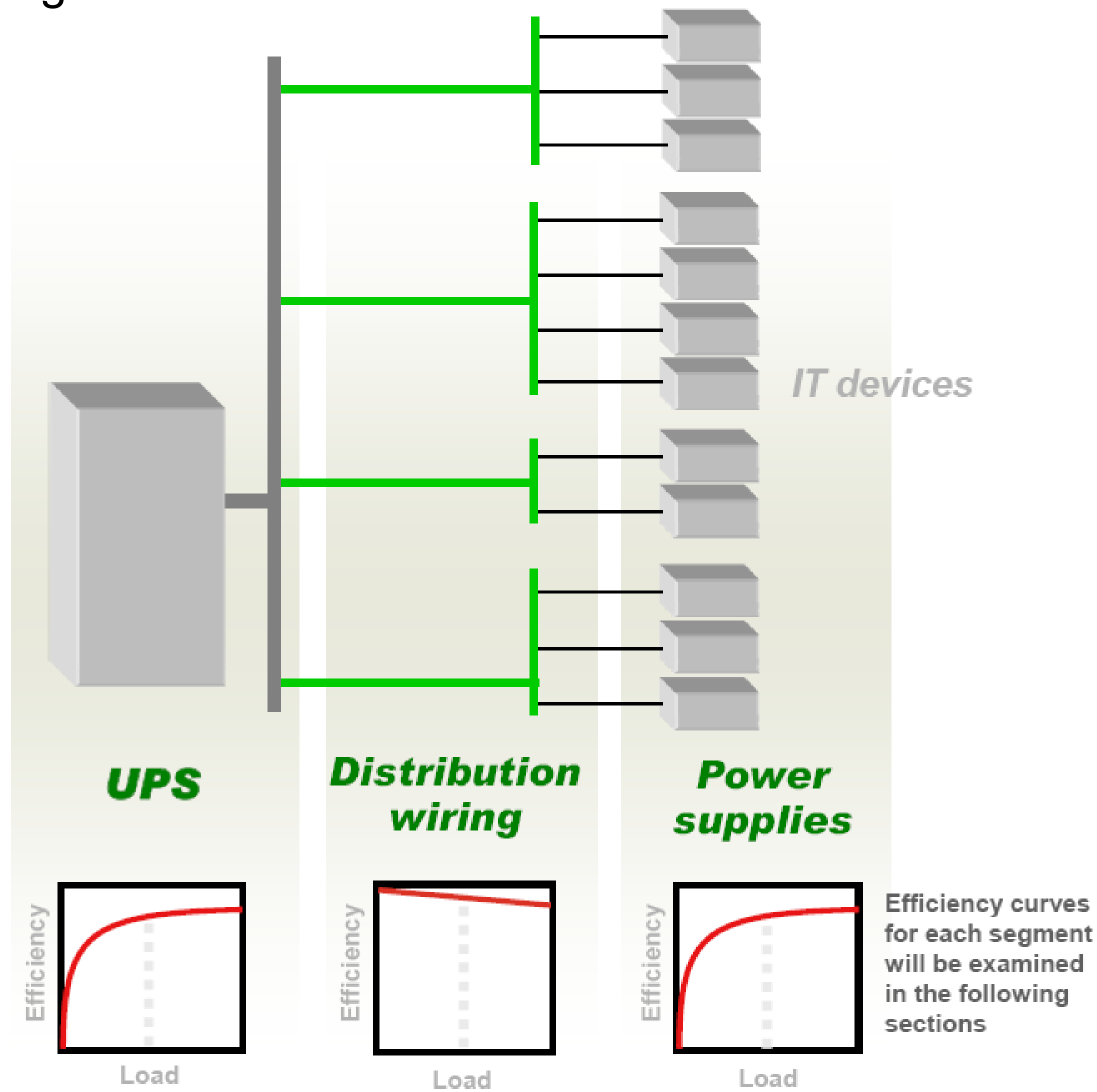

Figure 3

Data center power path: three segments, three efficiency curves

DC distribution has higher efficiency based on :

- It may be possible to build a DC UPS that is higher in efficiency than an AC UPS
- The elimination of power distribution unit (PDU) transformers will reduce electrical losses
- It may be possible to improve the efficiency of the IT equipment power supply itself, beyond the improvements possible in an AC input design

Figure 4

Efficiency of several commercially available AC and DC UPS systems

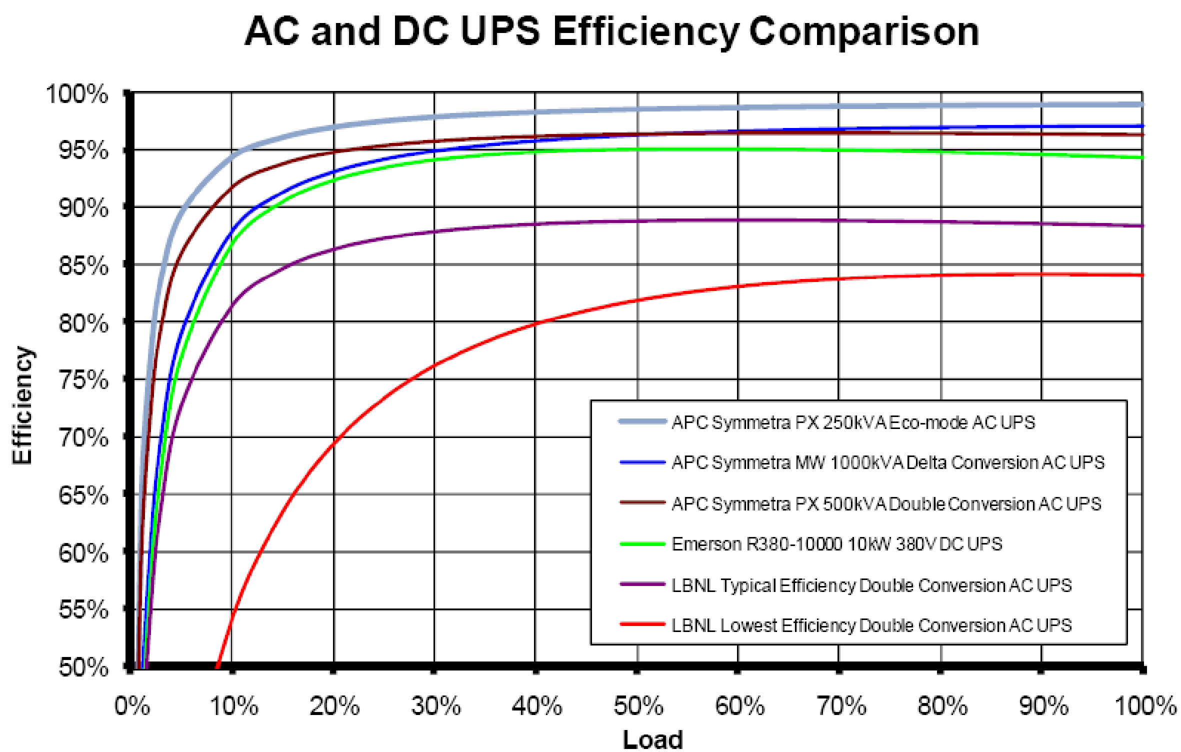

Figure 6

Distribution wiring efficiency

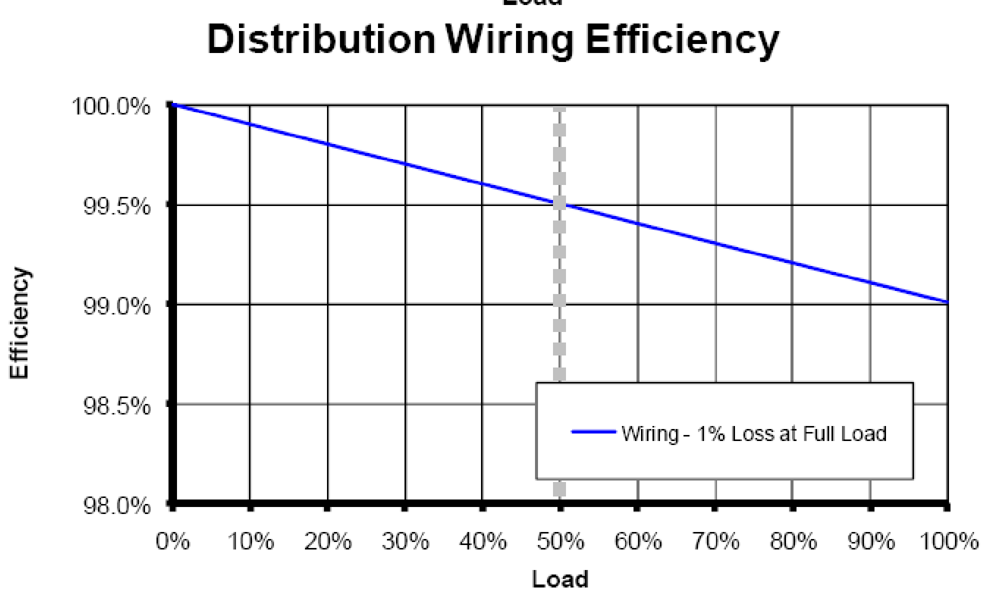

Use high voltage & high efficiency Power supply **(APC White Paper #128)**

- **Figure 3** illustrates a 1-3 percentage point improvement in power supply efficiency with increased voltage. This represents a savings of $4 to $31 per year per server in electrical cost by increasing the voltage from 120 to 240. In addition, approximately $2 to $16 per year per server is saved on air conditioning operating costs.

Figure 3

Power supply efficiency (source: HP)

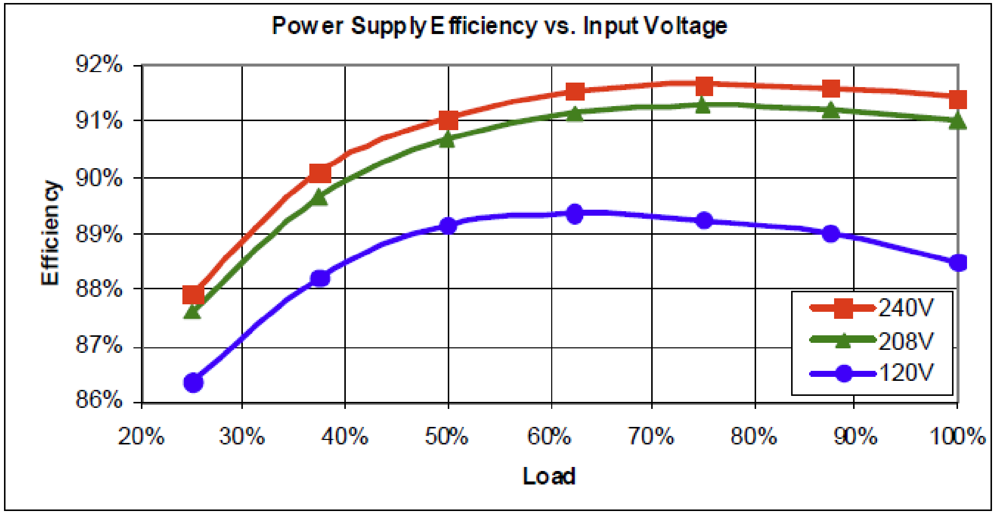

General benefits to increasing distribution voltage from 120 to 240

- More power capacity given the same branch circuit current
- Less current required given the same branch circuit power capacity
- Higher power distribution efficiency / lower energy cost
- Higher IT equipment power supply efficiency / lower energy cost
- Lower copper material cost (less copper required)

DC power distribution

Table 1

Data center efficiency improvement from DC distribution, compared to conventional design and to this paper's high-efficiency architecture

DC distribution element	Savings compared to conventional design	Savings compared to high-efficiency architecture described in this paper
Elimination of transformers	5-10%	None
Replacing UPS with AC-DC converter	5-15%	None
New IT equipment that accepts high voltage DC input	4%	2%
TOTAL IMPROVEMENT	**13-28%**	**2%**

ΔT across IT equipment

APC white paper #139

Table 1

ΔT by type of IT equipment

Type of IT	ΔT(˚F)	ΔT(˚C)
Blade servers	40	22
1U – 2U servers	30	17
Other	20	11

1. Cooling system – a health check

APC white paper #42

- Maximum Cooling Capacity. If there isn’t enough gas in the tank to power the engine then no amount of tweaking will improve the situation.
- CRAC (computer room air conditioning) units. Measured supply and return temperatures and humidity readings must be consistent with design values. Ensure that all filters are clean.
- Chiller water/ condenser loop. Check condition of the chillers and/or external condensers, pumping systems, and primary cooling loops. Ensure that all valves are operating correctly. Check that DX systems, if used, are fully charged.

- Room temperatures. Check temperature at strategic positions in the aisles of the data center. These measuring positions should generally be centered between equipment rows and spaced approximately every fourth rack position.
- Rack temperatures. Measuring points should at the center of the air intakes at the bottom, middle, and top of each rack. These temperatures should be recorded and compared with the manufacturer’s recommended intake temperatures for the IT equipment.
- Tile air velocity. If a raised floor is used as a cooling plenum, air velocity should be uniform across all perforated tiles or floor grilles.

- Condition of subfloors. Any dirt and dust present below the raised floor will be blown up through vented floor tiles and drawn into the IT equipment. Under-floor obstructions such as network and power cables obstruct airflow and have an adverse effect on the cooling supply to the racks.
- Airflow within racks. Gaps within racks (unused rack space without blanking panels, empty blade slots without blanking blades, unsealed cable openings) or excess cabling will affect cooling performance.
- Aisle & floor tile arrangement. Effective use of the subfloor as a cooling plenum critically depends upon the arrangement of floor vents and positioning of CRAC units.

2. Common causes for "poor cooling"

www.upsite.com

- Dirty or blocked coils choking airflow
- Undercharged DX systems
- Incorrectly located control points
- Un-calibrated or damaged sensors
- Reversed supply & return piping
- Faulty valves
- Faulty pumps
- Pumps running unnecessarily
- Free cooling systems not initiated

3. Install blanking panels and implement cable management regime

Figure 2

Effect of installation of blanking panels on server air inlet temperature

2A. (left)

Without blanking panels

2B. (right)

With blanking panels

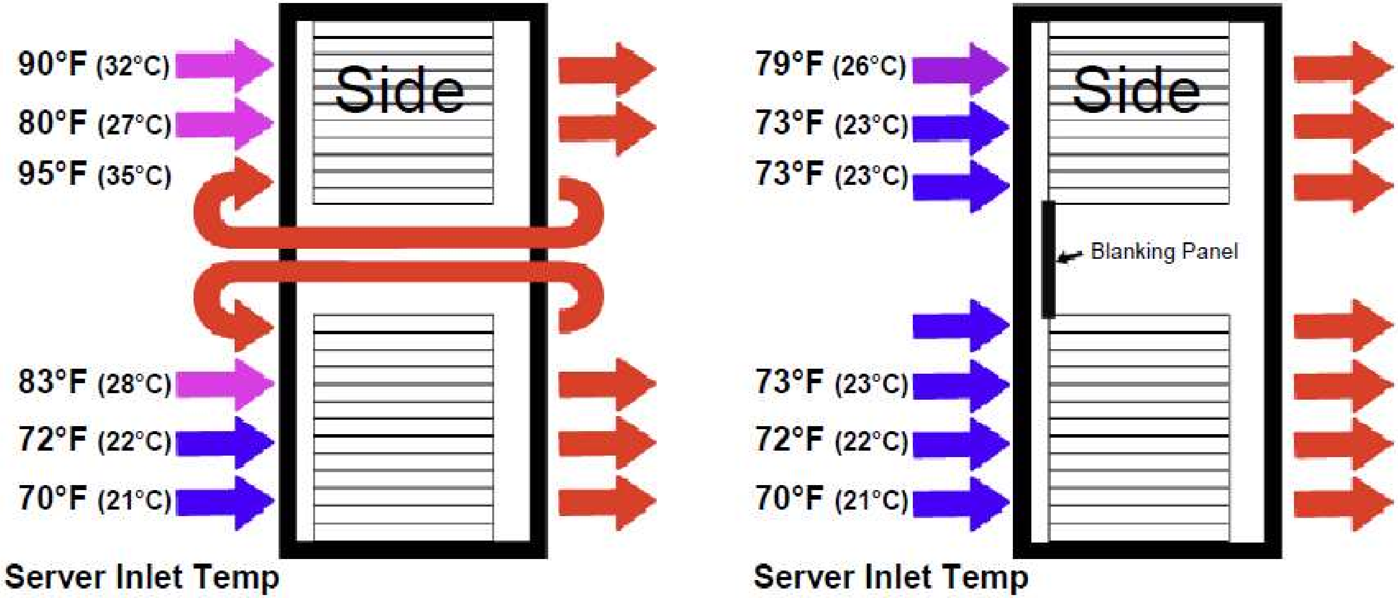

Figure 3

Example of unstructured cabling

- 4. Remove sub-floor blockages and seal floor
- 5. Separate high-density racks

Figure 5

Data center with all high-density racks together

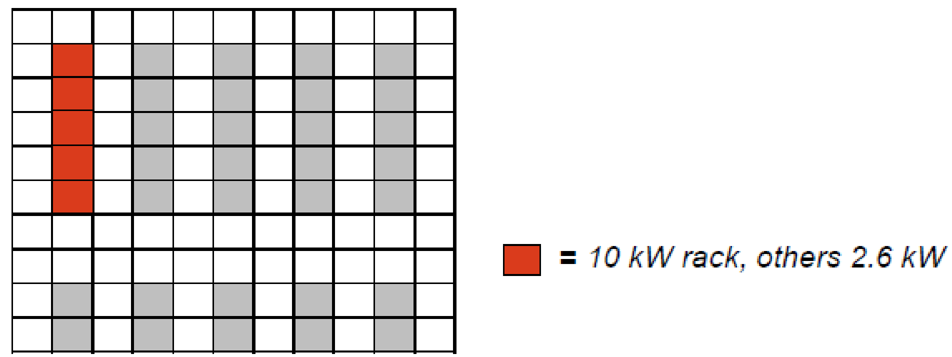

Figure 6

Data center with high-density racks spread out

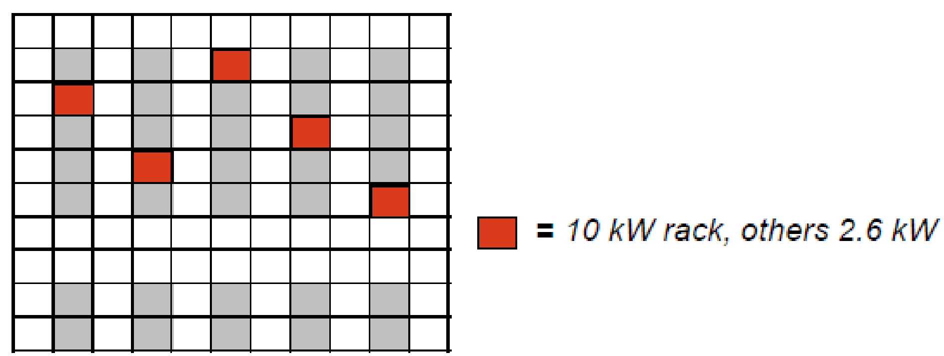

6. Set up hot-aisle/cold-aisle

Figure 7

Rack arrangement with no separation of hot or cold aisles

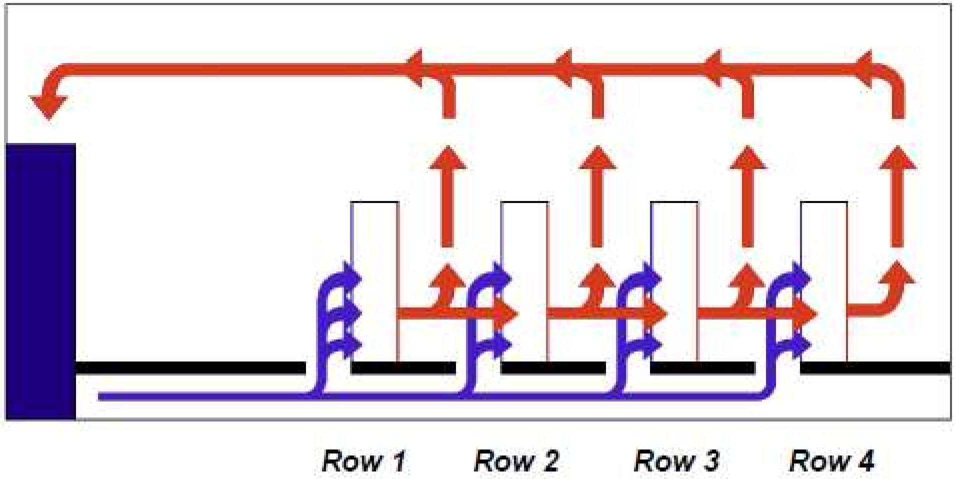

Figure 8

Hot aisle / cold aisle rack arrangement

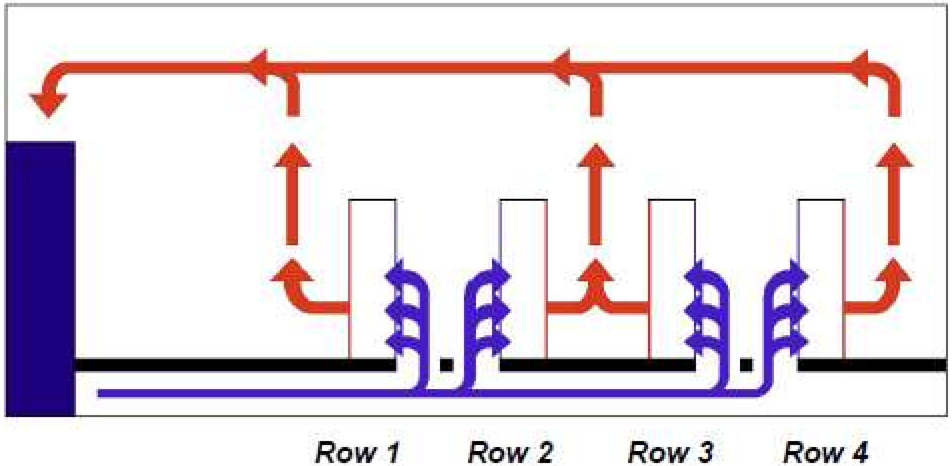

7. Align CRACs with hot aisles

7. Align CRACs with hot aisles

CRAC units must be aligned with hot aisles to optimize cooling efficiency. **Figure 9** shows a typical room layout where CRAC units have been evenly placed around the perimeter of the room to service a hot aisle/cold aisle arrangement.

In this example, the CRAC units along the two side walls are too close to the cold aisle, which causes the airflow to bypass the floor vents in that aisle. These CRAC units would be better positioned along the top and bottom walls to get better airflow along the aisles.

In **Figure 10** the CRAC units have been moved to the top and bottom walls and are now aligned with the hot aisles. Conventional wisdom would indicate that CRACs should be aligned with cold aisles to generate a flow of air to the floor vents. However, CFD (computational fluid dynamics) analysis has shown that hot air from the hot aisles crosses over the racks into the cold aisle when returning to the CRACs, causing a mixing of hot and cold air that increases the temperature of supply air to the rack fronts.

In summary, with a raised-floor cooling system it is more important to align CRAC units with the air return path (hot aisles) than with the subfloor air *supply* path (cold aisles).

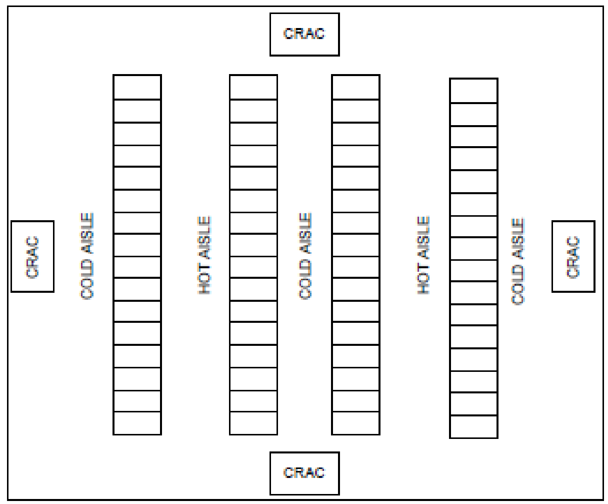

Figure 9

Typical CRAC

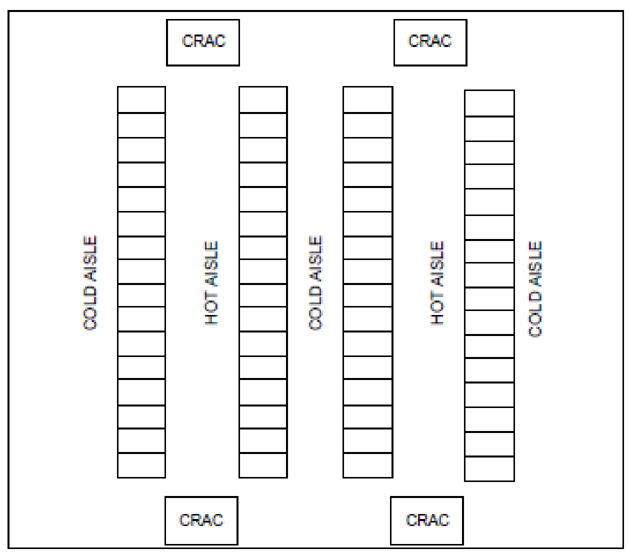

Figure 10

CRACs aligned with aisles

• 8. Manage floor vents

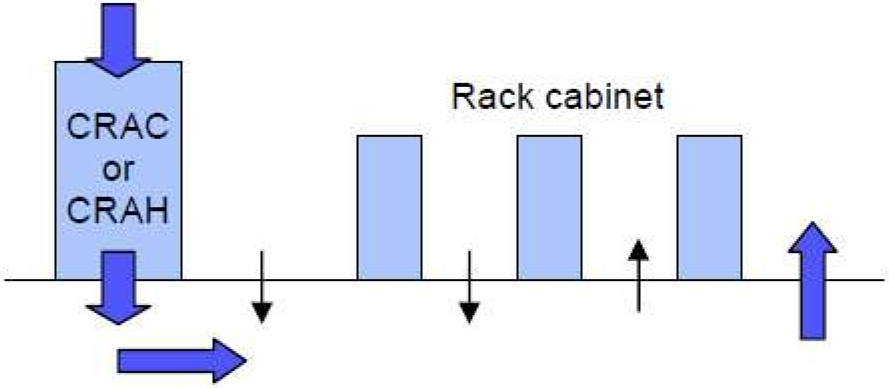

Figure 11

Relative air movement in high-velocity underfloor environments

9. Install row based cooling architecture

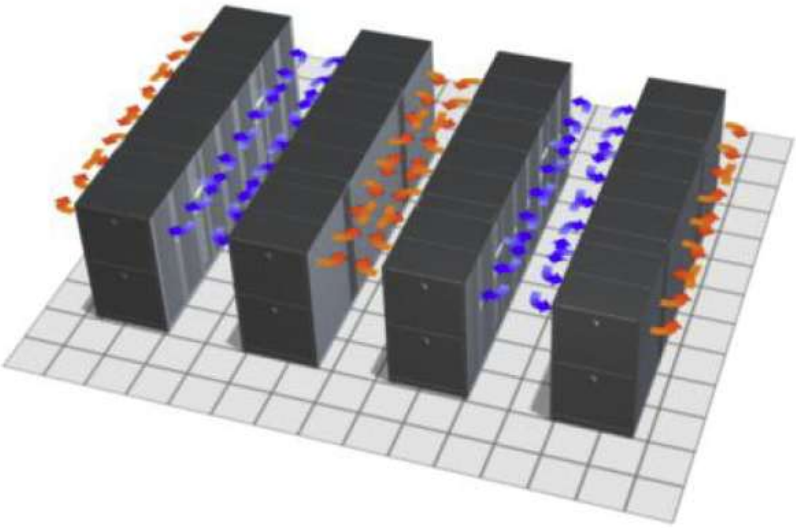

10. Install air flow assist devices

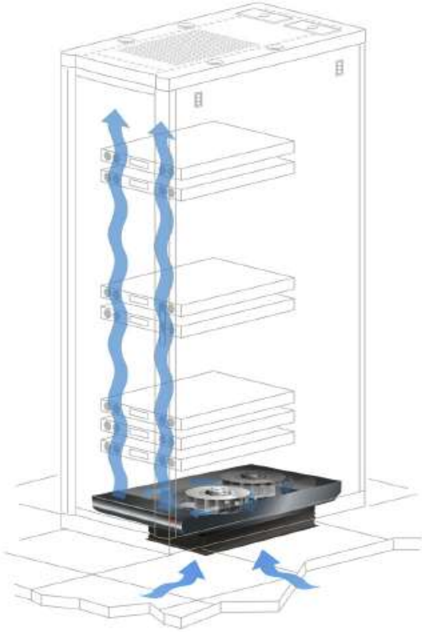

Figure 12

Rack-mounted fully ducted air supply unit

Fan-tray devices, such as Air Distribution Unit (ADU), fit into the rack's bottom U spaces and direct the airflow vertically to create a cold air "curtain" between the front door and the servers. Blanking panels (see solution #3 earlier in this paper) must be used to ensure the integrity of this newly created plenum.

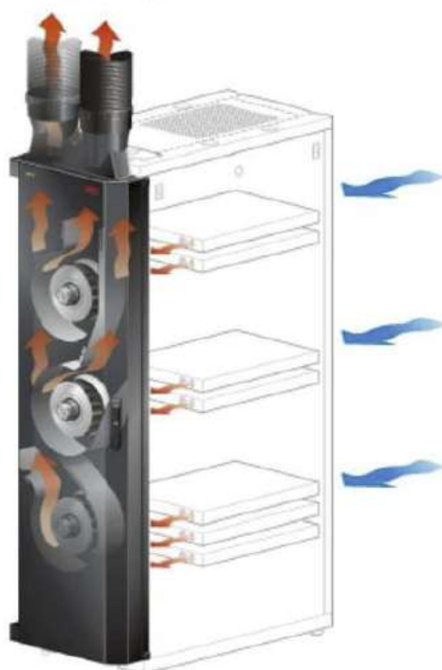

Figure 13

Rack-mounted fully ducted air return unit

For higher densities, the rear door of the cabinet can be removed and replaced with an air-moving device such as Air Removal Unit (ARU). Hot exhaust air that would normally be expelled into the hot aisle is gathered and propelled upwards, where it is ducted into the return air plenum. This eliminates recirculation at the rack and improves CRAC efficiency and capacity. *Blanking panels and rack side panels must be used with these devices.*

APC white paper #118

Figure 3

Typical data center infrastructure efficiency curve

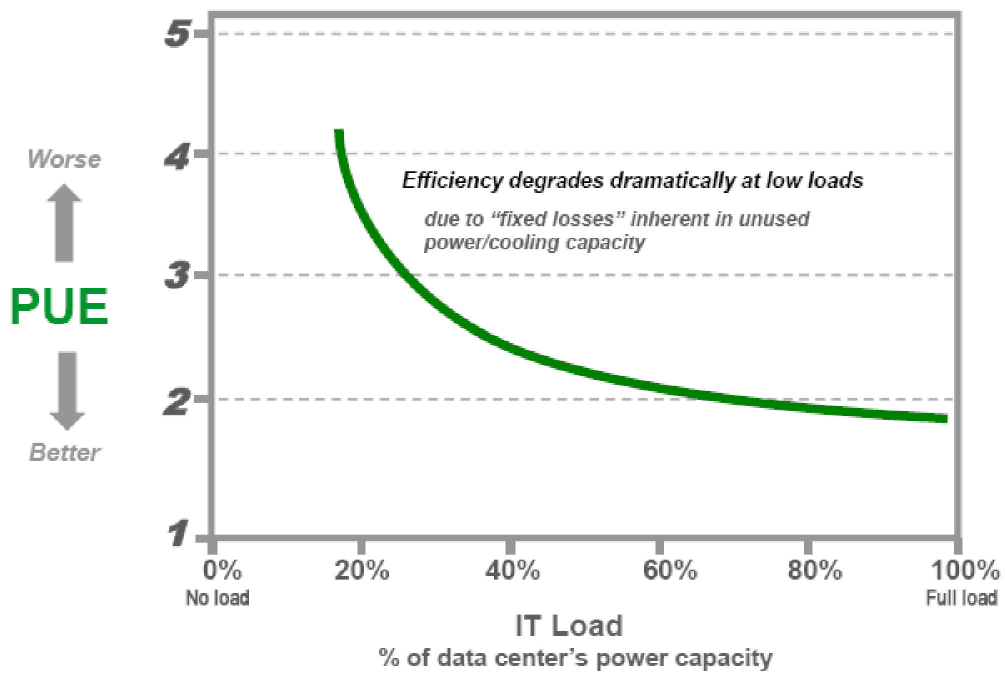

Floor Management
APC white paper #144

- Control of airflow using hot-aisle/cold-aisle rack layout

Figure 2

Basic hot-aisle/cold-aisle data center equipment layout plan

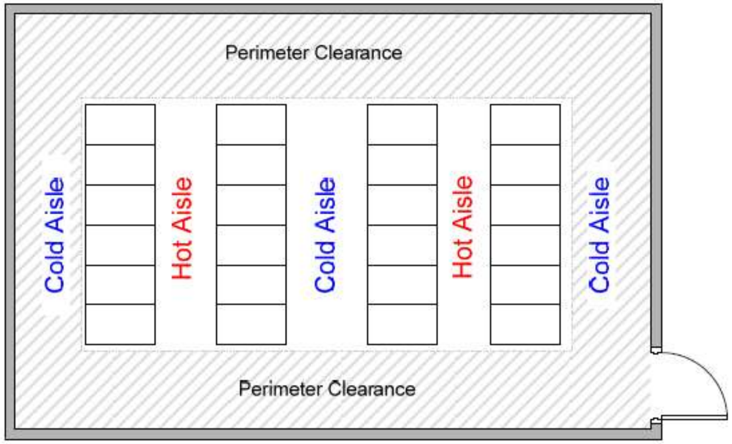

Figure 5

Pitch of a row layout

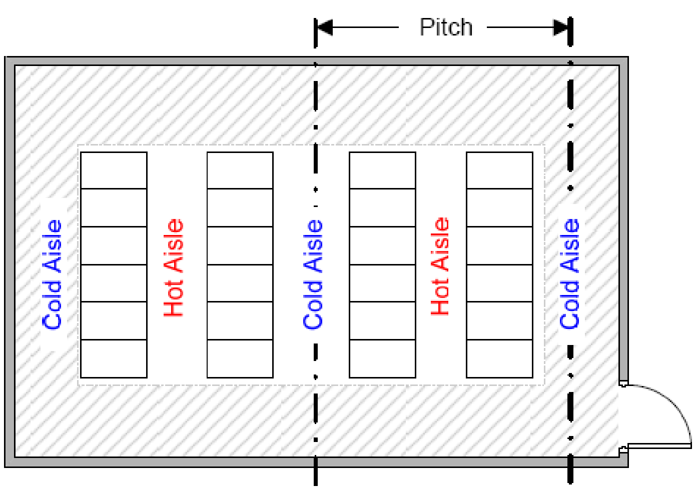

Figure 6

The four standard pitches of row layouts

A.

Compact

B.

Wide hot-aisle

C.

Wide cold-aisle

D.

Wide hot-aisle & cold aisle

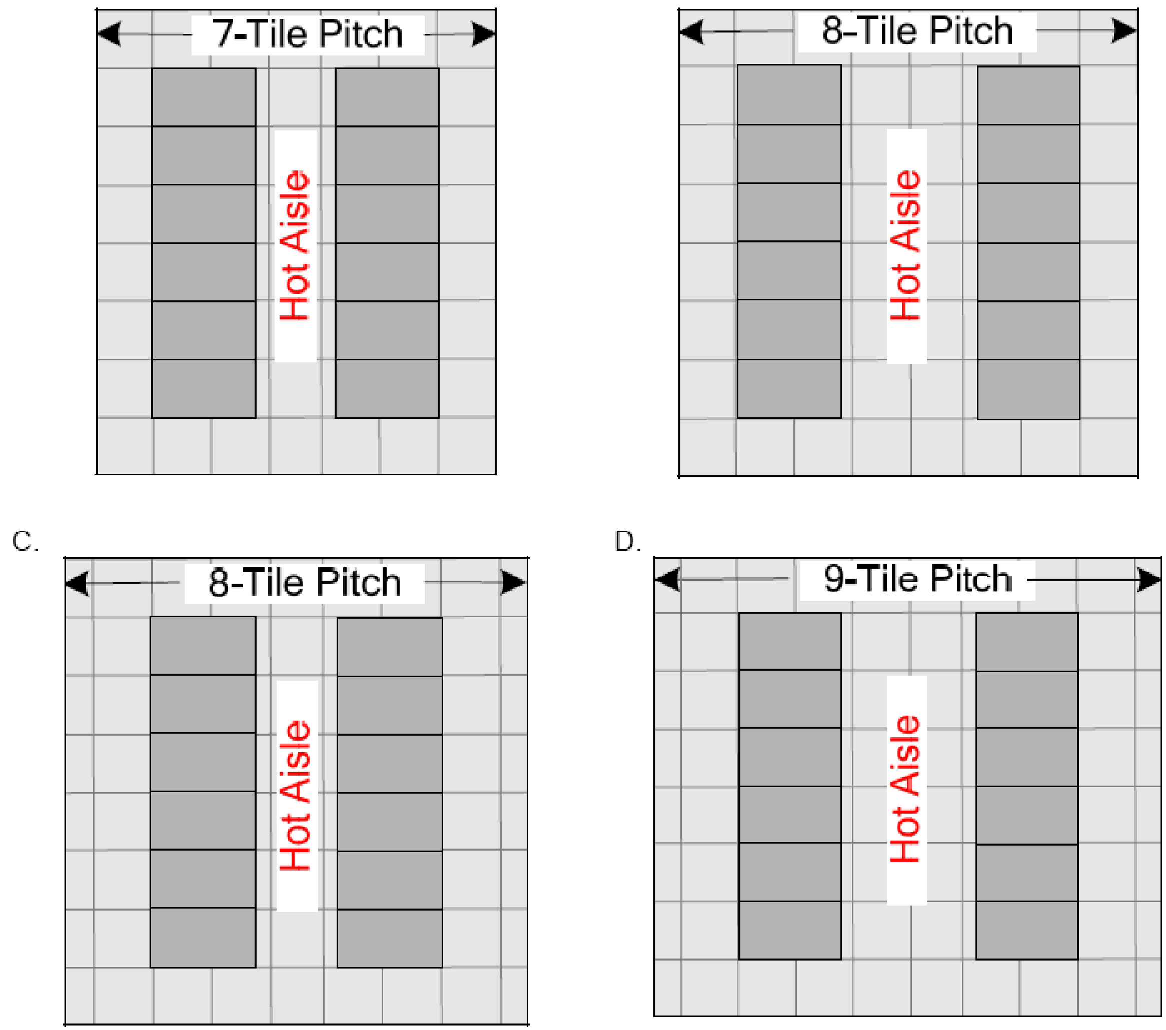

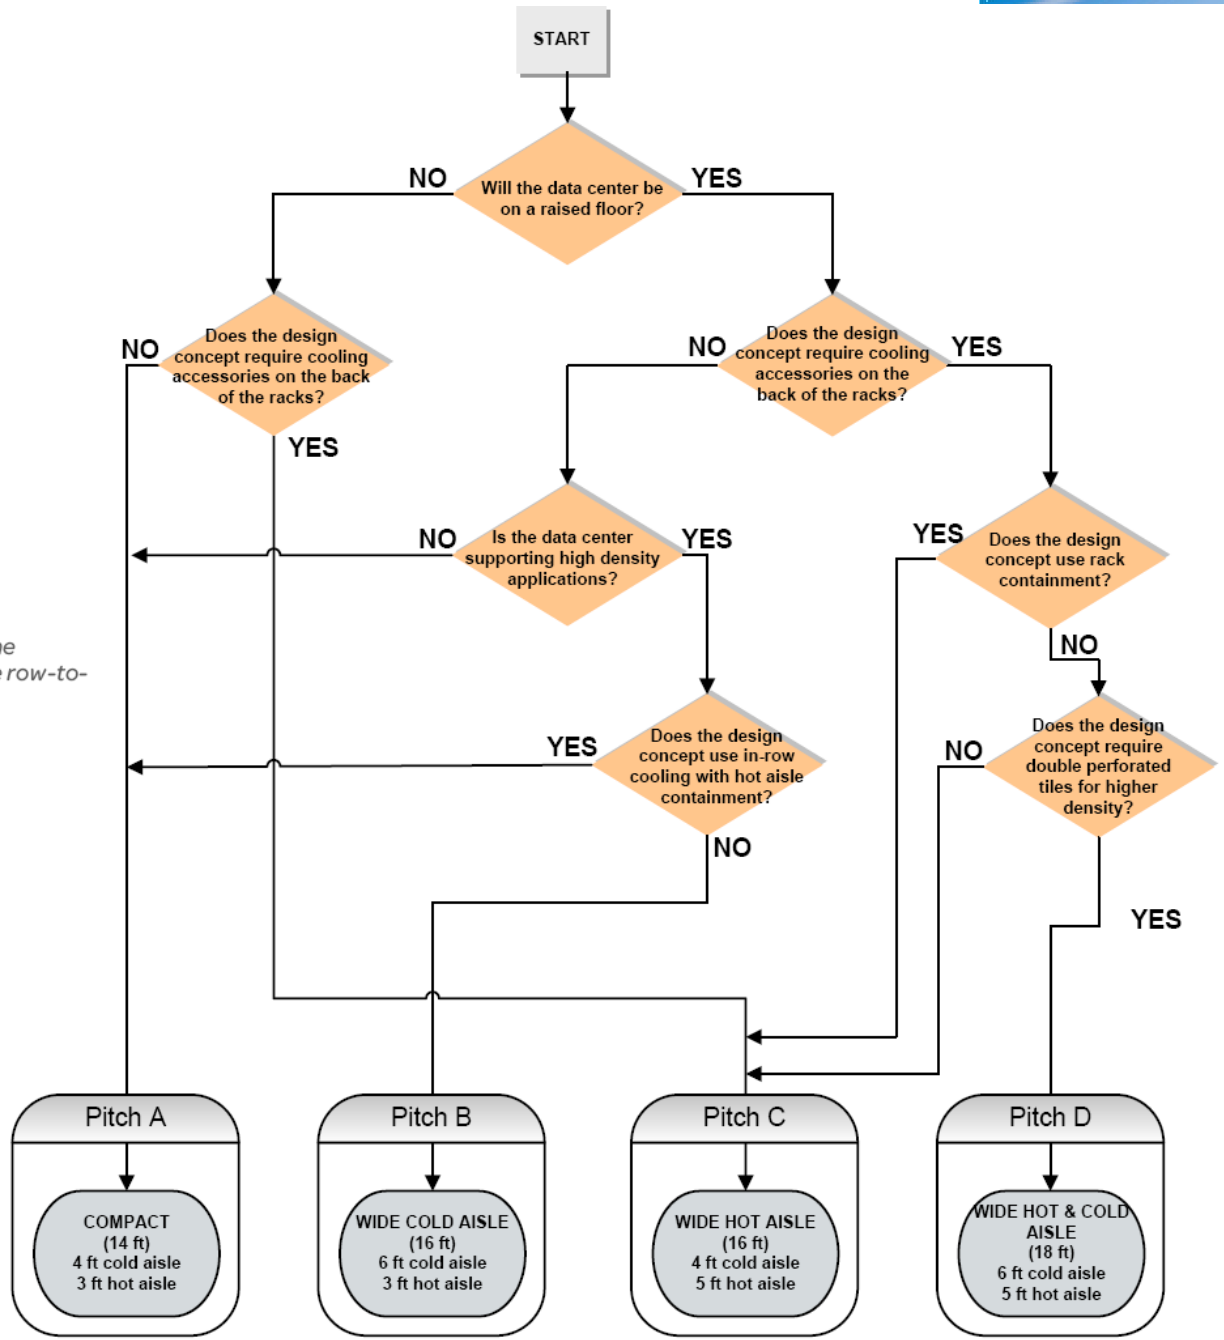

Figure 7

Choosing the appropriate row-to-row pitch

Provide access ways that are safe and convenient

- The room in Figure 3A fits 40 rack locations when no columns are present. When a column exists, but aligns with a row of equipment racks, as Figure 3B shows, only one rack may be impacted.

Figure 3

Sample impact of columns on number of rack locations when column aligns with row

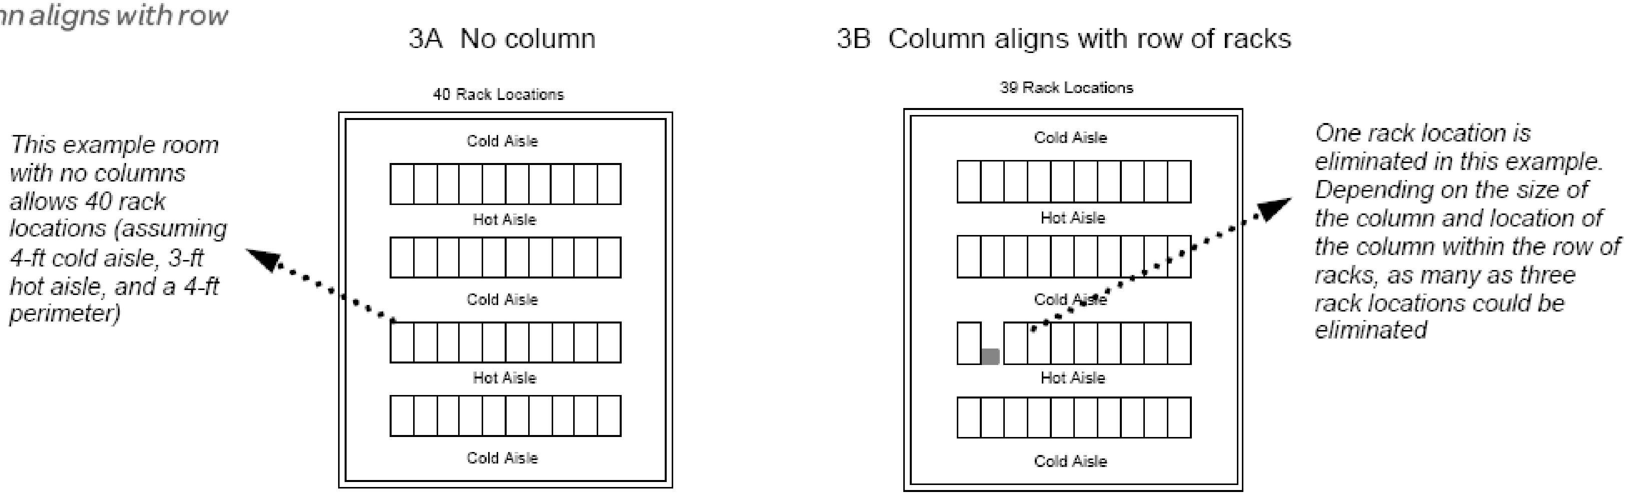

Figure 4

Sample impact of columns on number of rack locations when column aligns with aisle

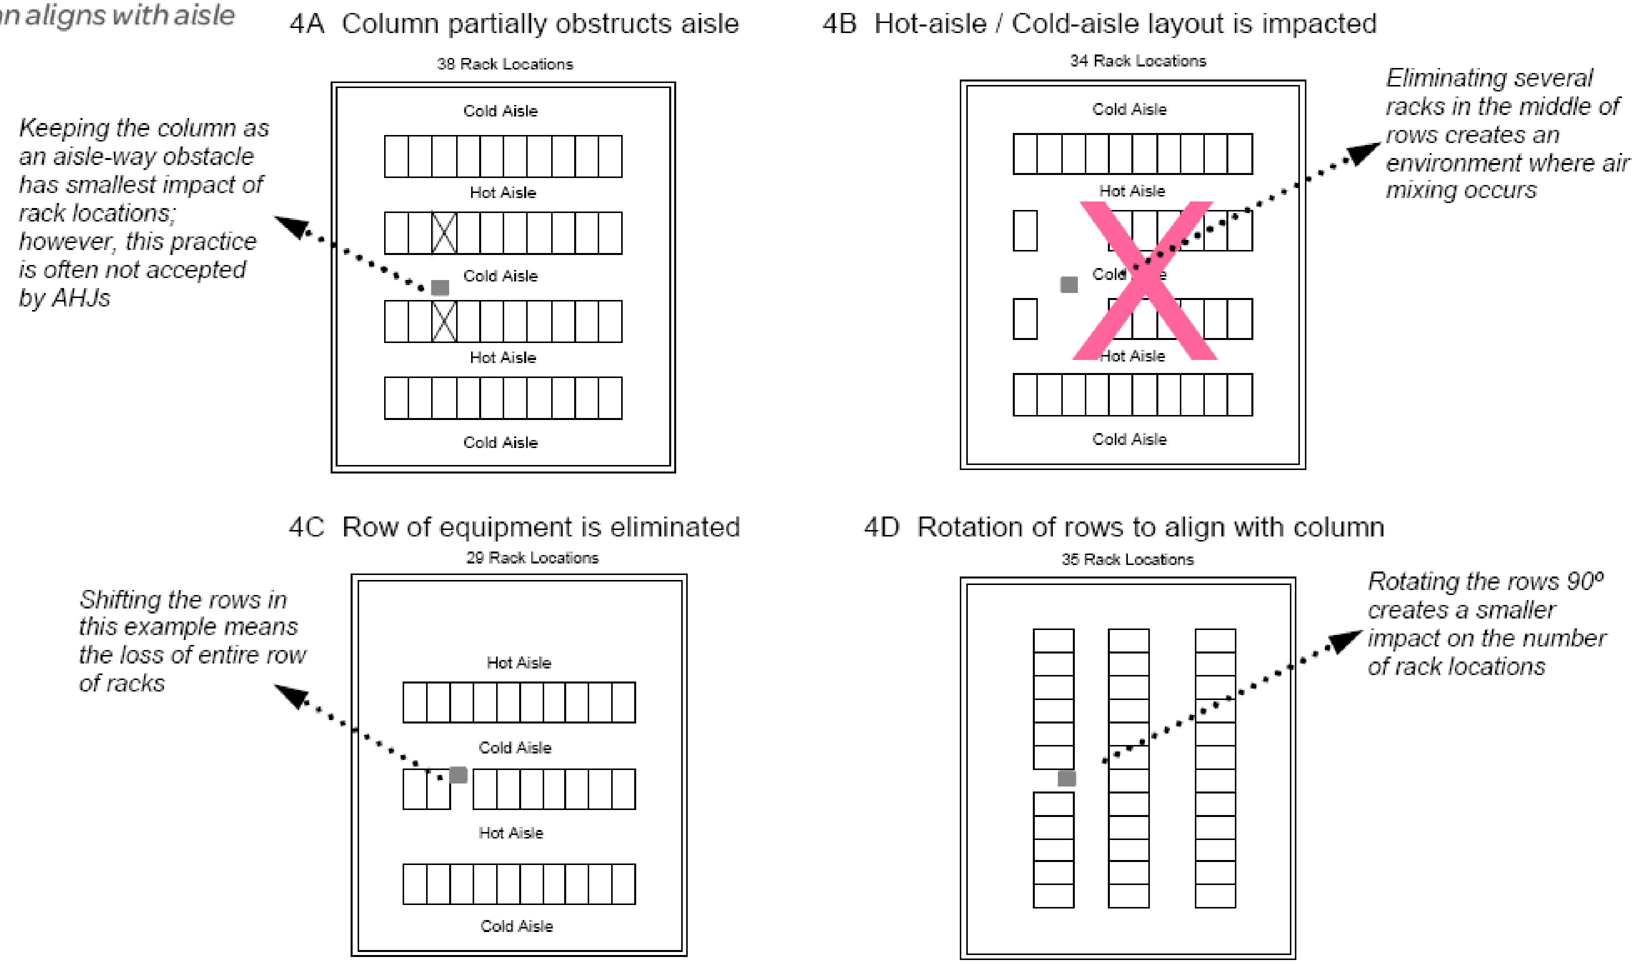

Basic principles of structural room layouts – Standard layout

Figure 8

Impact of room dimension on number of rows

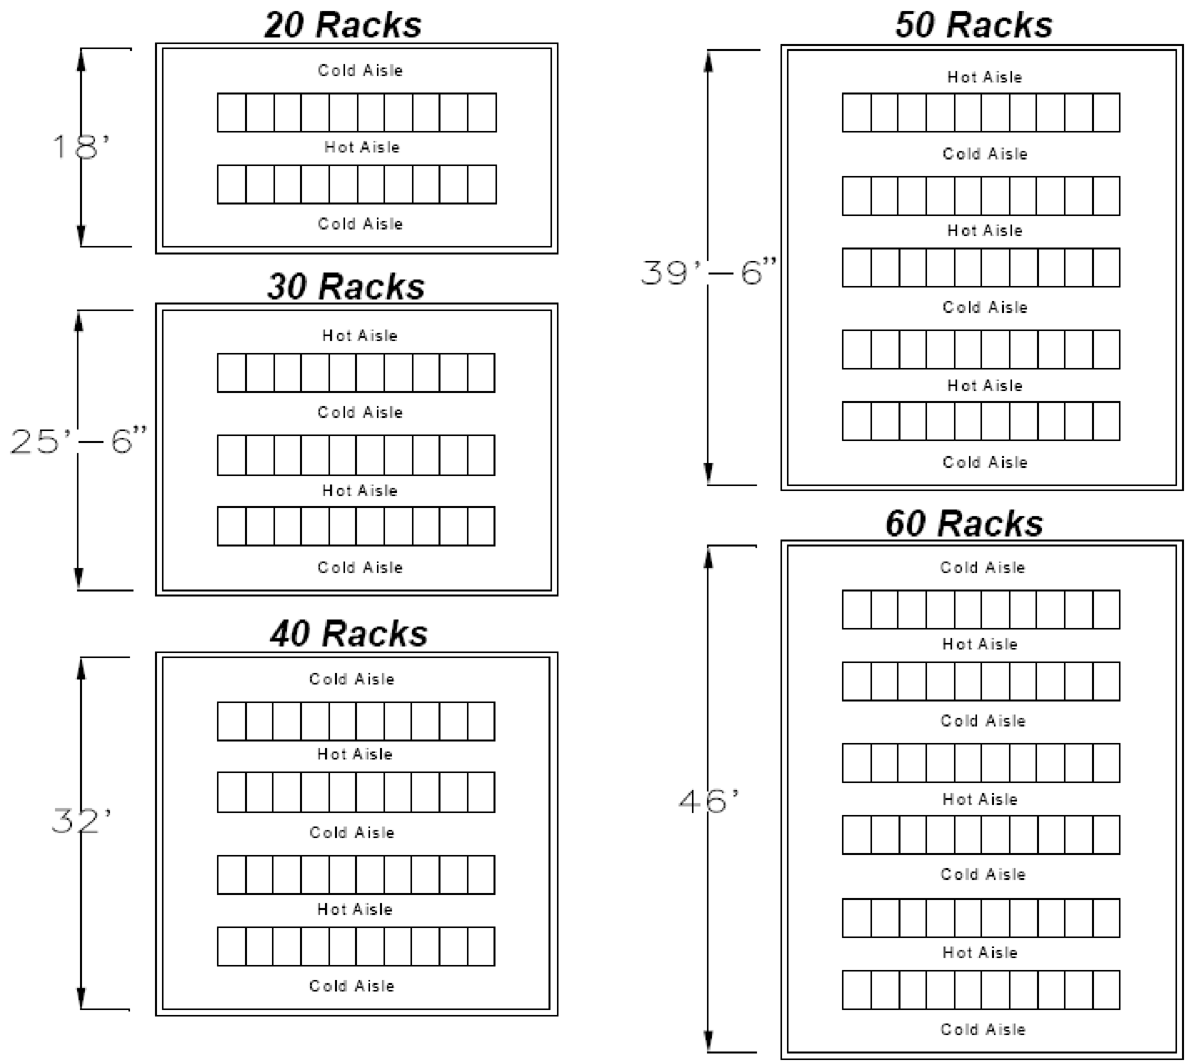

CRAC & Its selection

http://gefrc.com/thermal/Liebert_Challenger.asp

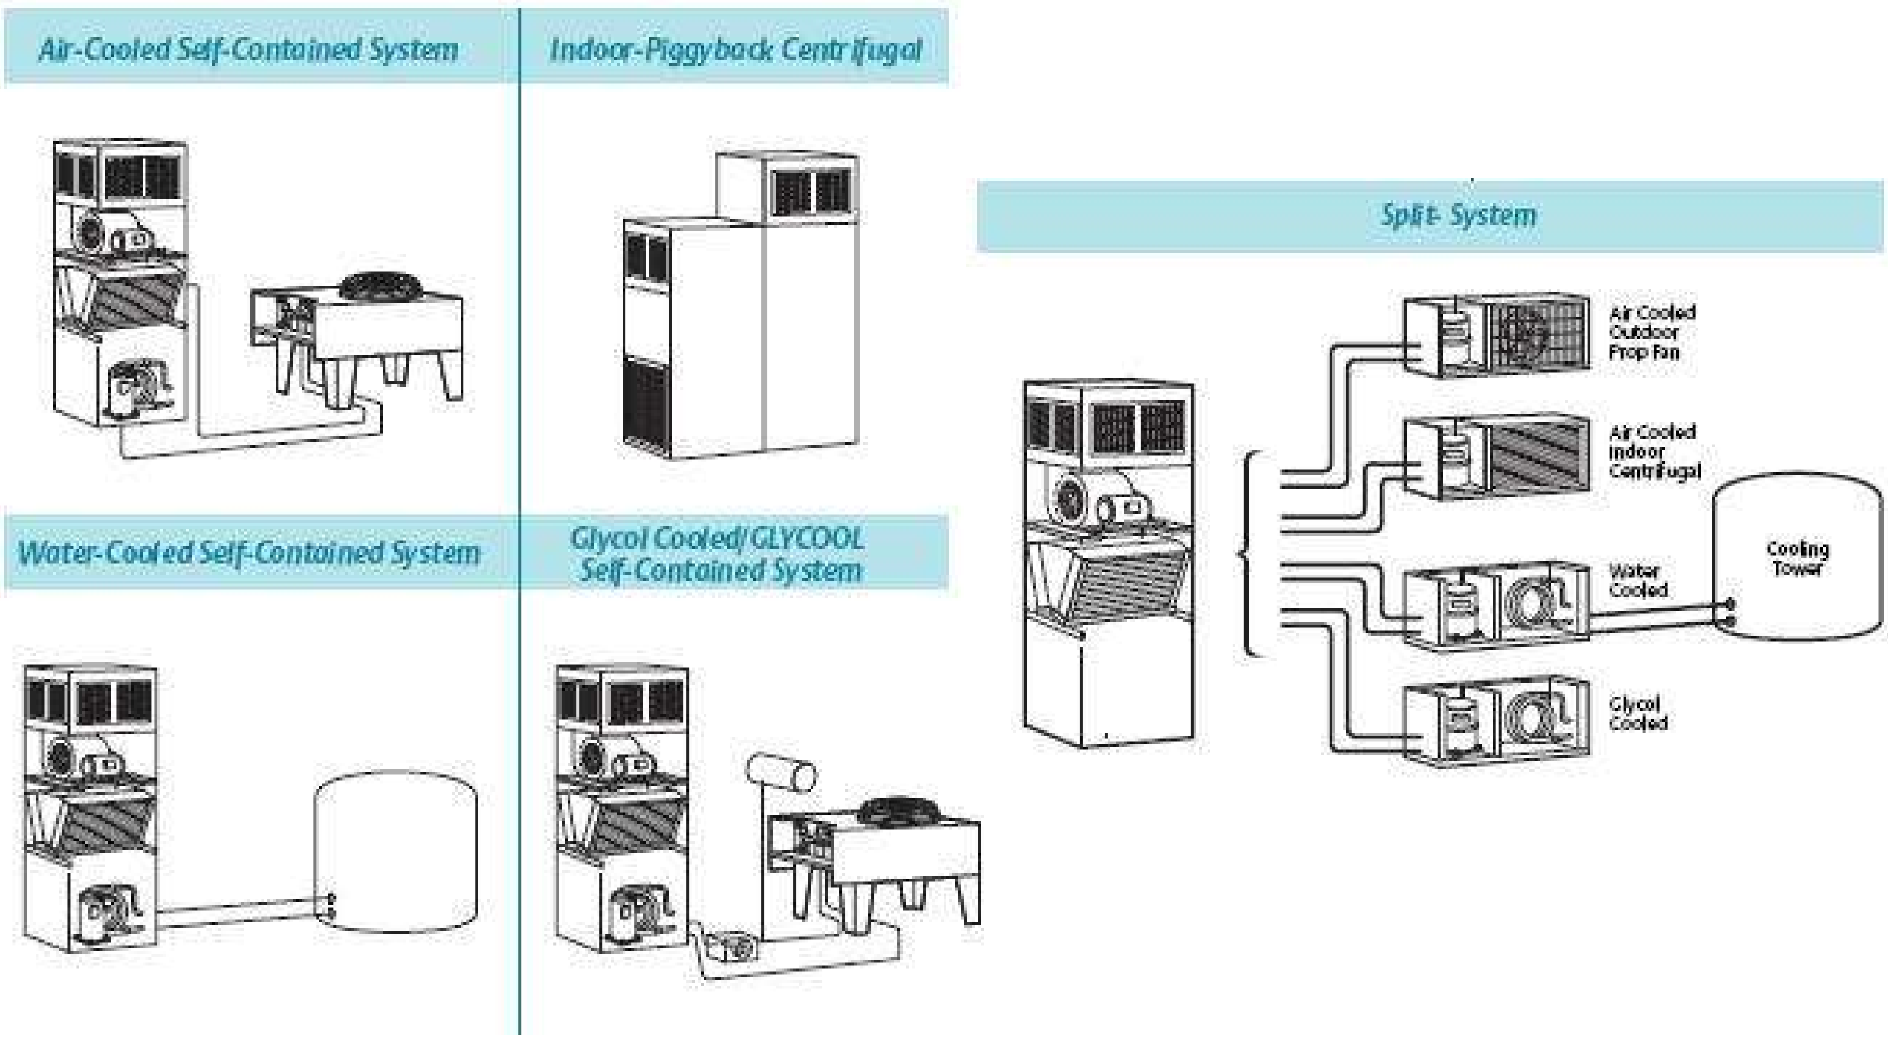

http://gefrc.com/thermal/Liebert_Challenger.asp

Up Flow Model for Non-Raised Floor

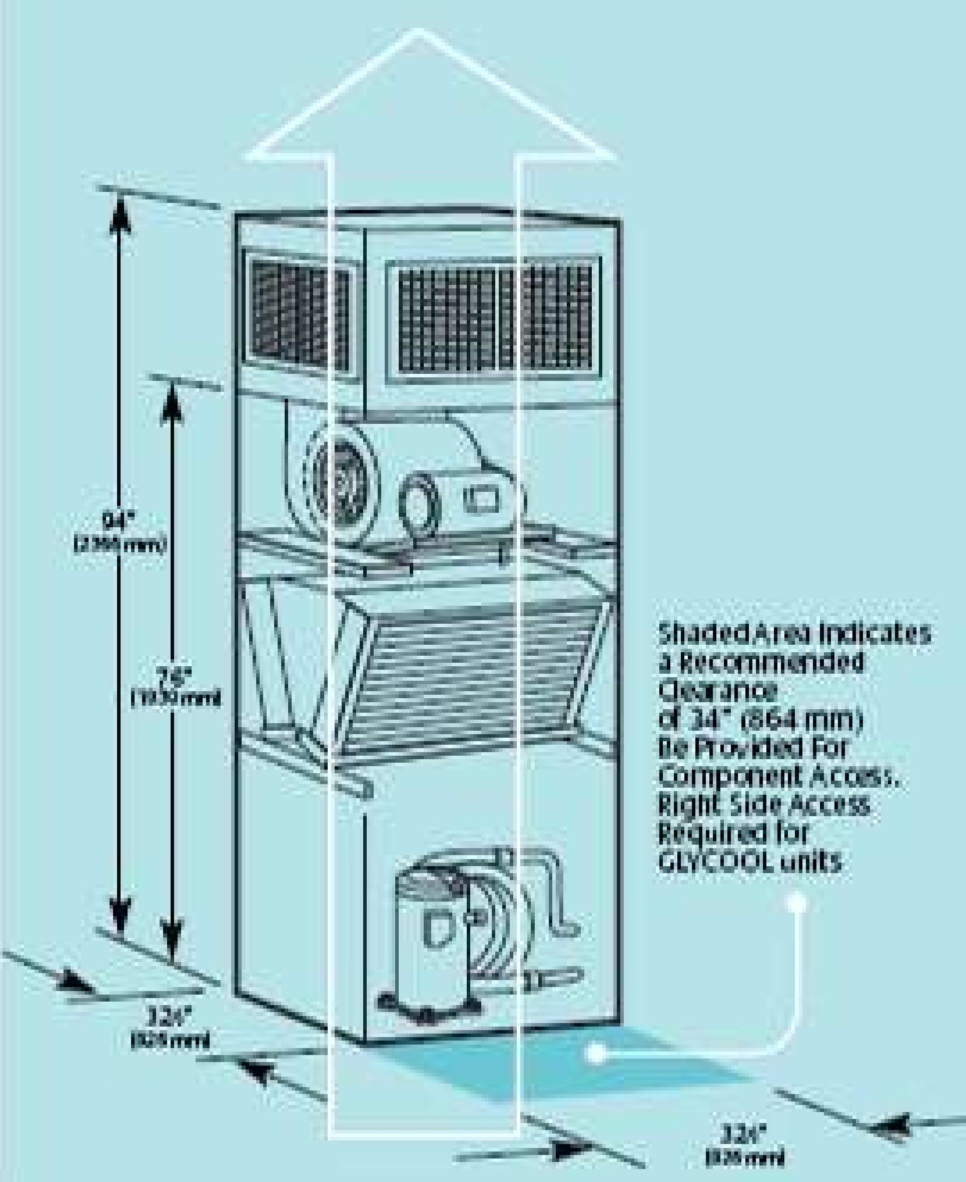

Down Flow Model for Raised Floor

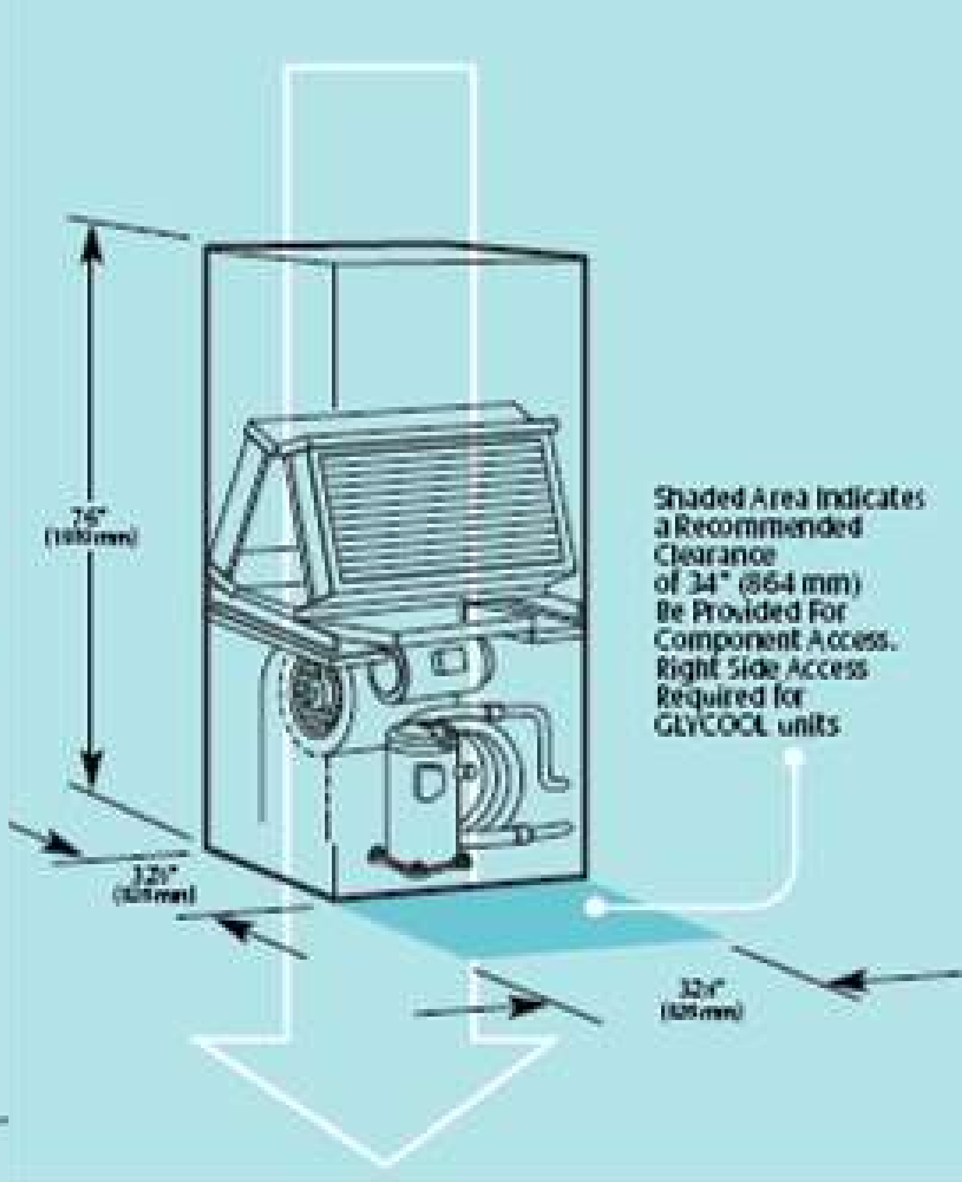

- For **smaller rooms** without raised floor or ductwork, upflow or downflow CRAC units are often **located in a corner or along a wall**. In these cases, it can be difficult to align cool air delivery with cold aisles and hot air return with hot aisles. Performance will be compromised in these situations. However, it is possible to improve the performance of these systems as follows:
  - For **upflow** units**, locate the unit near the end of a hot aisle and add ducts to bring cool air to points over cold aisles as far away from the CRAC unit as possible**.
  - For **downflow** units, **locate the unit at the end of a cold aisle oriented to blow air down the cold aisle**, and add either a dropped ceiling plenum return, or hanging ductwork returns with return vents located over the hot aisles.

The Different Technologies for Cooling Data Centers

APC white paper #59

Figure 1

Simplified breakdown of the 13 fundamental heat removal methods

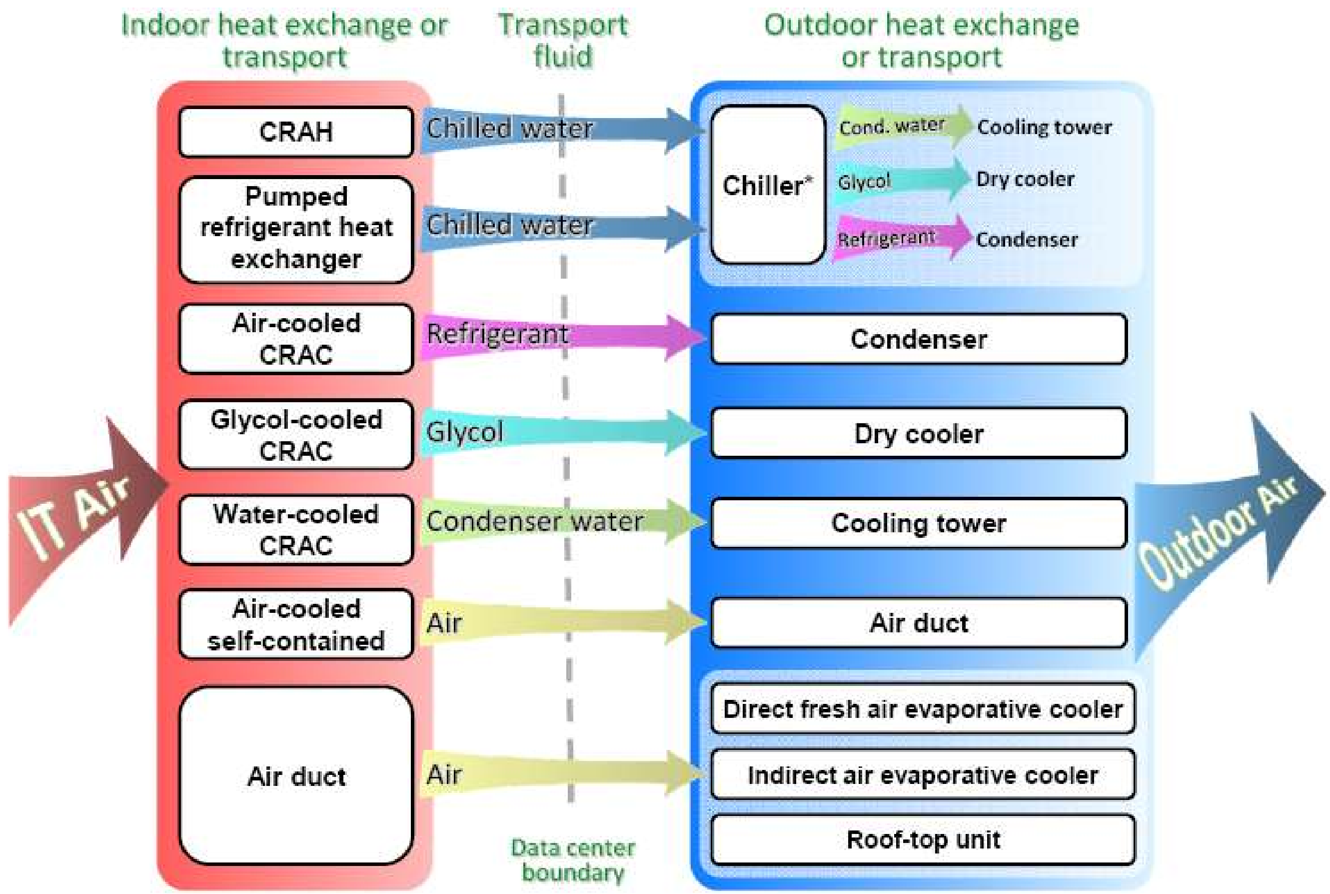

* Note that in some cases the chiller is physically located indoors.

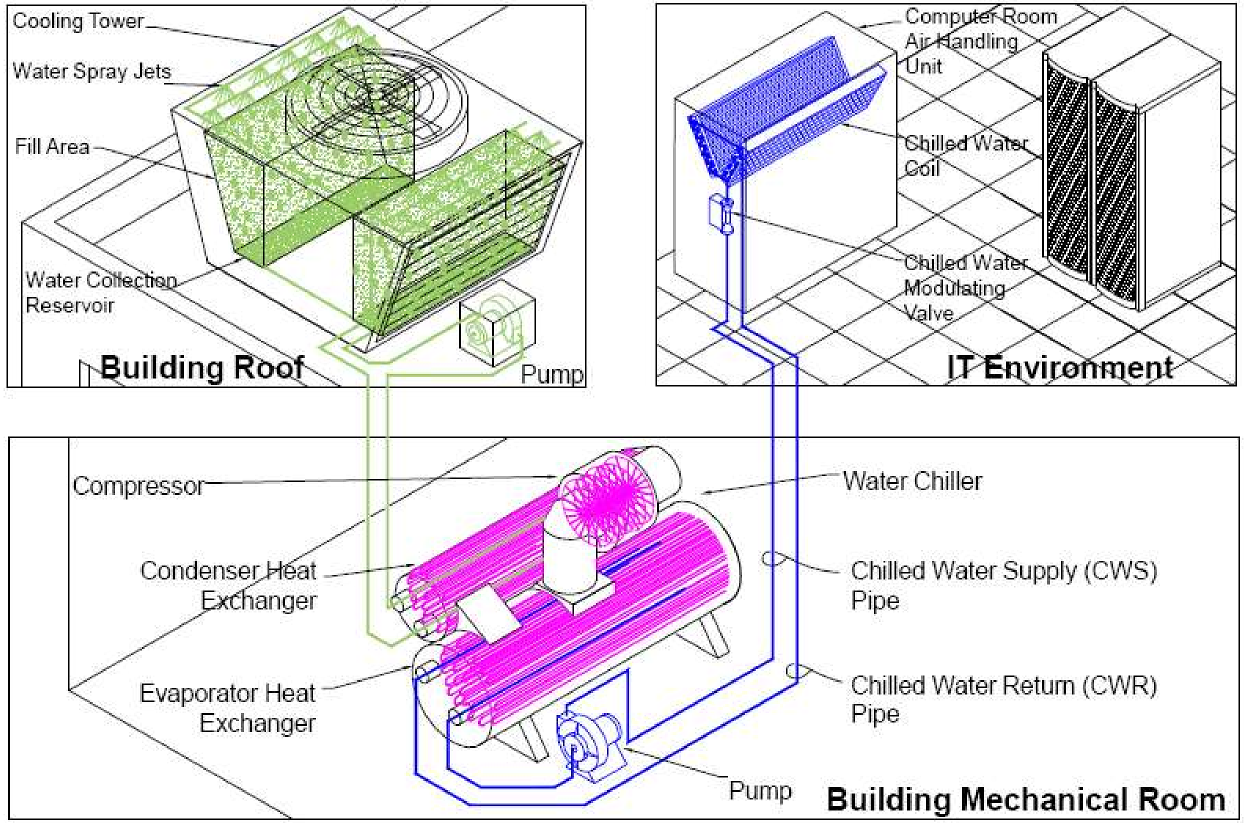

Figure 2

Water-cooled chilled water system

Water-cooled Chiller **Cooling Tower**

Figure 3

Example of a water-cooled chiller (left) and cooling tower (right)

Advantages

- Chilled water CRAH units generally cost less, contain fewer parts, and have greater heat removal capacity than CRAC units with the same footprint.
- Chilled water system efficiency improves greatly with increased data center capacity
- Chilled water piping loops are easily run very long distances and can service many IT environments (or the whole building) from one chiller plant.
- Chilled water systems can be engineered to be extremely reliable.
- Can be combined with economizer modes of operation to increase efficiency. Designing the system to operate at higher water temperatures 12-15°C [54-59°F]) will increase the hours on economizer operation.

Disadvantages

- Chilled water systems generally have the highest capital costs for installations below 100 kW of electrical IT loads.
- Introduces an additional source of liquid into the IT environment.

Usually used

- In data centers 200 kW and larger with moderate-to-high availability requirements or as a high availability dedicated solution. Water-cooled chilled water systems are often used to cool entire buildings where the data center may be only a small part of that building.

Pumped refrigerant for chilled water systems

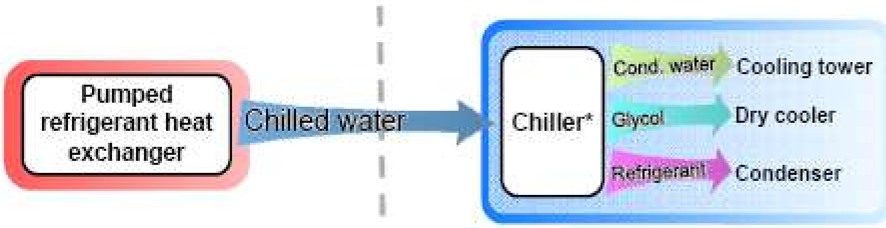

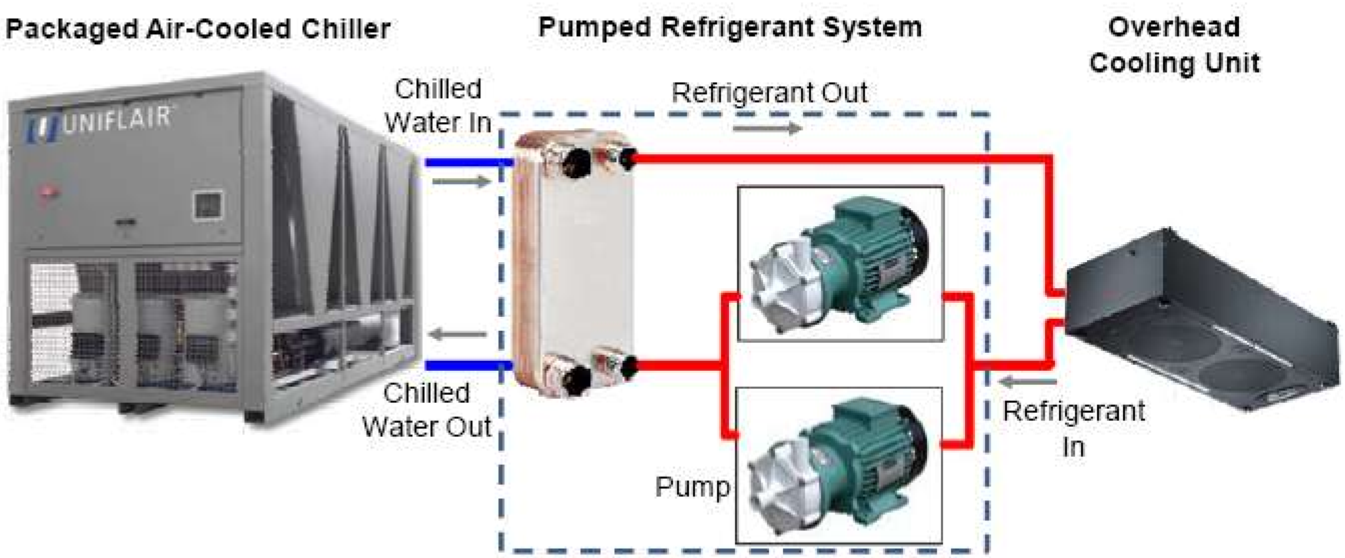

Advantages

- Keeps water away from IT equipment in chilled water applications
- Oil-less refrigerants and non-conductive fluids eliminate risk of mess or damage to servers in the event of a leak.
- Efficiency of cooling system due to close proximity to servers or direct to chip level.

Disadvantages

- Higher first cost as a result of adding additional pumps and heat exchangers into the cooling system.

Usually Used

- These systems are usually used for cooling systems that are closely coupled to the IT equipment for applications like row and rack based high density cooling.
- Chip Level Cooling where coolant is piped directly to the server

Air-cooled system (2-piece)

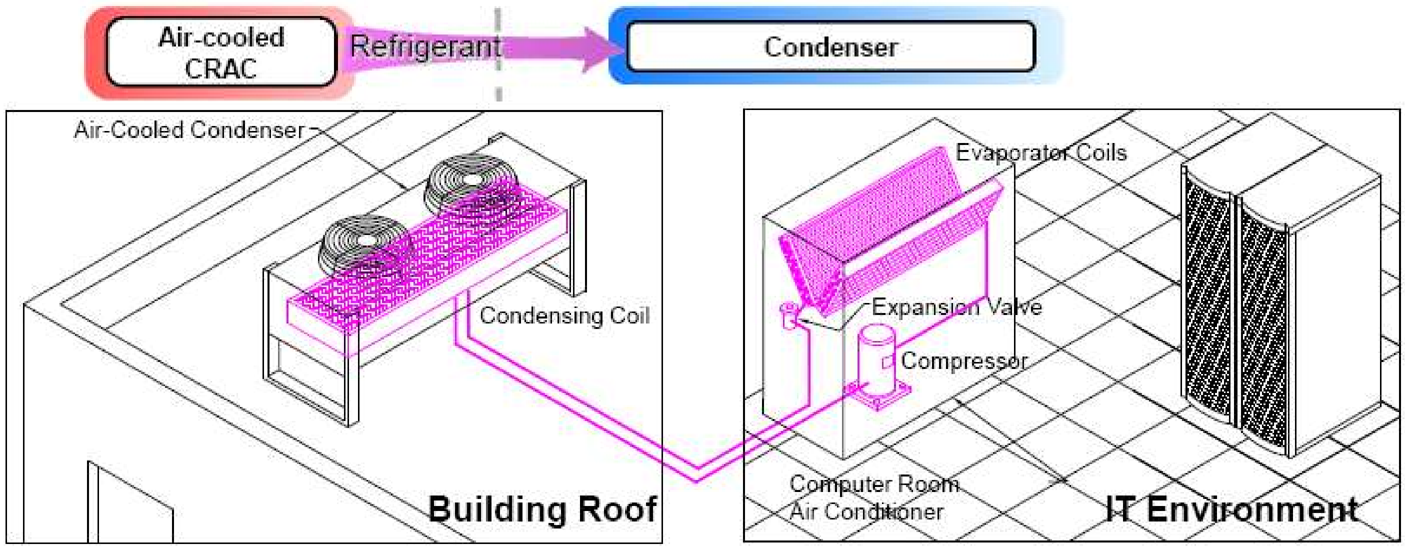

Figure 6

Air-cooled DX system (2-piece)

Advantages

- Lowest overall cost
- Easiest to maintain

Disadvantages

- Refrigerant piping must be installed in the field. Only properly engineered piping systems that carefully consider the distance and change in height between the IT and outdoor environments will deliver reliable performance.
- Refrigerant piping cannot be run long distances reliably and economically.
- Multiple computer room air conditioners cannot be attached to a single air-cooled condenser.

Usually used

- In wiring closets, computer rooms and 7-200kW data centers with moderate availability requirements.

Glycol-cooled system

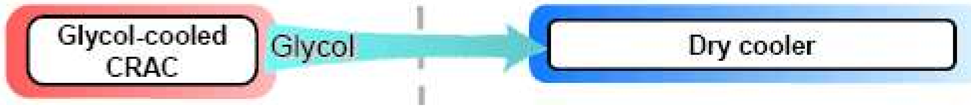

Figure 8

Glycol-cooled system

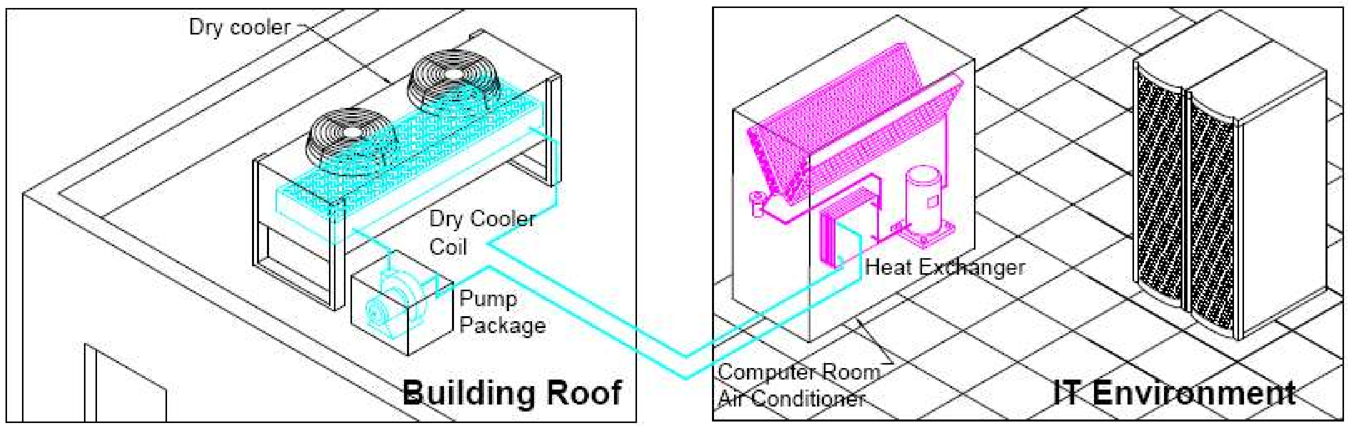

Advantages

- The entire refrigeration cycle is contained inside the CRAC unit as a factory-sealed and tested system for highest reliability with the same floor space requirement as a two piece air-cooled system.
- Glycol pipes can run much longer distances than refrigerant lines (air-cooled split system) and can service several CRAC units from one dry cooler and pump package.
- In cold locations, the glycol within the dry cooler can be cooled so much (below 10°C [50°F]) that it can bypass the heat exchanger in the CRAC unit and flow directly to a specially installed *economizer coil*. Under these conditions, the refrigeration cycle is turned off and the air that flows through the economizer coil, now filled with cold flowing glycol, cools the IT environment. This economizer mode, also known as "*free cooling*", provides excellent operating cost reductions when used.

Disadvantages

- Additional required components (pump package, valves) raise capital and installation costs when compared with air-cooled DX systems.
- Maintenance of glycol volume and quality within the system is required.
- Introduces an additional source of liquid into the IT environment.

Usually used

- In computer rooms and 30-1,000 kW data centers with moderate availability require-ments.

Water-cooled system

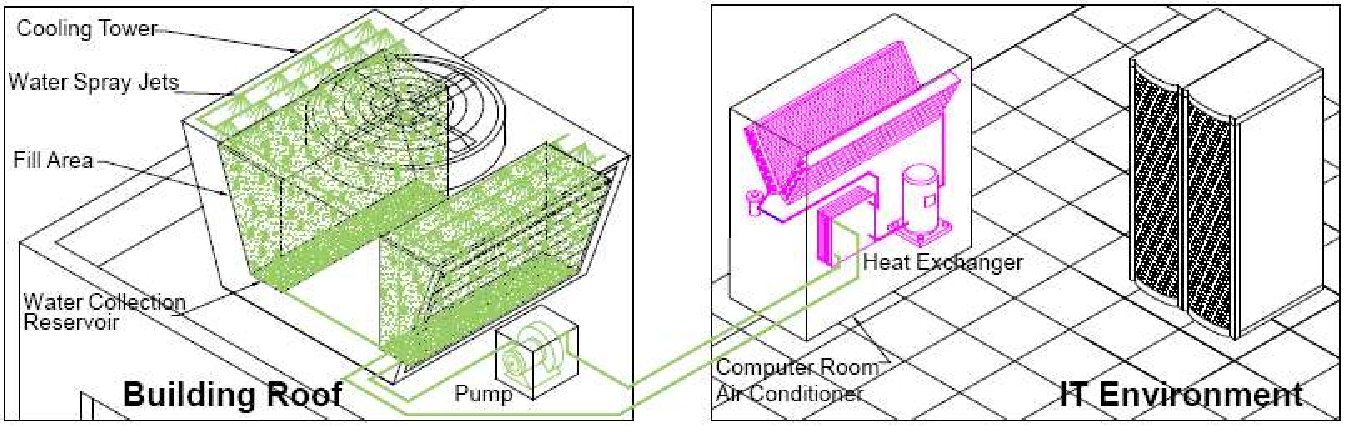

Figure 9

Water-cooled system

- A water (also called *condenser water*) loop is used instead of glycol to collect and transport heat away from the IT environment
- Heat is rejected to the outside atmosphere via a cooling tower instead of a dry cooler as seen in **Figure 9**.

Advantages

- All refrigeration cycle components are contained inside the computer room air conditioning unit as a factory-sealed and tested system for highest reliability.
- Condenser water piping loops are easily run long distances and almost always service many computer room air conditioning units and other devices from one cooling tower.
- In leased IT environments, usage of the building's condenser water is generally less expensive than chilled water (chilled water is explained in the next section).

Disadvantages

- High initial cost for cooling tower, pump, and piping systems.
- Very high maintenance costs due to frequent cleaning and water treatment requirements.
- Introduces an additional source of liquid into the IT environment.
- A non-dedicated cooling tower (one used to cool the entire building) may be less reliable then a cooling tower dedicated to the computer room air conditioner.

Usually used

- In conjunction with other building systems in data centers 30kW and larger with moderate-to-high availability requirements.

Air-cooled self-contained system (1-piece)

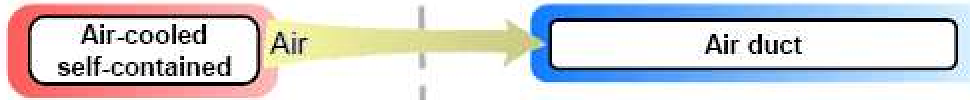

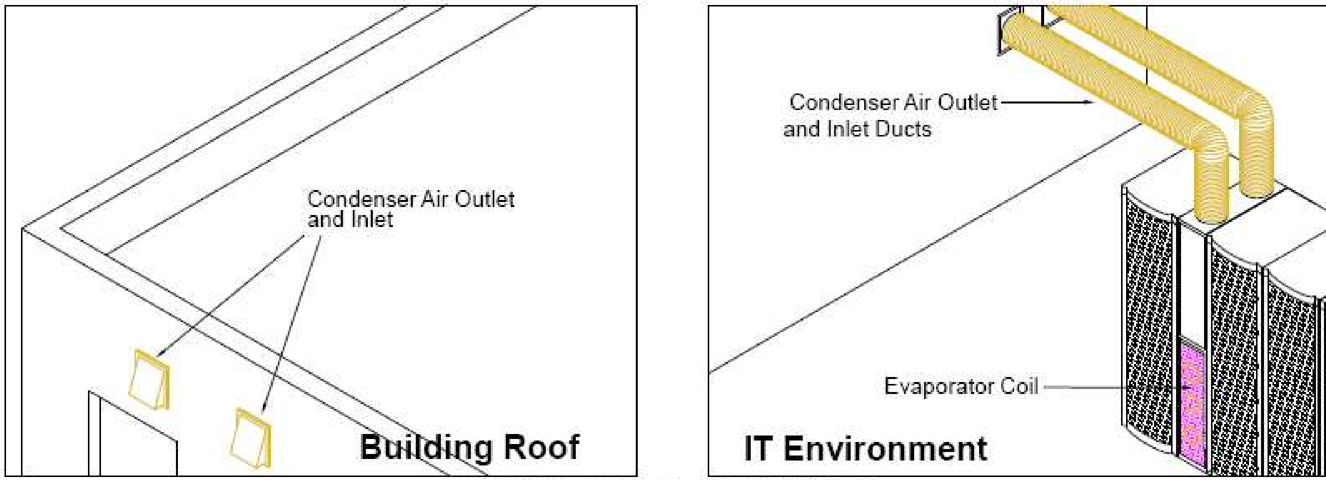

Air cooled self contained

Figure 10

Indoor air-cooled self-contained system

Portable Self Contained Cooling Unit

Figure 11

Examples of indoor air-cooled self-contained system

Direct fresh air evaporative cooling system

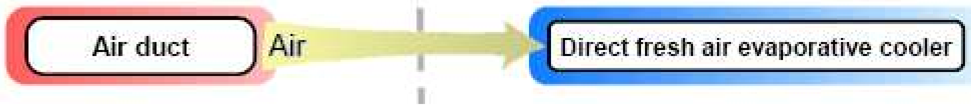

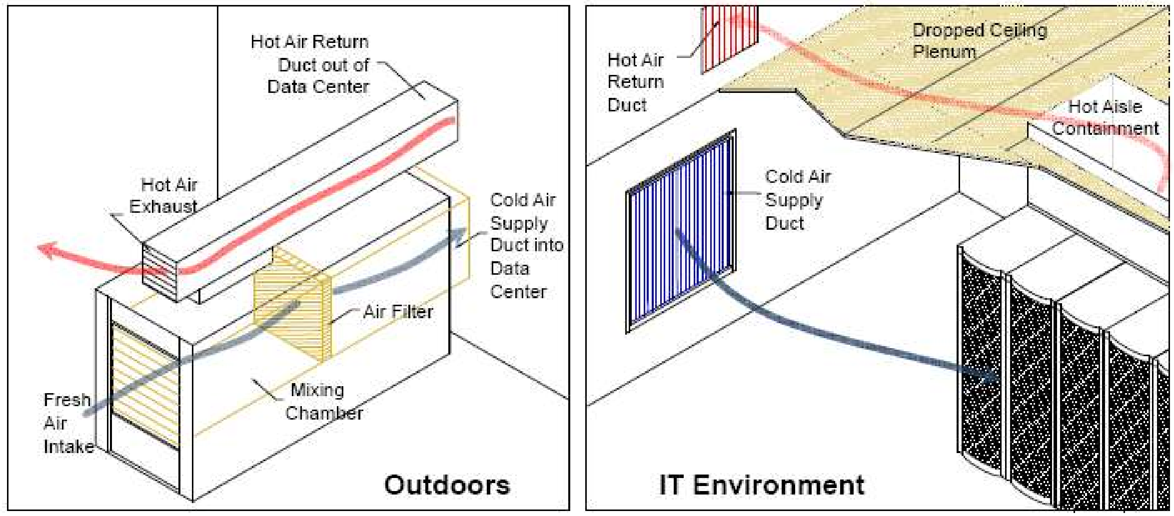

Figure 12

Example of an indirect air evaporative cooling system

Figure 13

Example of a direct fresh air evaporative cooling system

Advantages

- All cooling equipment is placed outside the data center, allowing for white space to be fully utilized for IT equipment.
- Significant cooling energy savings in dry climates (e.g. 75%) compared to systems with no economizer mode.

Disadvantages

- May be difficult to retrofit into an existing data center.
- Subject to frequent filter changes in locations with poor air quality.
- Evaporative cooling contributes to humidity in the data center.

Usually used

- In 1,000kW data centers and larger with high power density.

Indirect air evaporative cooling system

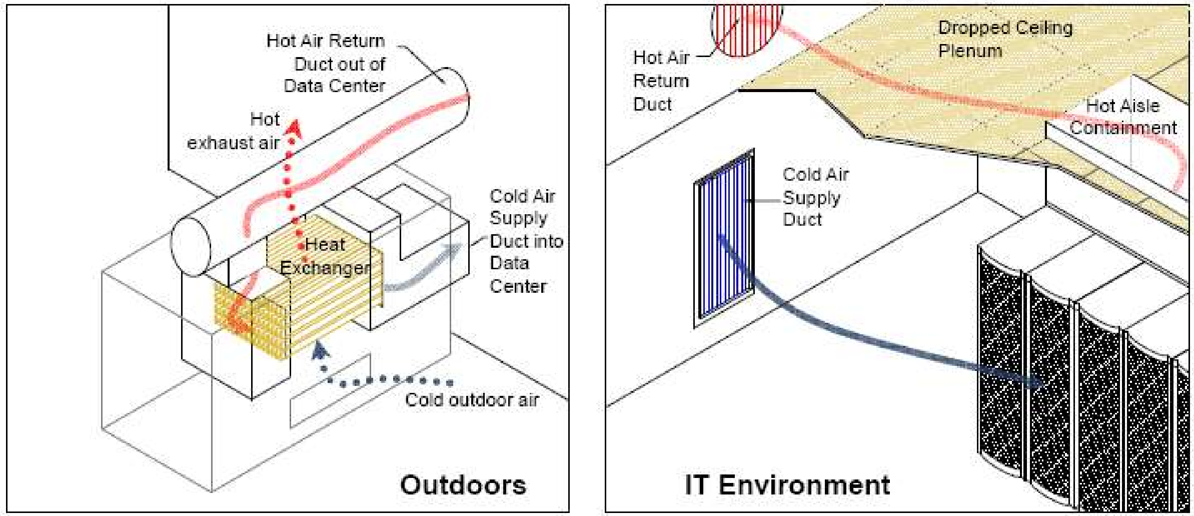

Figure 14

Indirect air economizer system

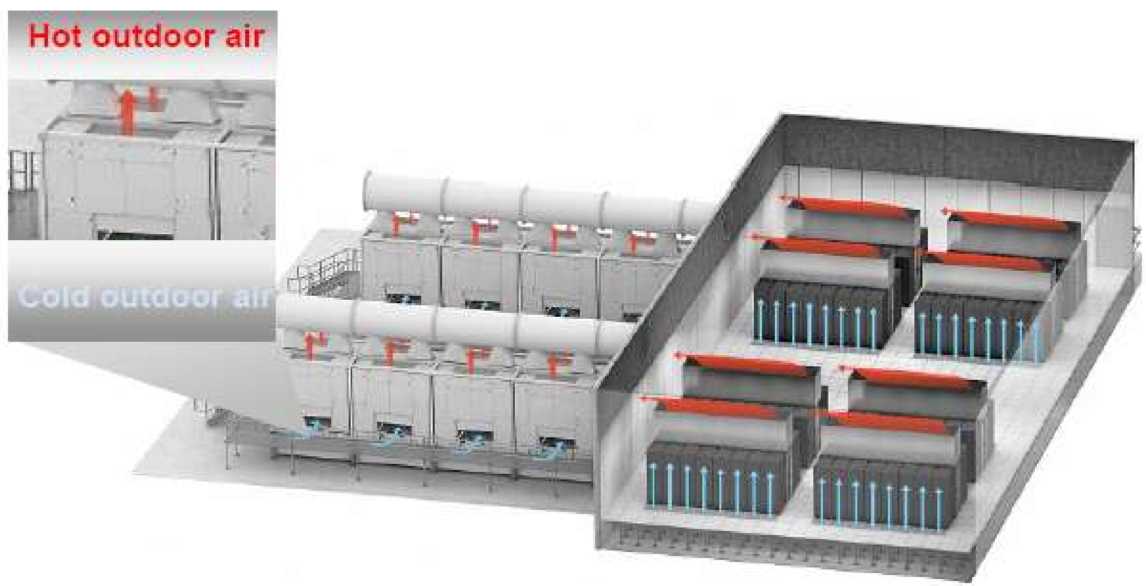

Figure 15

Example of an indirect air evaporative cooling system

Advantages

- All cooling equipment is placed outside the data center, allowing for white space to be fully utilized for IT equipment.
- Significant cooling energy savings in most climates (e.g. 75%) compared to systems with no economizer mode.

Disadvantages

- May be difficult to retrofit into an existing data center.

Usually used

- In 1,000kW data centers and larger with high power density.

Self-contained roof-top system

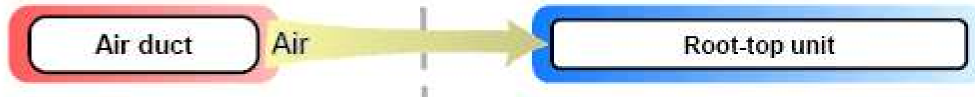

Figure 16

Self-contained roof-top system

Advantages

- All cooling equipment is placed outside the data center, allowing for white space to be fully utilized for IT equipment.
- Significant cooling energy savings in mild climates compared to systems with no economizer mode.

Disadvantages

- May be difficult to retrofit into an existing data center.

Usually used

- In data centers that are part of a mixed-use facility.

Advantages

- Indoor self-contained systems have the lowest installation cost. There is nothing to install on the roof or outside the building except for the condenser air outlet.
- All refrigeration cycle components are contained inside one unit as a factory-sealed and tested system for highest reliability.

Disadvantages

- Less heat removal capacity per unit compared to other configurations.
- Air routed into and out of the IT environment for the condensing coil usually requires ductwork and/or dropped ceiling.
- Some systems can rely on the building HVAC system to reject heat. Issues can arise when the building HVAC system shuts down in the evening or over the weekend.

Usually used

- In wiring closets, laboratory environments and computer rooms with moderate availability requirements.
- Sometimes used to fix hot spots in data centers.

Two fundamental physical arrangements of precision cooling equipment

- **Ceiling mounted systems**

Figure 6 – Typical ceiling mounted computer room air conditioner

Figure 7 – Typical floor mounted portable computer room air conditioner

Figure 8 – Typical floor mounted computer room air conditioner

- **Floor mounted system**

The 10 combinations of heat removal methods and equipment arrangements

Table 1 – The 10 basic cooling system configurations

	Ceiling Mounted	Floor Mounted
Air Cooled System (2Piece)	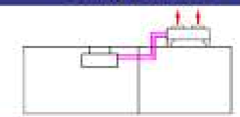 Requires roof access and a 10' (3m) floor to structural ceiling height. Roof should be within 2 stories of IT environment. Air cooled condenser and refrigerant piping required.	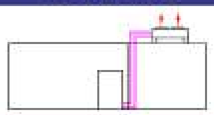 Requires roof access. Roof should be within two stories of IT environment. Requires air cooled condenser and refrigerant piping. Portable systems usually don't use outdoor components.
Air Cooled Self Contained System (1Piece)	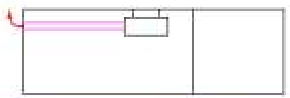 IT environment must have dropped ceiling or ducts should be installed for condenser air. Ensure 10' (3m) floor to structural ceiling height.	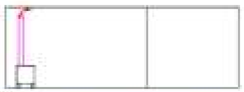 IT environment must have dropped ceiling for condenser air tubes. Large floor mounted systems require outdoor heat rejection components.
Glycol Cooled Systems	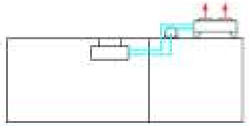 Building must have roof access and a 10' (3m) floor to structural ceiling height. Fluid cooler, pump package and glycol piping required.	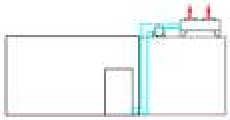 Requires roof access. Fluid cooler, pump package and glycol piping required. Portable systems usually don't use outdoor components.
Water Cooled Systems	 Building must have 10' (3m) floor to structural ceiling height. Hookup to building condenser water required.	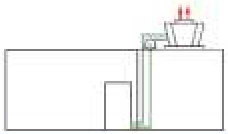 Building must have condenser water system with adequate capacity. Hookup required. Portable systems don't use condenser water.
Chilled Water Systems	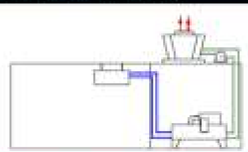 Building has 10' (3m) floor to structural ceiling height and reliable chilled water system. Chilled water hookup required.	 Building must have reliable chilled water system with adequate capacity. Chilled water hookup required. Portable systems usually don't use chilled water.

DATA Center Environment

Arguments for increased air temperatures:

Journal of Electronic Packaging, MARCH 2011, Vol. 133 / 011004-1

– estimates predict that 4–5% of data center energy costs could be saved for every 1°C increase in server inlet temperature

– higher temperature settings could allow more free-cooling

– servers have previously been deployed in high temperature environments with no failures.

Arguments against increased air temperatures

Journal of Electronic Packaging MARCH 2011, Vol. 133 / 011004-1

- raising inlet temperature could cause additional energy consumption in other components.
- costs of thermal-related equipment failures may exceed operational savings.
- higher temperatures can impact server performance and service life.
- higher room ambient temperatures may require more frequent server shutdown where cooling system failure occurs.

The Criterion for Acceptable Cooling in a Data Center

- Manufacturers of computer servers design their equipment with a certain allowable maximum inlet temperature. This value is around 24°C (75°F).
- The air-conditioners in a data center usually supply cold air at 13°C (55°F). If the 13°C cold air enters the servers, there is no difficulty in satisfying the manufacturer’s criterion. But the cold air does not always enter at all the inlet locations on the server rack.
- Often, the hot air exhausted by the rack finds its way to the inlet of the same rack or some other rack. This is how the cooling in a data center is compromised.

Conventional data center air conditioning flow diagram

ASHRAE Transactions 2010, Vol. 116, Part 1.2010, 98-108

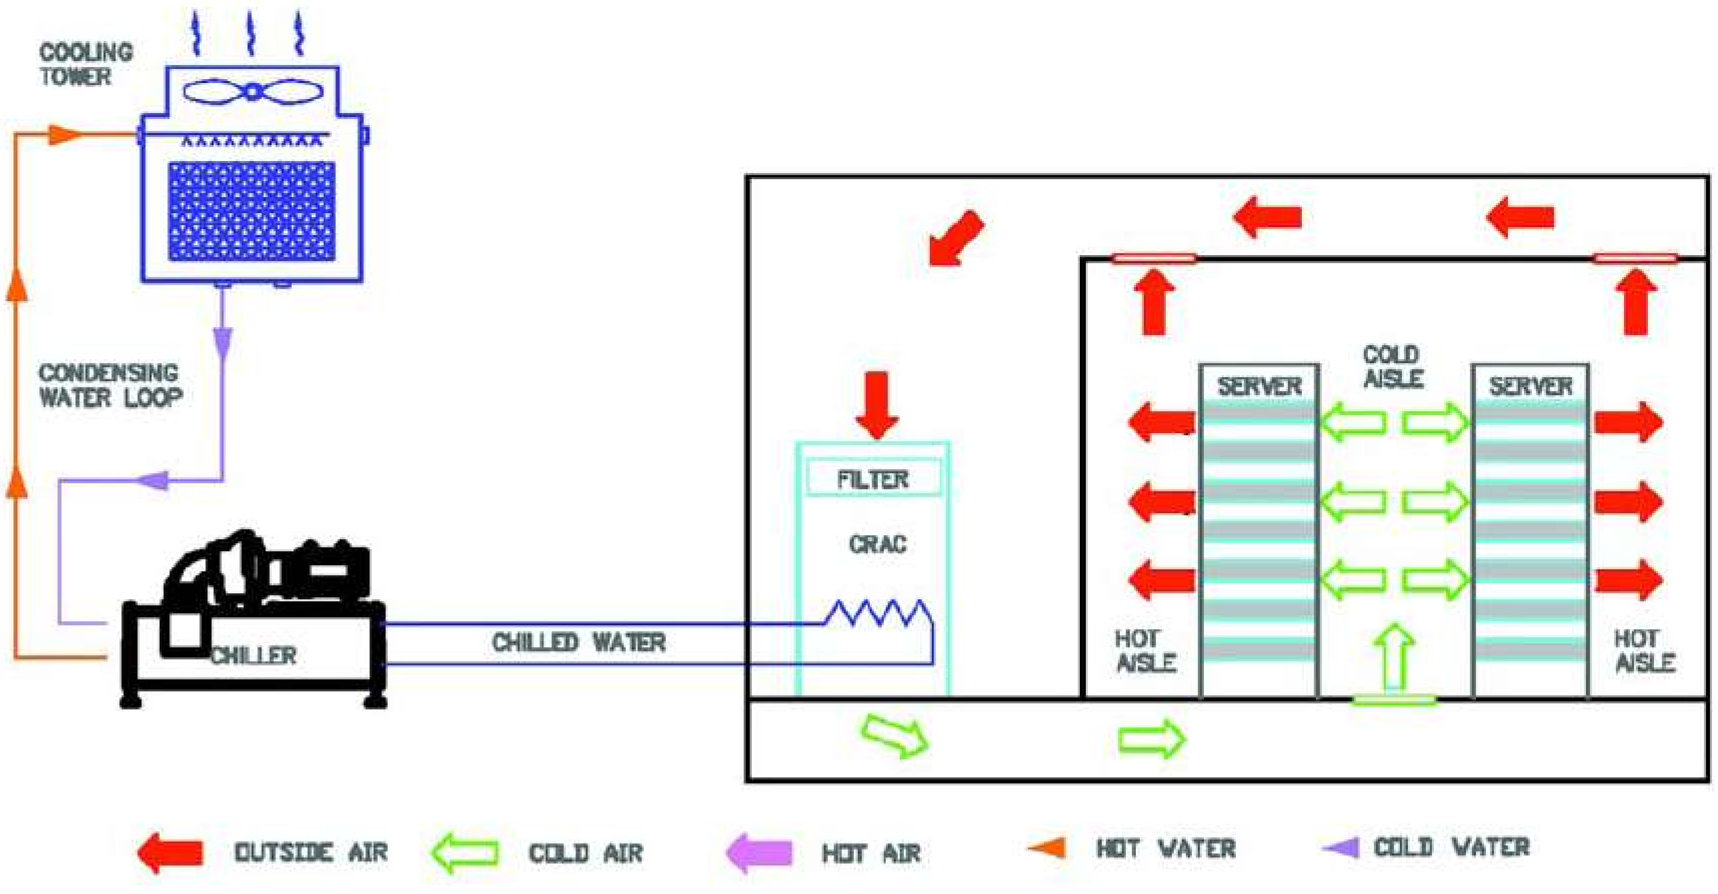

Which cooling solutions are appropriate for use in different size IT environments?

- **Wiring closets** (1-3 rack enclosures or equivalent using 1-18 kW of electricity)
- **Computer rooms** (1-5 rack enclosures or equivalent using 3-30 kW of electricity)
- **Small data centers** (5-20 rack enclosures or equivalent using 7-100 kW of electricity)
- **Medium data centers** (20-100 rack enclosures or equivalent using 28-500 kW of electricity)
- **Large data centers** (> 100 rack enclosures or equivalent using >200 kW electricity)

Wiring closets

- If closet temperatures are high, first try to increase ventilation to IT and communications equipment. If temperatures remain high and a precision cooling solution is required, ensure equipment ventilation and clearance requirements can be met for the proposed solution.

Table 2 – Basic cooling system configurations for wiring closets

	Ceiling Mounted	Floor Mounted
Air Cooled DX System (2Piece)	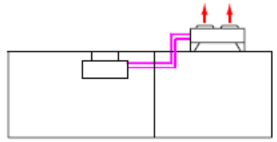 **Use if**: Building has roof access and enough room for solution. Roof is within 2 stories of wiring closet. Use for 3-17kW of equipment.	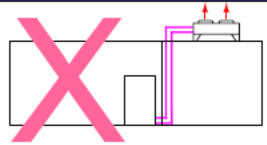 Take up significant floor space in a wiring closet. Typically sized system much larger than wiring closet requirements.
Air Cooled Self Contained System (1Piece)	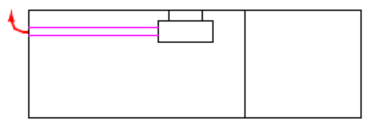 **Use if:** Ducts can be installed to supply and return condenser air. Use for 3-17kW of equipment.	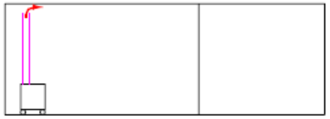 **Use if:** Ease of install and portability are advantages. Hot air exhaust ducts can be run outside the wiring closet. Use for 3-6 kW of equipment.
Glycol Cooled Systems	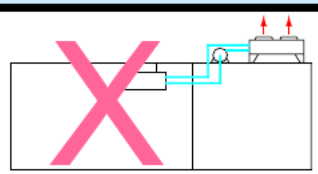 Not commonly used in this power range	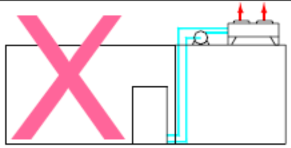 Take up significant floor space in a wiring closet. Typically sized system much larger than wiring closet requirements.
Water Cooled Systems	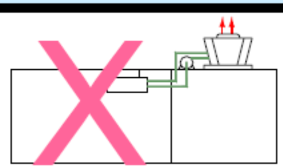 Not commonly used in this power range	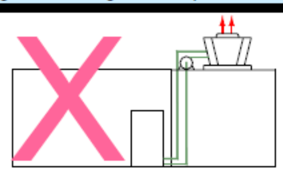 Take up significant floor space in a wiring closet. Typically sized system much larger than wiring closet requirements.
Chilled Water Systems	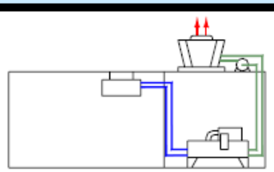 **Use if:** The only cooling source is chilled water, no possible location for outdoor condensers. Some hi-rise buildings.	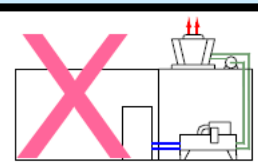 Take up significant floor space in a wiring closet. Typically sized system much larger than wiring closet requirements.

Suitable operating temperature for wiring closets (APC white paper #68)

- For active IT equipment typically found in a wiring closet, this temperature is usually 104°F (40°C). This is the maximum temperature at which the vendor is able to guarantee performance and reliability for the stated warranty period. It is important to understand that although the maximum published operating temperature is acceptable per the manufacturer, operating at that temperature will not generally provide the same level of availability or longevity as operating at lower temperatures. Because of this, some IT equipment vendors also publish recommended operating temperatures for their equipment in addition to the maximum allowed. Typical recommended operating temperatures from IT equipment vendors are between 70°F (21°C) and 75°F (24°C).

Heat dissipates in a small confined space like an office or closet in five different ways

- Conduction: Heat can flow through the walls of the space
- Passive Ventilation: Heat can flow into cooler air via a vent or grille, without an air moving device
- Fan-assisted Ventilation: Heat can flow into cooler air via a vent or grille that has an air moving device
- Comfort Cooling: Heat can be removed by a building’s comfort cooling system
- Dedicated Cooling: Heat can be removed by a dedicated air conditioner

General guideline for cooling strategies based on room power and target

Figure 1

Cooling method guide based on power load and target room temperature

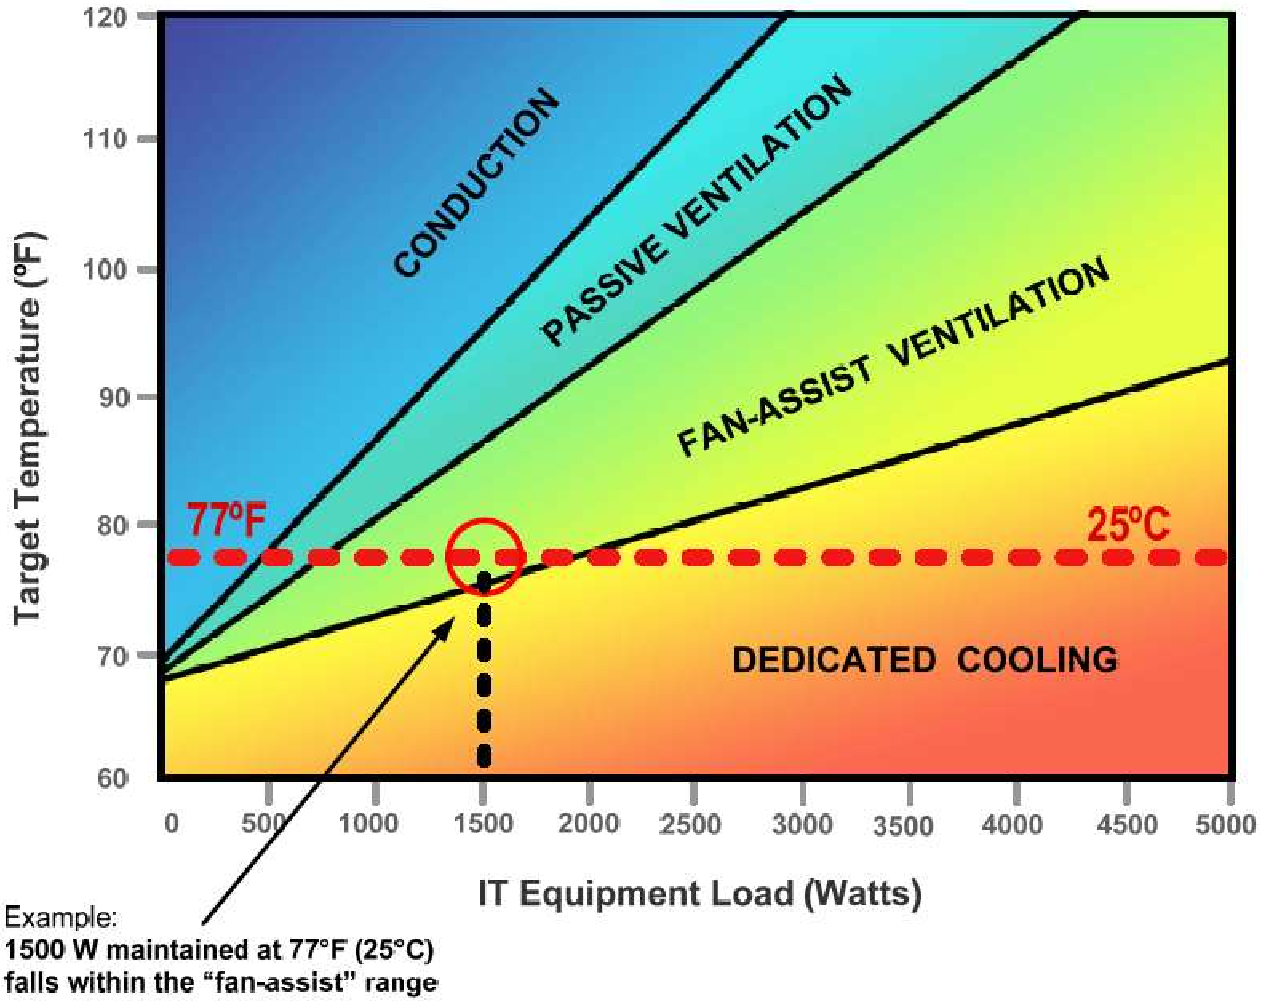

Figure 2

Cooling method selection to maintain recommended ASHRAE temperature range of 68 - 77°F (20 - 25°C)

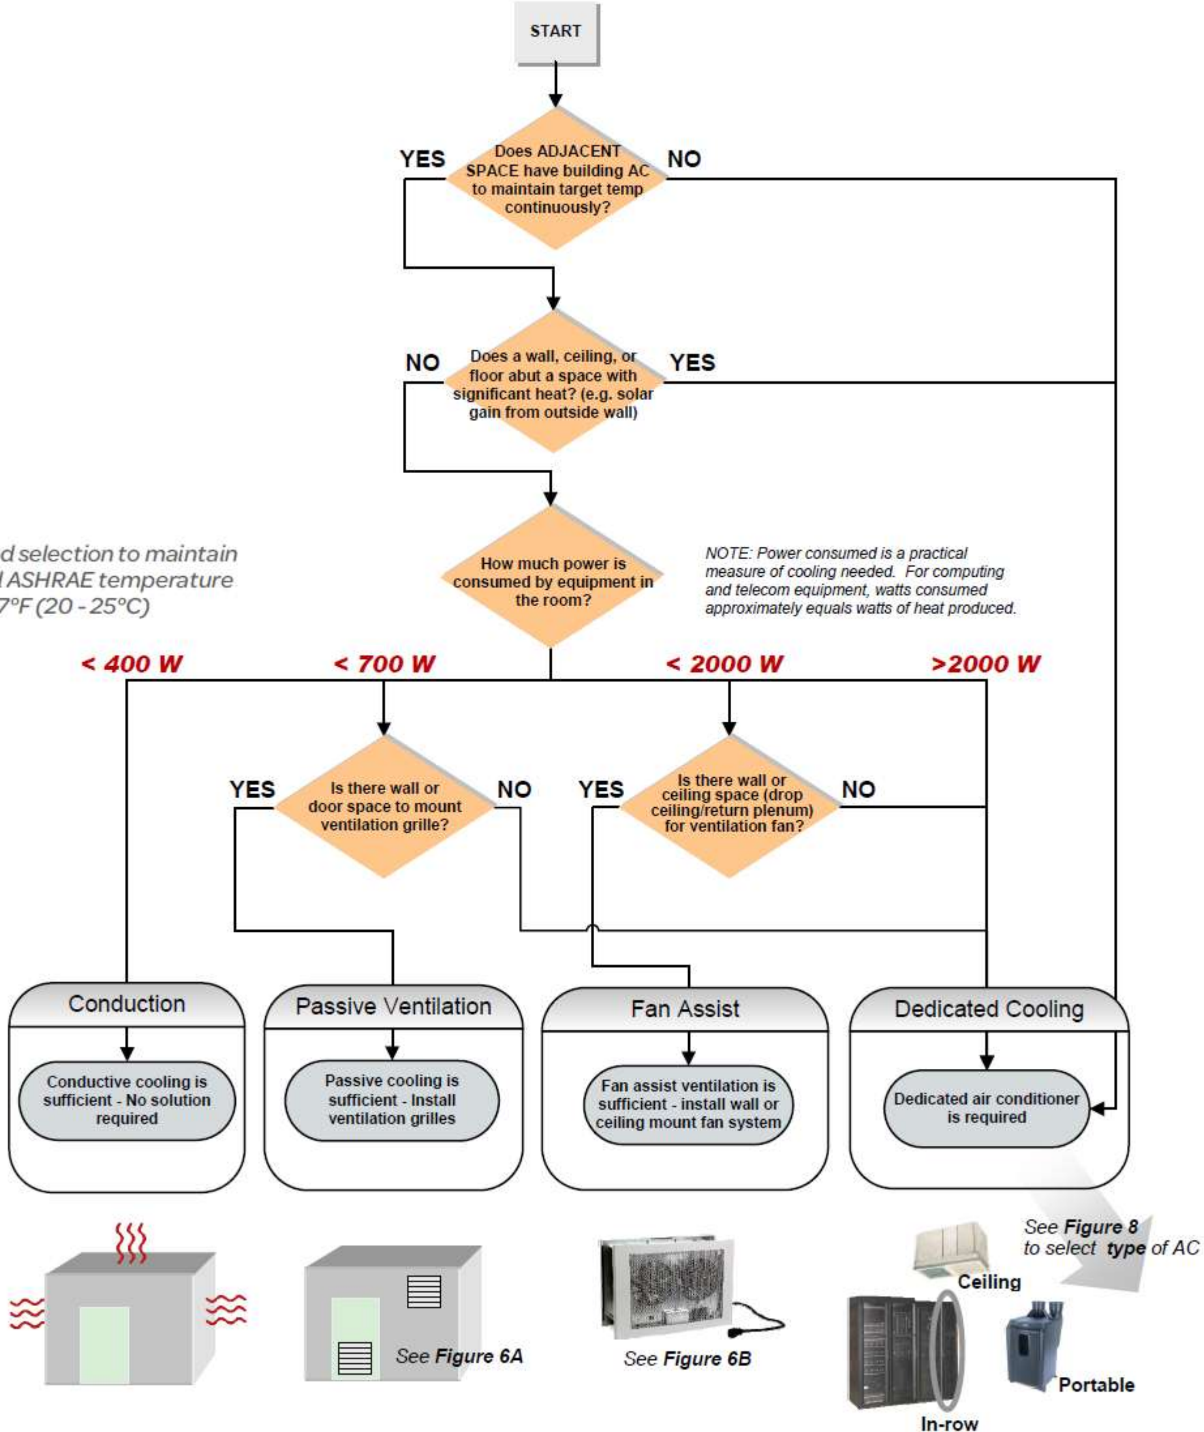

Table 2

Factors that can influence closet temperature vs. load relationship and expected impact

Factor	Expected impact on closet temperature
Room dimensions	Temperature increases as room dimensions decrease
Wall, ceiling, floor material	Temperature increases as construction material thermal resistance increases
Setback of building air conditioner on nights / weekends	Every degree increase in building air conditioner increases closet temperature by same amount
One wall subject to sun exposure / outdoor temperature on hot sunny day	Temperature increases as wall area exposed to outdoor temperature and sun increase

Figure 3

Closet temperatures versus IT equipment load: conduction performance

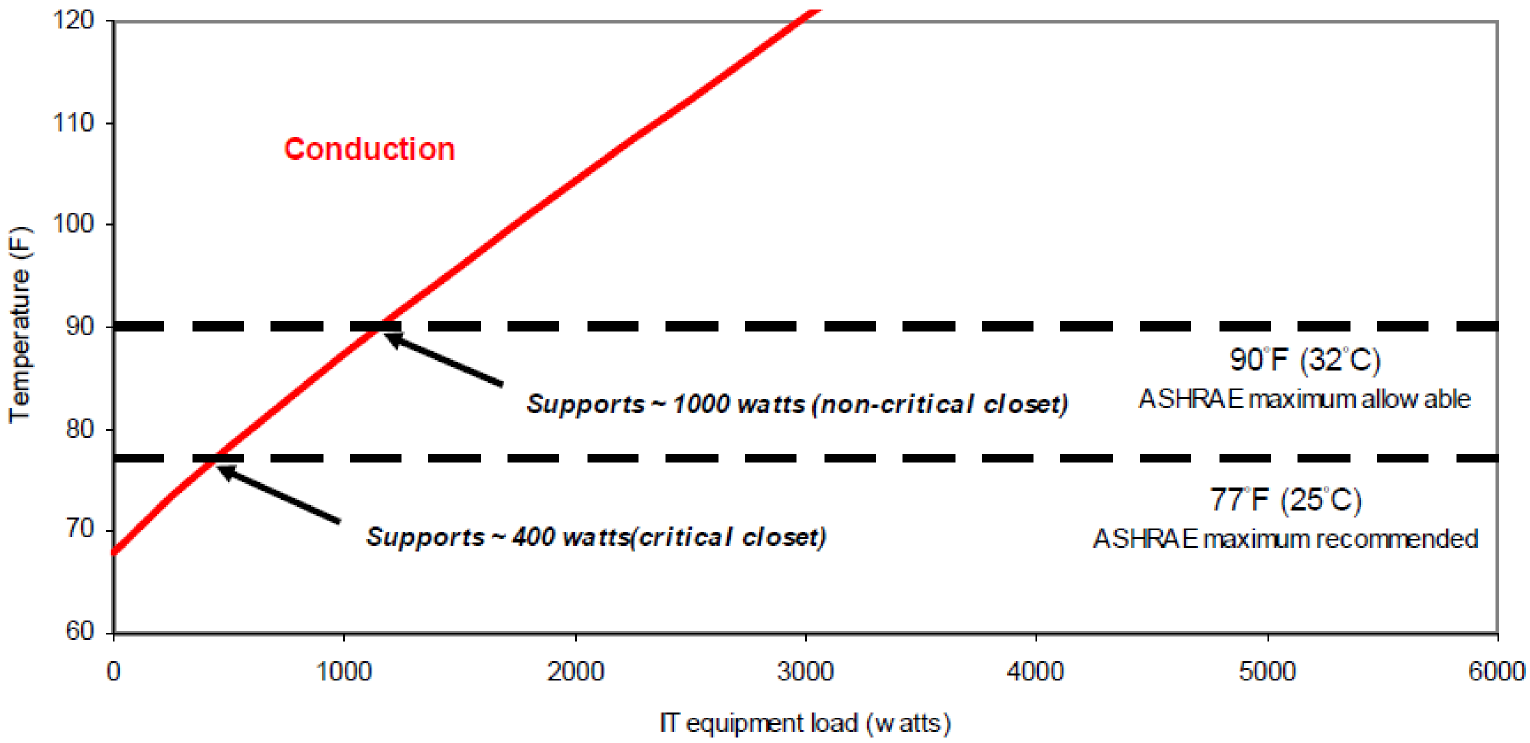

Figure 4

Effect of closet dimensions on conduction cooling performance

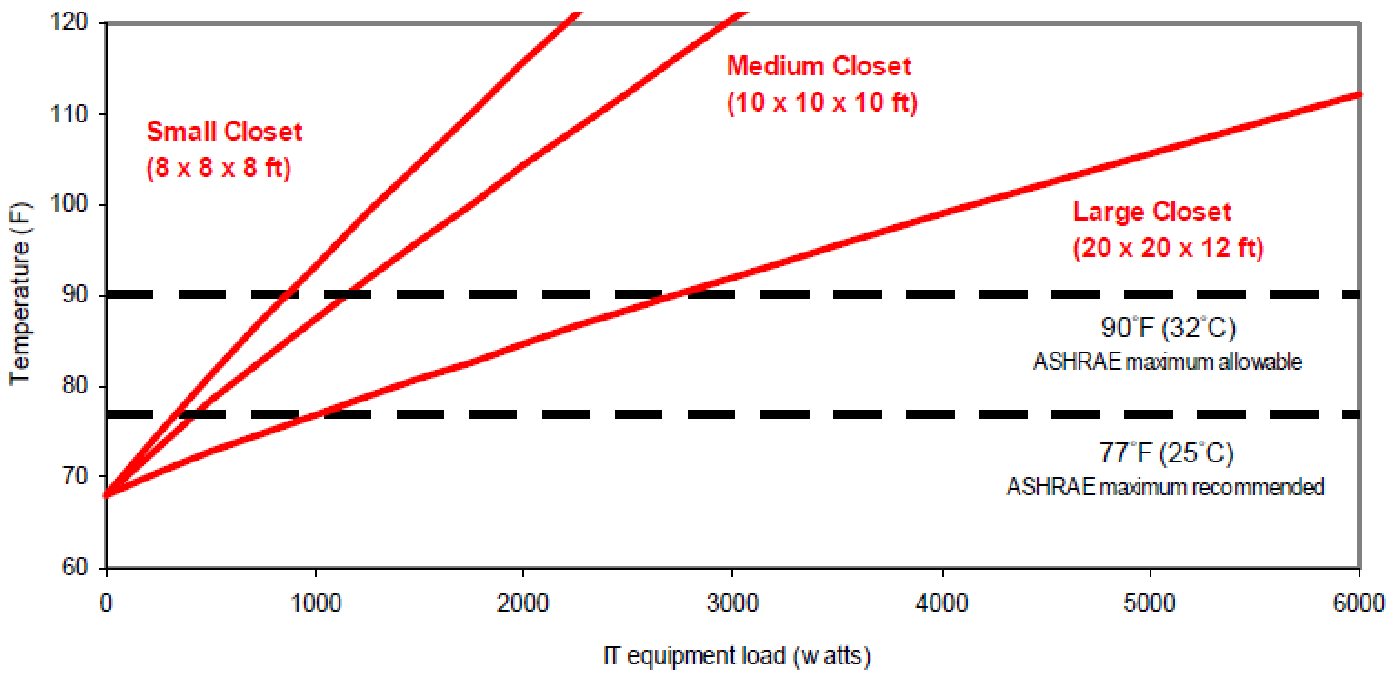

Figure 5

Effect of construction material on conduction cooling performance

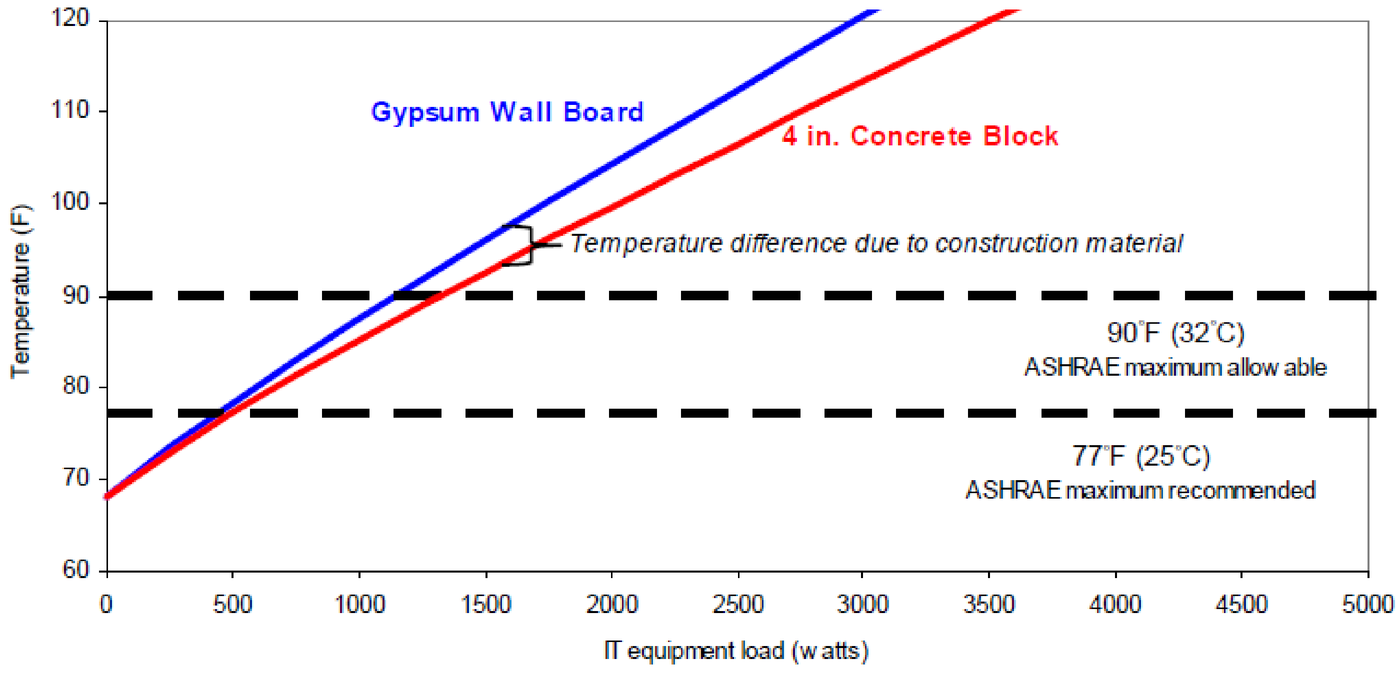

Comfort cooling

- Many buildings have an existing air conditioning system or combined heat and air conditioning system for creating a comfortable environment for personnel. These comfort cooling systems typically have air handling ductwork. It appears attractive to take advantage of this system by installing additional ducting to closets, in the same way that ducts are added when new offices or rooms are added. **However, simply adding ducts rarely solves closet cooling problems and often makes them worse.**
- Comfort cooling systems cycle on and off. A thermostat placed somewhere in the zone, but not in the closet.
- Furthermore, best practice for comfort cooling systems involves turning up the temperature set points on week nights and weekends to help conserve electricity. Some are actually shut off completely.
- By simply adding ductwork, one is forced to choose between wasting electricity on nights and weekends and making the temperature swings in the wiring closet even worse.

Dedicated air conditioner is appropriate When.. (even when ventilation appears)

- The ventilation air outside the closet contains significant dust or other contaminants
- The ventilation air outside the closet is subject to excessive temperature swings
- Practical constraints such as leases or cosmetics make it impossible to add ventilation ducts

Figure 8

Dedicated air conditioner selection

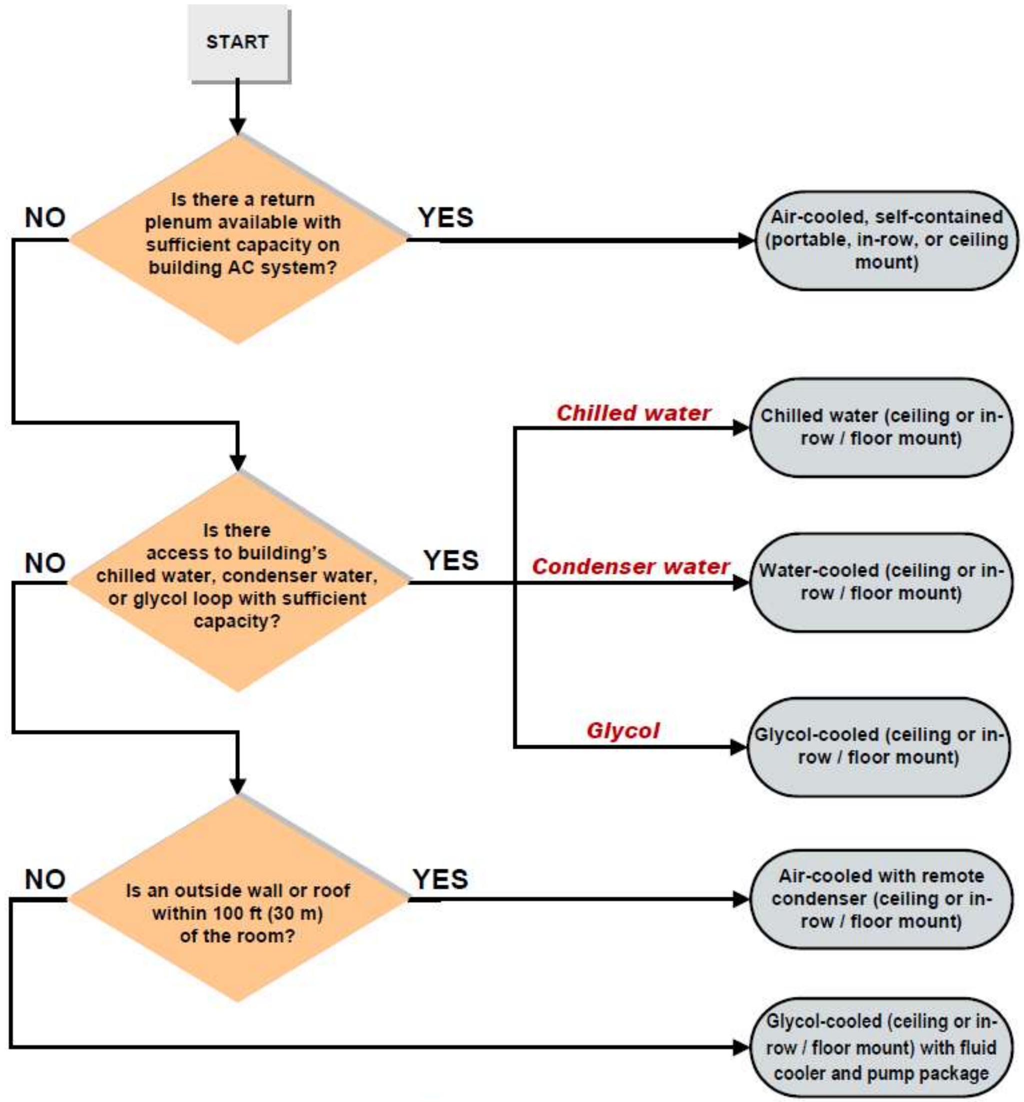

Ceiling-mount air conditioning unit

In-row air conditioning unit

Portable air conditioning unit

- When a UPS is installed, the IT equipment will continue to create heat during a power outage. Therefore the cooling system must continue to operate. **If the backup time of the UPS is less than 10 minutes, the thermal mass of the air and wall surfaces within the closet will keep the temperature within reasonable limits and no precautions need to be taken**. How ever, if the UPS is designed to provide runtime in excess of 10 minutes, then the cooling system must continue to operate during this period. This means that if fan-assisted ventilation or air conditioning is used, the fan or air conditioner must be powered by the UPS. The need to power the fan or air conditioner must be taken into consideration when sizing the UPS. In the case of fan-assisted ventilation, this is not a significant problem, but for air conditioners this may require a much larger UPS and battery.

Table 3

Ventilation system features and benefits

Feature	Benefit
Wall or ceiling mountable	More flexibility as one solution is compatible with many different closet types
Specified for calculated IT loads	Higher confidence that the solution will perform as expected
Remotely manageable	Lower mean time to recover (MTTR)
Multiple fan speeds	Ability to lower acoustic noise when max airflow is not required
More than one fan	Fan redundancy for fault tolerance
Tamper-proof mounting	Higher level of security
Easy installation	Requires minimal modification to closet environment and reduces the need for outside contractor involvement
Minimal assembly required	Fast, easy installation
Plug or hard wire configurations	Simple compliance with local electrical regulations
Broad capacity range	Ability to standardize on a single device for different installations
Specified and characterized for use with a UPS system	Higher overall system availability

Figure 9

Closet fan-assisted ventilation unit

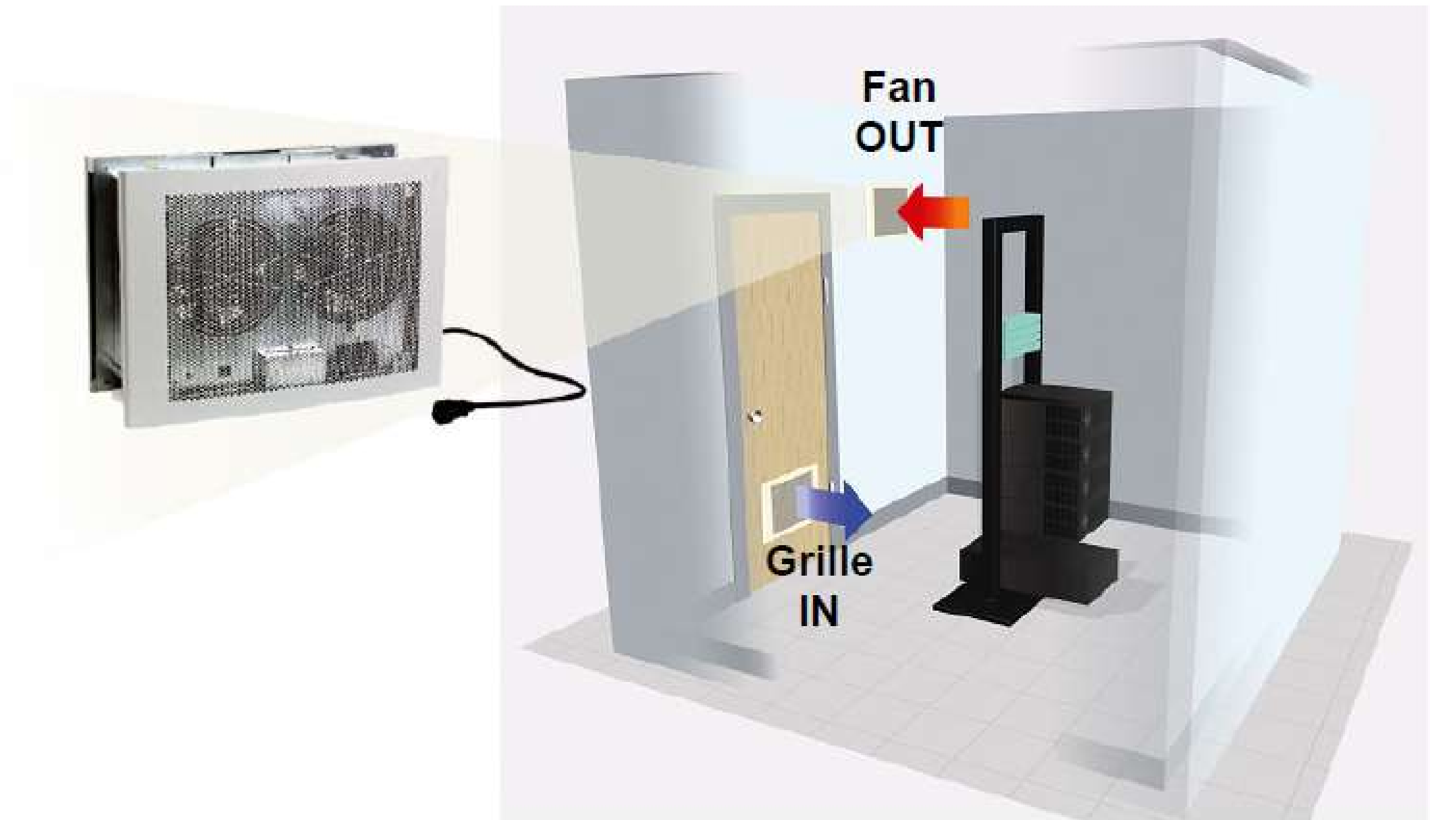

Computer rooms

- Computer rooms (**1-5** rack enclosures or equivalent using **3-30 kW** of total electricity) are often reused office space with varying levels of available space and ventilation. For rooms with very small electrical loads the building air conditioning system may be sufficient provided adequate ventilation is provided to the room.
- Most computer rooms require multiple portable or ceiling mounted systems and some heavily loaded rooms work well with a large floor mounted system if floor space is available.

Table 3 – Basic cooling system configurations for computer rooms

	Ceiling Mounted	Floor Mounted
Air Cooled DX System (2Piece)	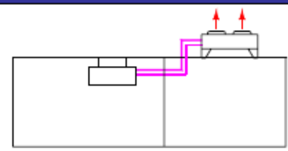 **Use if**: Building has roof access and a 10' (3m) floor to structural ceiling height. Roof is within 2 stories of IT environment. Single or multiple systems are a good choice for 6-30 kW loads	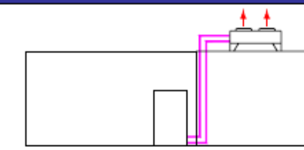 **Use if:** Building has roof access and available space for solution. Roof is within two stories of IT environment. Single large system OK for loads greater than 25kW if it fits.
Air Cooled Self Contained System	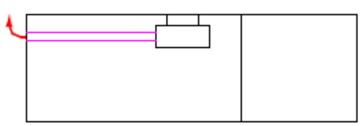 **Use if:** IT environment has dropped ceiling or ducts can be installed for condenser air. Ensure 10' floor to structural ceiling height. Single or multiple ceiling mounted systems are a good choice for 6-30 kW loads	 **Use if**: Dropped ceiling for condenser air ducts exists. Single or multiple portable systems are OK for less than 12 kW of equipment
Glycol Cooled Systems	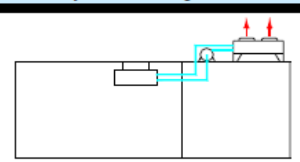 **Use if**: Building has roof access and a 10' (3m) floor to structural ceiling height. Computer room is a long distance to outdoors.	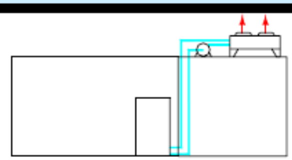 **Use if:** Building has roof access and available space for solution. Single large system OK for loads greater than 25kW if it fits. Use if free cooling is desired in areas with cold winters.
Water Cooled Systems	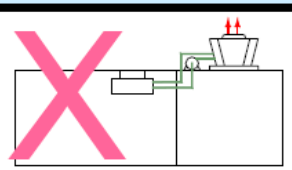 Condenser water is usually not routed far from mechanical room. Use is a possibility in some hi-rise buildings.	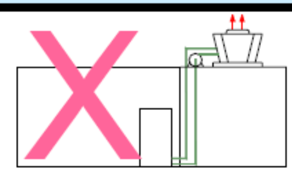 Condenser water is usually not routed far from mechanical room. Use is a possibility in some hi-rise buildings.
Chilled Water Systems	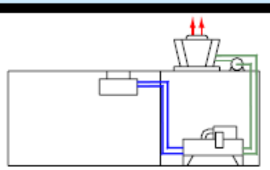 **Use if:** The only cooling source is chilled water, no possible location for outdoor condensers. Some hi-rise buildings.	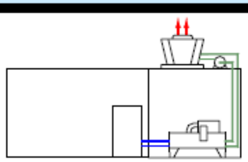 **Use if:** The only cooling source is chilled water, no possible location for outdoor condensers. Some hi-rise buildings.

Small data centers

- Small data centers (**5-20** rack enclosures or equivalent using **7-100 kW** of total electricity) are usually purpose-built rooms with sufficient space and ventilation for IT equipment. If the amount of electricity consumed by IT equipment is high (more than 3 kW per rack enclosure) then more space may have to be devoted to the room’s cooling solution. Most small data centers use ceiling mounted and large floor mounted cooling systems. Portable systems are used as needed for hot spots and temporary cooling.

Table 5 – Basic cooling system configurations for medium data centers

	Ceiling Mounted	Floor Mounted
Air Cooled DX System (2Piece)	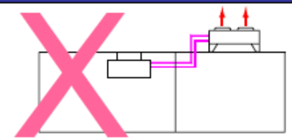 Insufficient capacity per unit	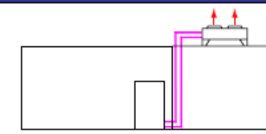 **Use if:** Building has roof access and roof is adjacent the data center. For smaller and low-density data centers. Lowest cost solution.
Air Cooled Self Contained System	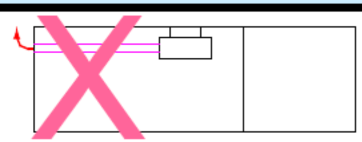 Insufficient capacity per unit	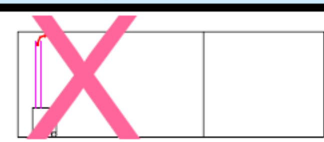 Insufficient capacity per unit
Glycol Cooled Systems	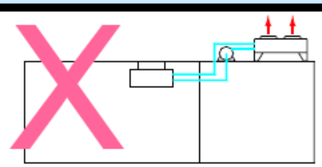 Insufficient capacity per unit	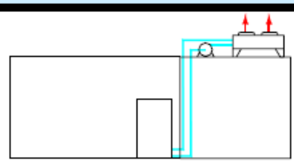 **Use if:** Building has roof access but data center is far away. No chilled water or condenser water already piped into data center. Use if free cooling is desired in areas with cold winters.
Water Cooled Systems	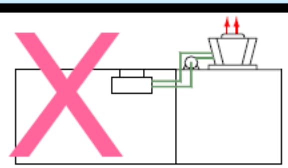 Insufficient capacity per unit	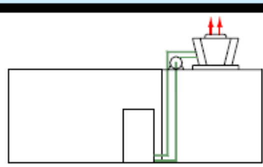 **Use if:** Building condenser water system has more available capacity or lower usage cost (high rise) than chilled water
Chilled Water Systems	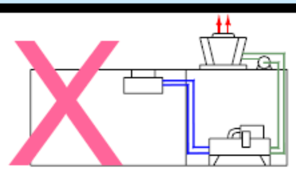 Insufficient capacity per unit	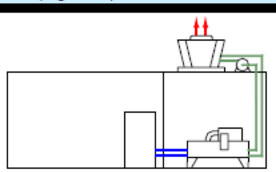 **Use if:** Building has reliable chilled water supply without setbacks.

Large data centers

- Large data centers (> **100 rack** enclosures or equivalent using >**200 kW** total electricity) are purpose-built rooms optimized for the availability of IT equipment. Rising power densities are forcing increased space allocation for cooling solutions. Large data centers use multiple large floor mounted cooling systems or very large custom rooftop central cooling systems (not shown in the diagrams below).

Table 6 – Basic cooling system configurations for large data centers

	Ceiling Mounted	Floor Mounted
Air Cooled DX System (2Piece)	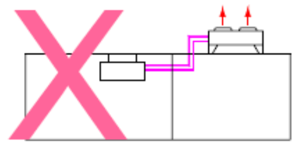 Insufficient capacity per unit	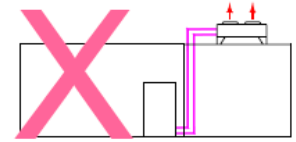 Large data centers usually have dedicated chilled water systems. Only use if no chilled water system exists or no capacity is available.
Air Cooled Self Contained System (1Piece)	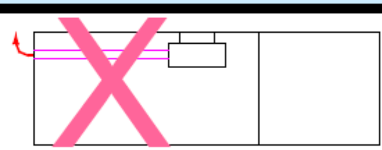 Insufficient capacity per unit	 Insufficient capacity per unit
Glycol Cooled Systems	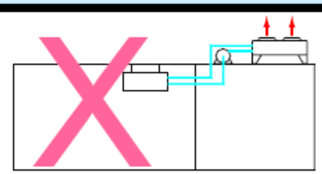 Insufficient capacity per unit	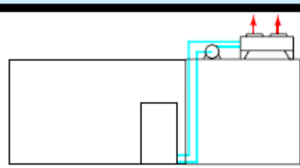 **Use if:** Waterside free cooling is desired in areas with cold winters.
Water Cooled Systems	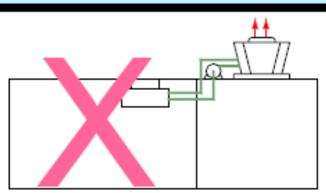 Insufficient capacity per unit	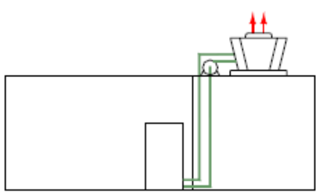 **Use if:** Building condenser water system has more available capacity or lower usage cost than chilled water
Chilled Water Systems	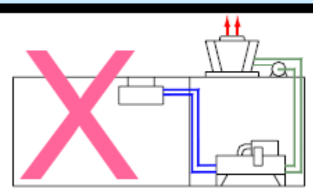 Insufficient capacity per unit	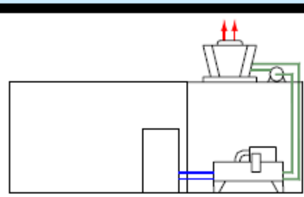 **Use if:** Data center has mission critical chilled water supply.

The nine types of cooling systems

- Every cooling distribution system has a supply system and a return system. The supply system distributes the cool air from the CRAC unit to the load, and the return system takes the exhaust air from the loads back to the CRAC. For both the supply and the return, there are three basic methods used to convey air between the CRAC and the load, which are:
  - Flooded
  - Locally Ducted
  - Fully Ducted

- In a Flooded distribution system, the CRAC and the loads eject or draw in bulk air from the room, without any special ductwork in between them. In a Locally Ducted distribution system, air is provided or returned via ducts which have vents located near the loads. In a Fully Ducted system, supply or return air is directly ducted into or out of the loads.
- A critical goal of a data center cooling system is to separate the equipment exhaust air from the equipment intake air in order to prevent equipment from overheating. This separation also significantly increases the efficiency and capacity of the cooling system. When equipment power density increases, the corresponding increase in exhaust air volume and intake air volume makes it more difficult to prevent equipment from drawing exhaust air from itself or neighboring equipment into its intake. For this reason partial or complete ducting of the supply air to the equipment intake or return air from the equipment exhaust becomes necessary as power density increases.

Table 1

The 9 types of air distribution (traditional room-based cooling implementations)

	Flooded return	Targeted return	Contained return
Flooded supply	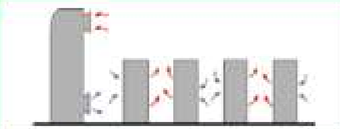		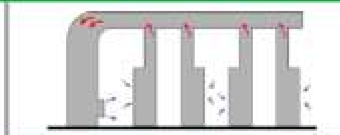
	Small LAN rooms < 40kW Not recommended for most data centers Low cost, simple installation Least energy efficient of all air distribution architectures because 100% of the cold supply air is allowed to mix with hot return air. Supply air temperature extremely unpredictable above Distribution type can cool up to 3kW per rack	General use Not recommended for most data centers Low cost, ease of install More energy efficient than flooded return since 40-70% of IT hot exhaust air is captured and delivered back to the cooling unit. Supply air more predictable than flooded supply since less hot air is allowed to mix with cold supply air. Distribution type can cool up to 6kW per rack	Large data center / colocation Upgradeable (vendor specific) Most energy efficient of all air distribution architectures since it allows increased cooling unit supply temp resulting in increased economizer hours. 70-100% of IT equipment hot exhaust air is captured and delivered back to the cooling unit. Supply air is most predictable since no hot air is allowed to mix with cold supply air. Distribution type can cool up to 30kW per rack
Targeted supply	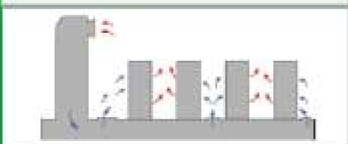		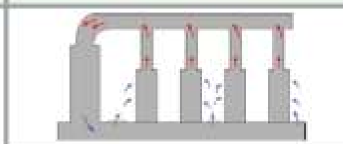
	Data centers with static power densities Not recommended for new designs – unable to keep up with power density projections More energy efficient than flooded supply since more IT equipment hot exhaust air is diverted back to the cooling unit. Distribution type can cool up to 6kW per rack	Small to medium data centers More energy efficient than flooded return since 60-80% of IT equipment hot exhaust air is captured and delivered back to the cooling unit. Supply air more predictable since less hot air is allowed to mix with cold supply air. Distribution type can cool up to 8kW per rack	Hot spot problem solver Upgradeable (vendor specific) More efficient than targeted supply and return since 70-100% of IT equipment hot exhaust air is captured and delivered back to the cooling unit. Supply air is most predictable since no hot air is allowed to mix with cold supply air. Allows increased cooling unit supply temp resulting in increased economizer hours. Distribution type can cool up to 30kW per rack
Contained supply		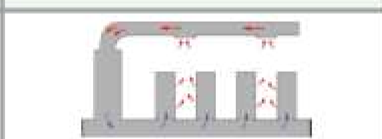	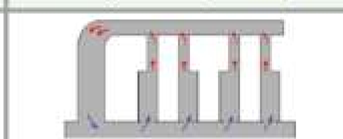
	Mainframes / racks with vertical airflow More energy efficient than targeted supply but less efficient than contained return. Containing the supply air, forces the rest of the room to become the hot aisle which limits the number of economizer hours. Supply air is more predictable since little hot air is allowed to mix with cold supply air. Distribution type can cool up to 30kW per rack	Mainframes / racks with vertical airflow More energy efficient than targeted supply but less efficient than contained return. Containing the supply air, forces the rest of the room to become the hot aisle which limits the number of economizer hours. Supply air is most predictable since no hot air is allowed to mix with cold supply air. Distribution type can cool up to 30kW per rack	Harsh non-data center environments Slightly less efficient than contained return with flooded or targeted supply - requires more fan energy. Allows increased cooling unit supply temp resulting in increased economizer hours. Distribution type can cool up to 30kW per rack

- APC white paper #55

Table 2

Non-traditional implementations

	Flooded return	Targeted return	Contained return
Flooded supply	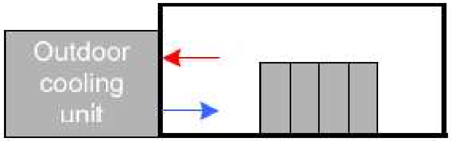 **Hard floor, cooling unit located outdoors** Not recommended for most data centers. Not effective because air mixing prevents predictable IT inlet temperatures.	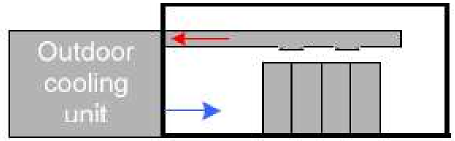 **Hard floor, cooling unit located outdoors** Not recommended for most data centers. Not effective because air mixing prevents predictable IT inlet temperatures.	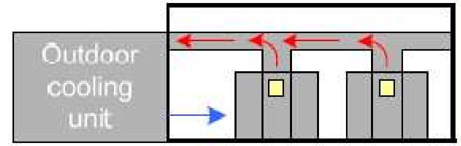 **Hard floor, cooling unit located outdoors** Recommended for new data centers. Variable speed fans on cooling units controlled by IT temperature.
Targeted supply	**No non-traditional alternative**	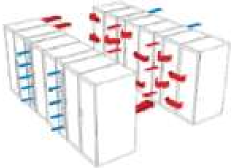 **Hard floor, row-based cooling units** Recommended for data centers below 1MW. Variable speed fans on cooling units controlled by IT temperature.	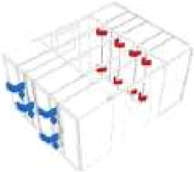 **Hard floor, row-based cooling units** Recommended for data centers below 1MW. Variable speed fans on cooling units controlled IT temperature.
Contained supply	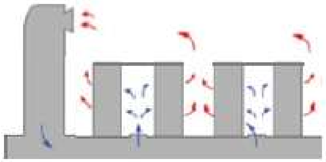 **Raised floor, perimeter cooling units** Not recommended for new data centers. Good solution for existing data centers. Variable speed fans on cooling units controlled by pressure and active tiles controlled by IT temperature.	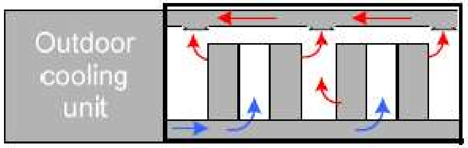 **Raised floor, cooling unit located outdoors** Targeted return doesn't add much value since supply is contained therefore not recommended. Variable speed fans on cooling units controlled by pressure and active tiles controlled by IT temperature.	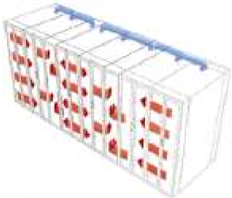 **Hard floor, row-based cooling units** Only recommended for harsh environments or existing data centers where complete containment is required for a single row of racks (e.g. squeezing a row into an existing hot aisle). Variable speed fans on cooling units controlled by IT temperature.

Figure 1

Examples of containment for new data centers - Flooded supply and contained return

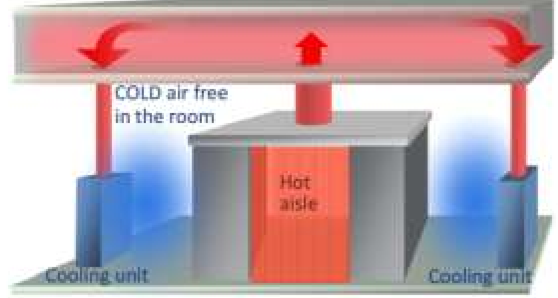

Flooded supply and contained return with cooling units located on perimeter

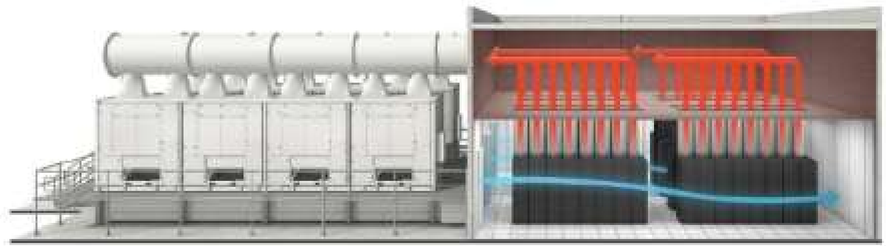

Flooded supply and contained return with cooling units located on outdoors

Figure 2

Examples of containment for new data centers - Targeted supply and contained return

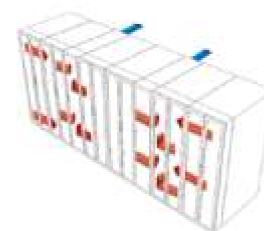

Targeted supply and contained return with cooling units located in the row

Figure 3

Examples of containment for new data centers - Contained supply and contained return

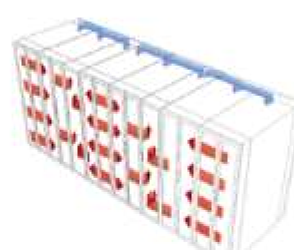

Contained supply and contained return with cooling units located in the row

Types of cooling in the raised-floor environment

- The raised floor approach is applicable when:
  - There is already an existing raised floor in the facility that can be reused
  - Mainframe computers with under floor air intake are being installed
  - There is a need to run significant amounts of water piping throughout the computing area

Cooling system design considerations

- Once the appropriate type of cooling system is selected, there are other elements that must be integrated into the system design. These include the following factors
  - Layout of racks in alternating rows
  - Location of CRAC units
  - Quantity and location of vents
  - Sizing of ductwork
  - Proper internal configuration of racks

Figure 1A – Rack-mounted Fully Ducted air return unit

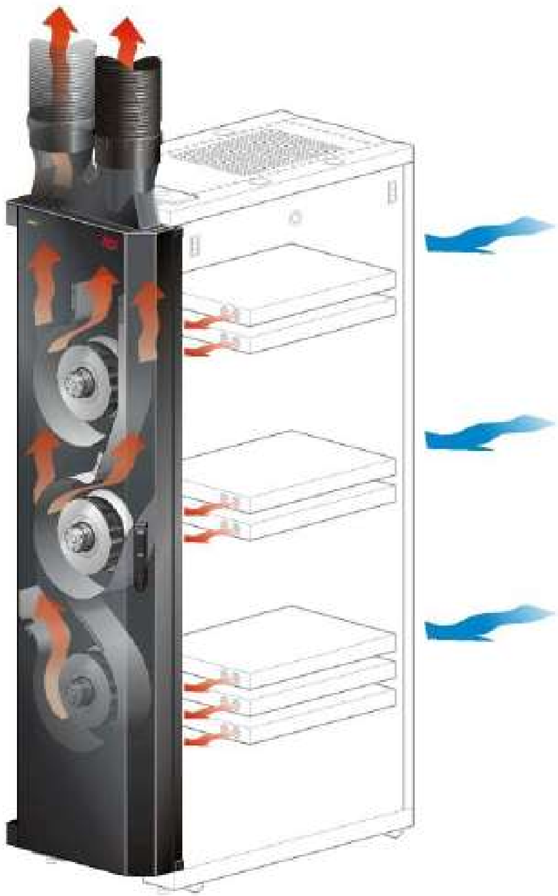

APC model ACF101BLK Air Removal Unit

Figure 1B – Rack-mounted Fully Ducted air supply unit

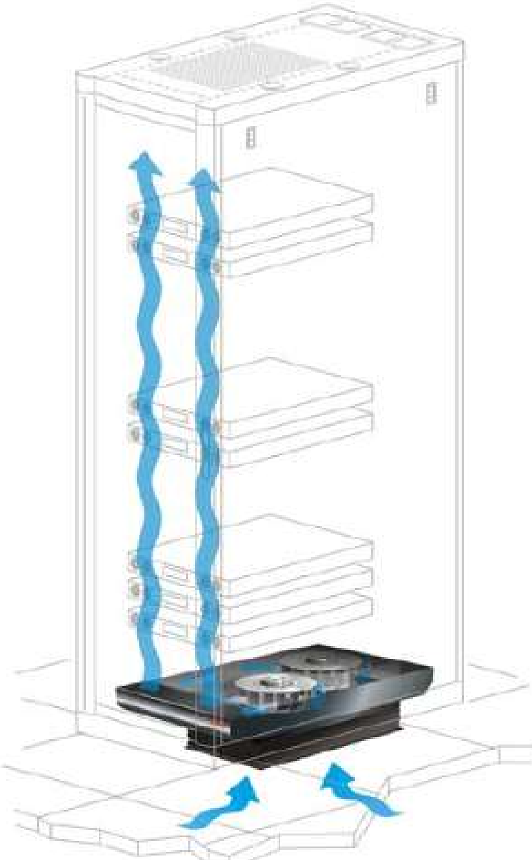

APC model ACF001 Air Distribution Unit

Flooded supply ducted return components

Figure 2 – *Integrated rack cooling system (multi rack)*

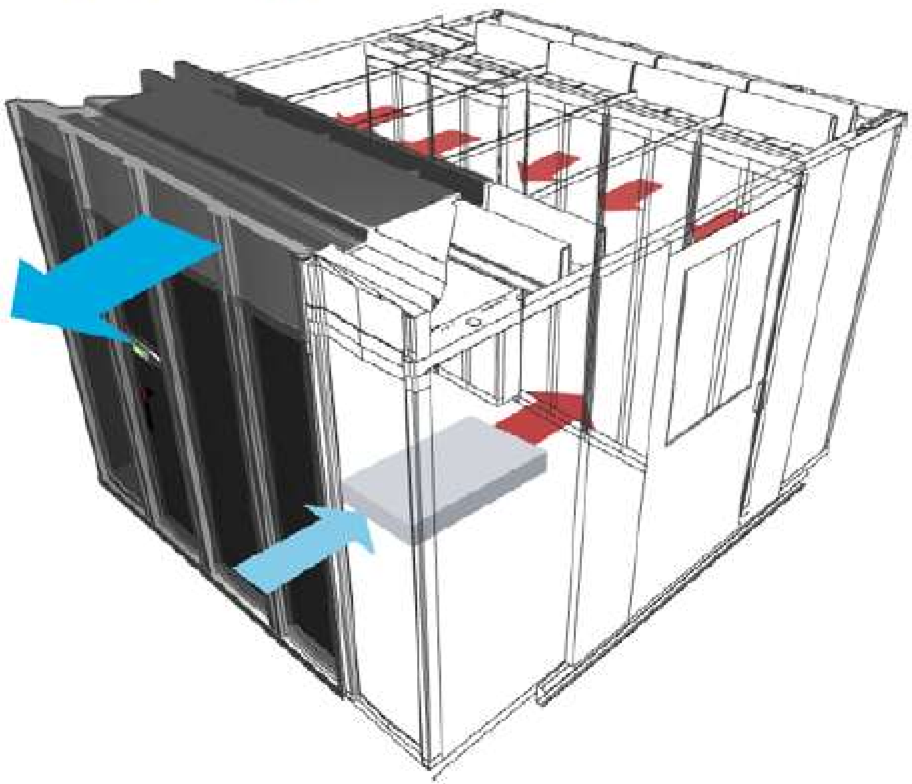

Figure 3 – *Integrated rack cooling system (single rack)*

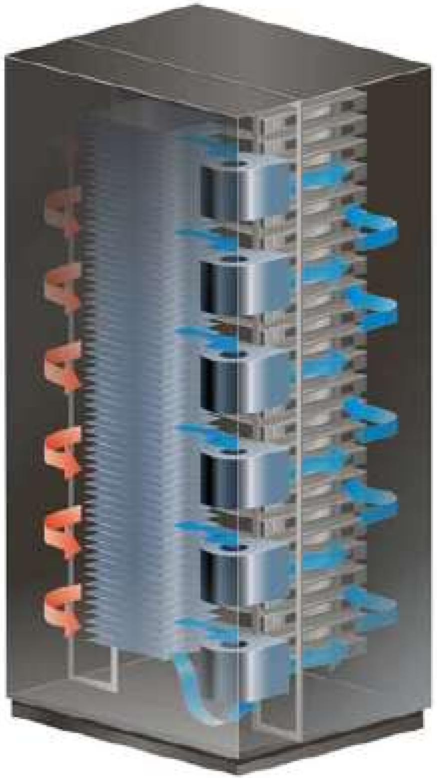

Custom CRAH Unit (Large)

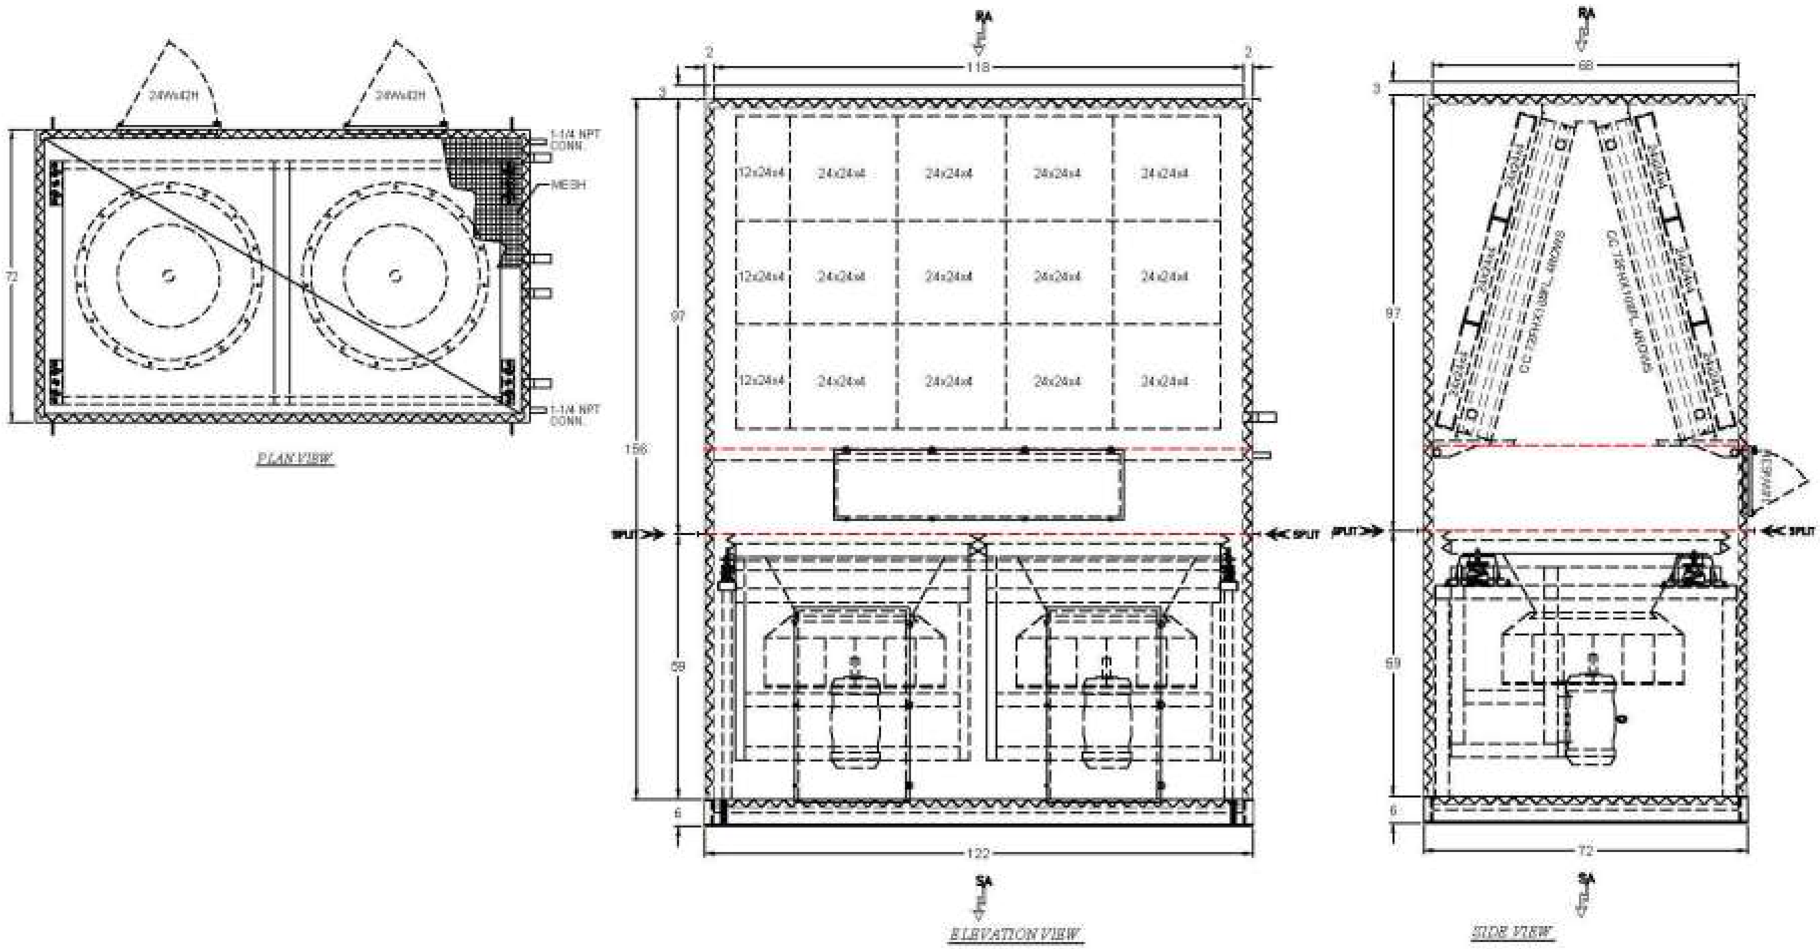

Example CRAH Unit Comparison

- 34% less water flow
- 13% less fan energy
  - More if you consider the supply air temperature and airflow issues
- Excess fan capacity on new units
- 36% higher cost for units, but
  - Fewer piping connections
  - Fewer electrical connections
  - Fewer control panels
  - No need for control gateway
  - Can use the existing distribution piping and pumps (case study)
  - Can use high quality sensors and place them where they make sense
- Possibly less turbulence at discharge?

Efficiency improvements of AC system

On refrigeration effect and refrigeration capacity

- The refrigeration capacity of the compressor Q is given by: ($q_e = h_1 - h_4$)

$$Q_e = \dot{m} . q_e$$

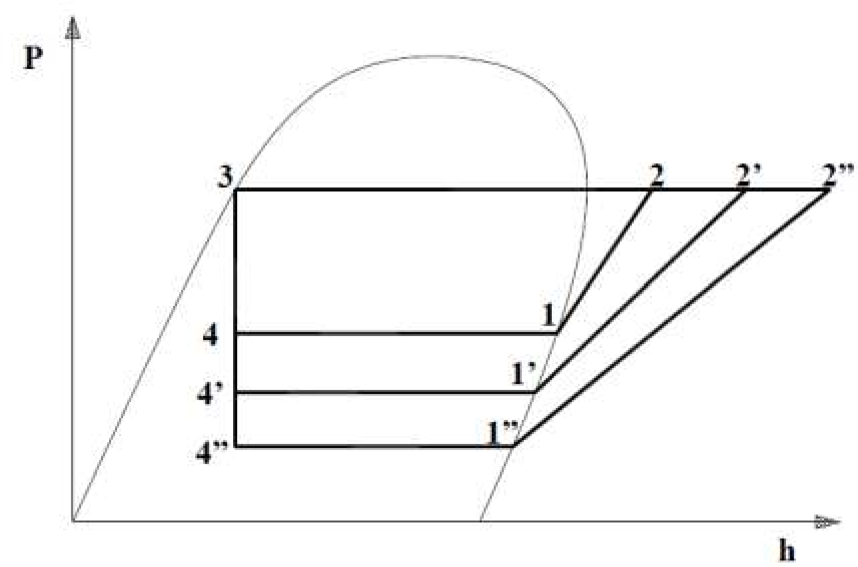

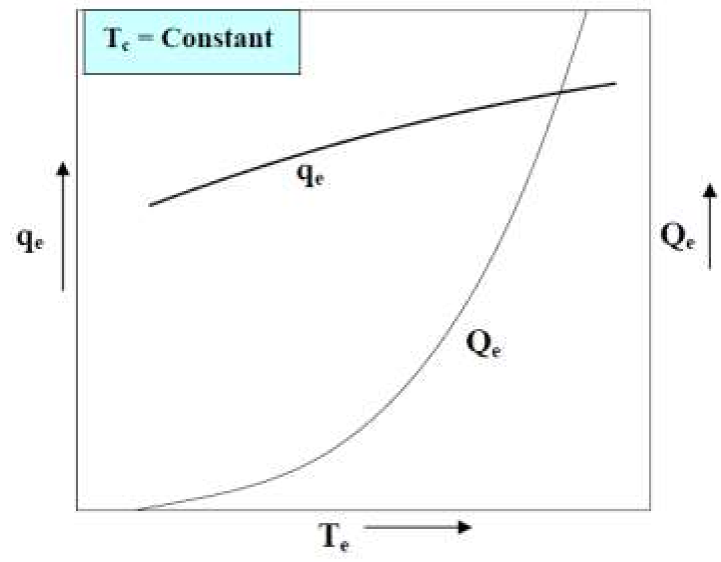

On work of compression and power requirement

- The power input to the compressor is given by: $W_c = \dot{m}.\Delta h_c$
- $(\Delta h_c = h_2 - h_1)$
- For a given clearance ratio and condenser temperature, the volumetric efficiency and hence the mass flow rate becomes zero at a lower limiting value of evaporator temperature ($T_e = T_{e,lim}$).

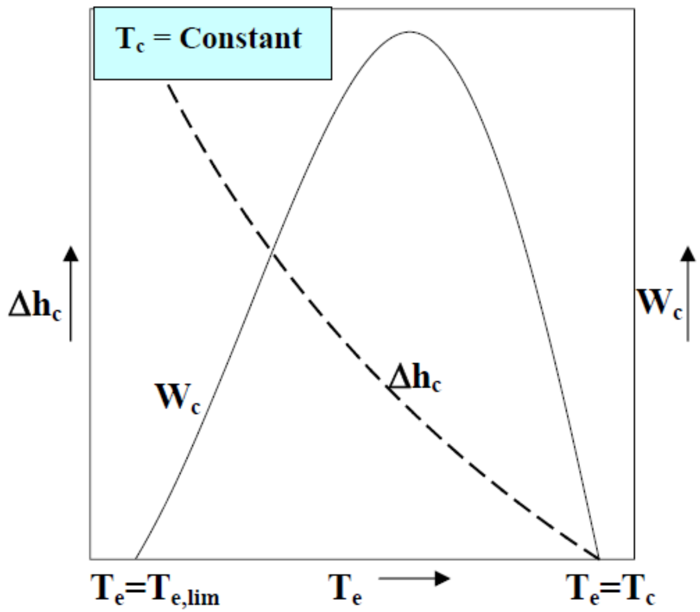

On COP and volume flow rate per unit capacity

- As evaporator temperature increases the specific volume of the refrigerant at compressor inlet reduces rapidly and the refrigerant effect increases marginally. Due to the combined effect of these two (and volumetric efficiency), the volume flow rate of refrigerant per unit capacity reduces sharply with evaporator temperature as shown in the following. This implies that for a given refrigeration capacity, the required volumetric flow rate and hence the size of the compressor becomes very large at very low evaporator temperatures.

$$V = \frac{\eta_{V,cl} \cdot \dot{V}_{sw}}{Q_e} = \frac{v_e}{q_e}$$

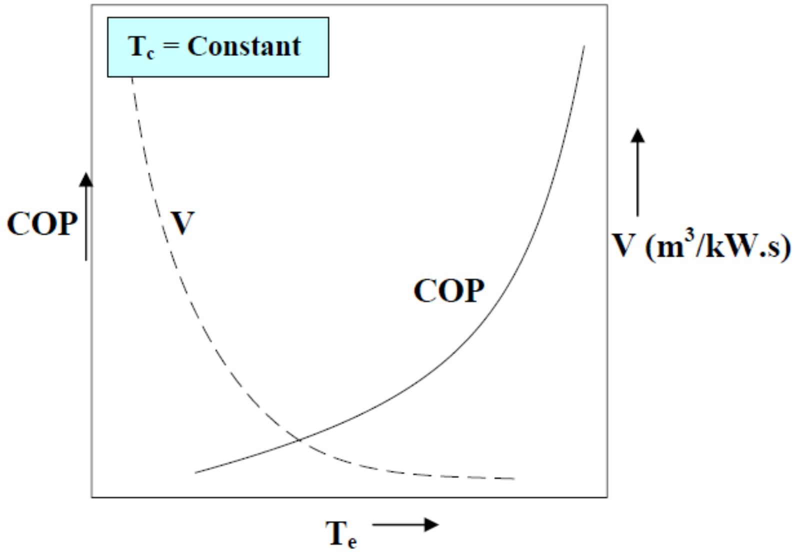

Effect of condenser temperature

- Atmospheric air is the cooling medium for most of the refrigeration systems. Since the ambient temperature at a location can vary over a wide range, the heat rejection temperature (i.e., the condensing temperature) may also vary widely. This affects the performance of the compressor and hence the refrigeration system. The effect of condensing temperature on compressor performance can be studied by keeping evaporator temperature constant.

On volumetric efficiency and refrigerant mass flow rate

- At a constant evaporator temperature as the condensing temperature increases, the pressure ratio increases, hence, both the volumetric efficiency and mass flow rate decrease as shown in the figure. However, the effect of condensing temperature on mass flow rate is not as significant as the evaporator temperature as the specific volume of refrigerant at compressor inlet is independent of condensing temperature.

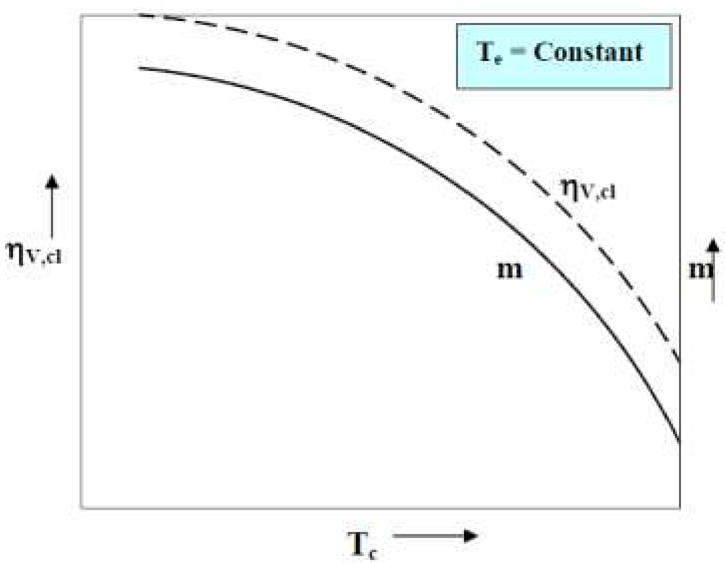

On refrigeration effect and refrigeration capacity

- Since the evaporator enthalpy remains constant at a constant evaporator temperature, the refrigeration effect decreases with increase in condensing temperature as shown in Figure. The refrigeration capacity (Q_e) also reduces with increase in condensing temperature as both the mass flow rate and refrigeration effect decrease as shown.

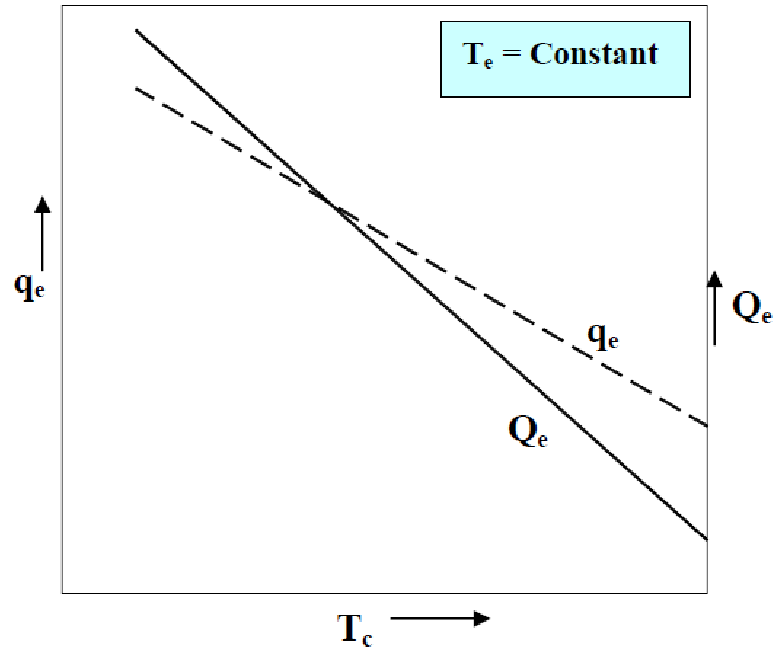

On work of compression and power requirement

- The work of compression is zero when the condenser temperature is equal to the evaporator temperature, on the other hand at a limiting condensing temperature the mass flow rate of refrigerant becomes zero as the clearance volumetric efficiency becomes zero as explained before.

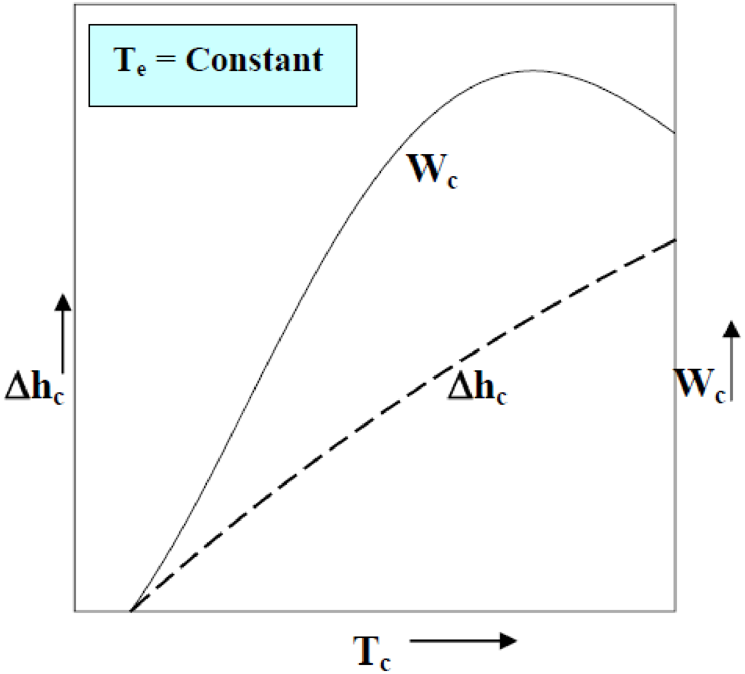

On COP and volume flow rate per unit capacity

- As condensing temperature increases the refrigeration effect reduces marginally and work of compression increases, as a result the COP reduces as. Even though the specific volume at compressor inlet is independent of condensing temperature, since the refrigeration effect decreases with increase in condensing temperature, the volume flow rate of refrigerant per unit capacity increases as condenser temperature increases as shown.

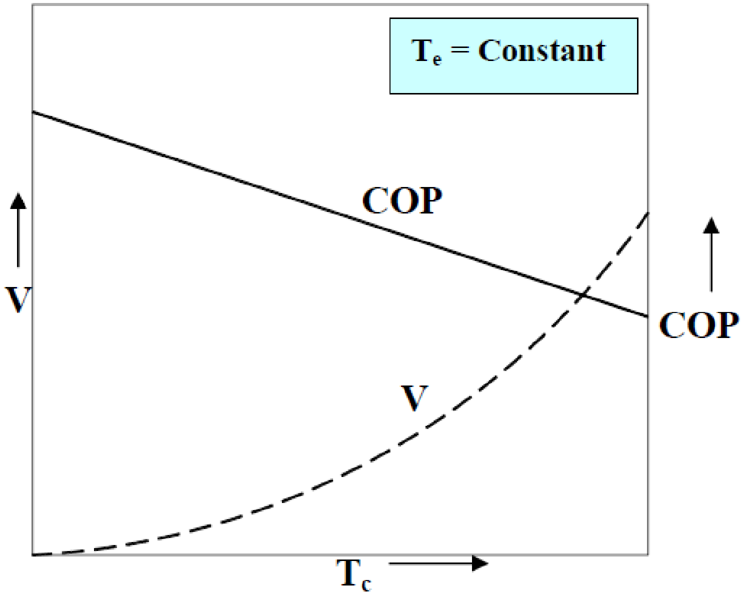

Compressor discharge temperature

- If the compressor discharge temperature is very high then it may result in breakdown of the lubricating oil, causing excessive wear and reduced life of the compressor valves (mainly the discharge valve). In hermetic compressors, the high discharge temperature adversely affects the motor insulation (unless the insulation is designed for high temperatures).

$$Pv^{\gamma} = \text{constant} \quad \text{and} \quad Pv = RT$$

- Then the discharge temperature, T is given by:

$$T_d = T_e \left(\frac{P_c}{P_e} \right)^{\frac{\gamma - 1}{\gamma}}$$

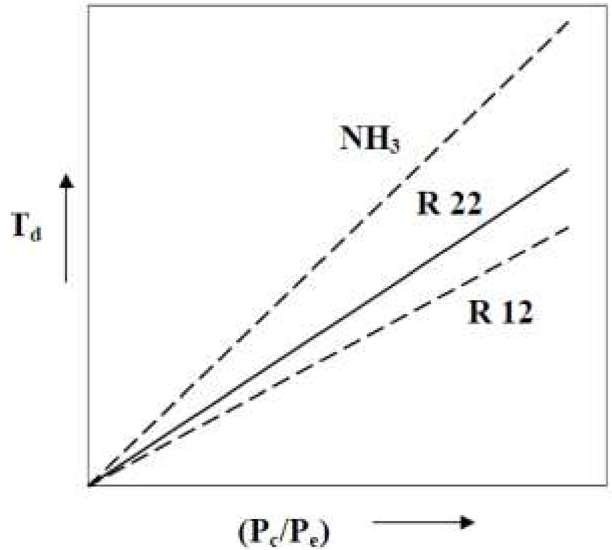

Actual compression process

- Actual compression processes deviate from ideal compression processes due to:
  - Heat transfer between the refrigerant and surroundings during compression and expansion, which makes these processes non-adiabatic
  - Frictional pressure drops in connecting lines and across suction and discharge valves
  - Losses due to leakage

Monitoring Index

Data center Energy Monitoring Index

ASHRAE Transactions 2010, Vol. 116, Part 1. pp. 9-23

- DCiE (Data center infrastructure efficiency

$$\text{DCiE} = \left[\frac{\text{IT-Equipment Power}}{\text{Total Facility Power}}\right]100\%$$

  - One benefit of using the DCiE rather than the PUE is that it has an easily understood scale of 0-100%

- PUE (Power usage effectiveness)

Rating of the DCiE

Rating	DCiE
Ideal (maximum)	100
State-of-the-Art	85
Best Practice	70
Improved Operations	60
Current Trend	55
Typical (average)	50

Average PUE for Different Size Data Center

Data Center Size	Average PUE
RFA < 10,000 ft^2	~2.8
10,000 ft^2 < RFA < 30,000 ft^2	~2.20
RFA > 30,000 ft^2	~2.1

In this context, *all other power* consumed in the data center – in other words, all power *not* consumed by the IT load – is considered "loss." These non-IT power consumptions, or losses, include:

- The internal inefficiencies of the power system (power path devices such as UPS, PDUs, wiring, etc.), dispersed as heat
- All power consumed by the cooling system
- All power consumed by other data center physical infrastructure subsystems (small by comparison, and not shown in **Figure 7**)

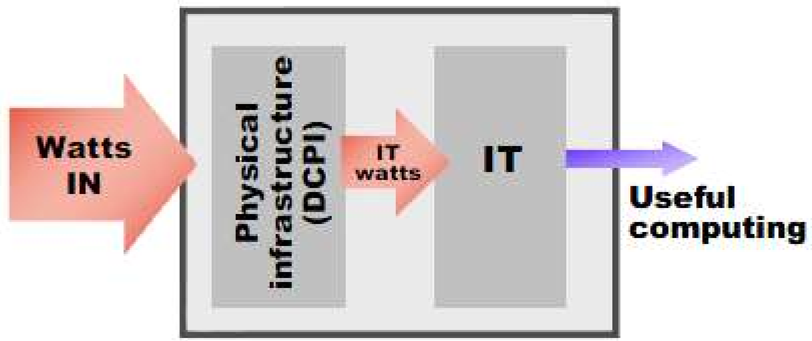

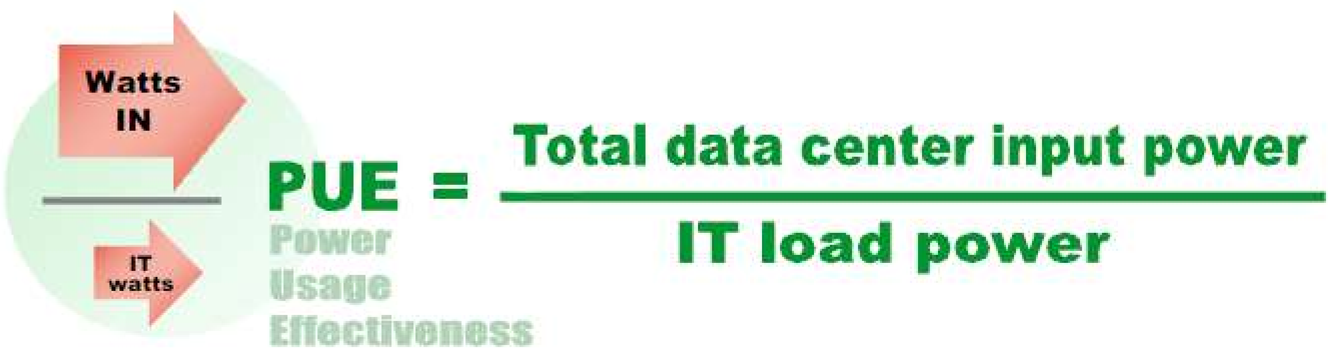

Losses..

Figure 8

Definition of "losses" in data center power consumption

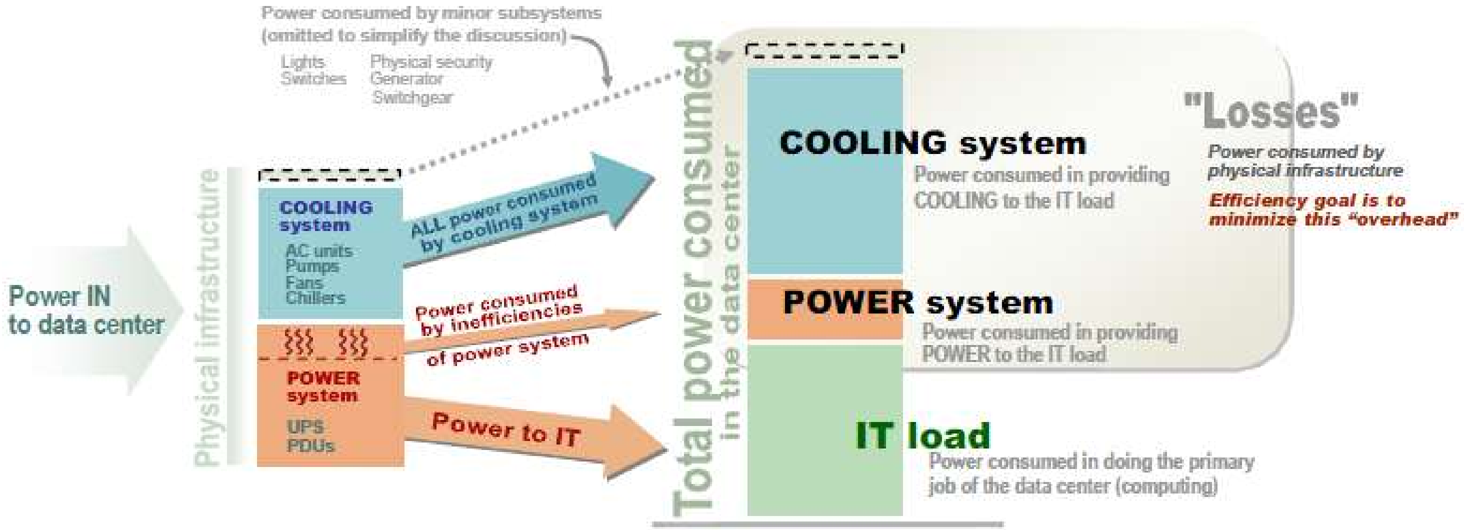

Fixed vs. proportional loss

Of the power consumed by the power and cooling systems – the "losses" in **Figure 8** – some stays the same no matter how large or small the IT load is, and some varies in proportion to the size of the IT load. These two components of consumed power (loss) are called **fixed loss** and **proportional loss**.[2]

- **Fixed loss** – Always the same amount no matter what the load. Fixed loss is power that is consumed whenever the device or system is running, regardless of how much load is present. Reducing the load does not change this fixed component of loss. Examples of devices with a large component of fixed loss are transformers and fixed-speed fans. The presence of fixed loss is the reason efficiency is better at high loads (where fixed loss is a small proportion of total power) and worse at low loads (where fixed loss is a large proportion of total power) – see **Figure 11** below. **Reducing fixed loss, by improved device efficiencies and/or better system configuration, is the most effective way to increase power and cooling efficiency**.
- **Proportional loss** – Directly proportional to the load on the device. Doubling the load will double the proportional loss. Reducing the load 75% will reduce the proportional loss 75%. Examples of devices with a large component of proportional loss are variable-speed fans and pumps.

Table 1

Availability issues addressed by optimized power and cooling infrastructure

Availability threat	Why?	How solved
Human Error	Human error has historically been a significant cause of downtime in data centers. With more complexity and change – including changes that can't be seen – comes a greater risk of mistakes and oversights.	The systems described in this paper – such as cooling that responds to local demand, rack-level instrumentation, and capacity management – build intelligence into power and cooling, minimizing the need for human interpretation and intervention.
Unpredictable cooling	Traditional room-based cooling is not agile enough to handle the dynamic high-density loads of a virtualized environment.	Row-based, managed cooling tightly controls delivery, not only in location but also in amount. Load conditions that threaten to exceed local cooling capacity can be proactively identified by the capacity management system, and corrective action can be taken.
Loss of cooling redundancy	Dynamic loads can locally raise cooling demand to the point where there is not enough capacity to cover scheduled or unscheduled downtime of cooling equipment.	Row-based cooling units can be installed to provide the desired redundancy for that row. The capacity management system can warn of insufficient or lost redundancy due to subsequent changes in loading.
Power overloads	With increased frequency of changes in power demand – due to reconfiguration of physical servers or the migration of virtual servers – comes increased risk of branch circuit loads creeping up close to the breaker-trip limit.	The capacity management system can warn of load imbalances before they pose an availability risk.

Data center Energy Monitoring Index Conti..

- RCI (RACK COOLING INDEX)
  - RCI is a measure of how effectively equipment racks are cooled within a given thermal guideline, both at the high end and at the low end of the temperature range.
  - ASHRAE recommended the allowable temperature ranges 64.4–80.6°F (18–27°C) and 59.0–89.6°F (15–32°C), respectively.
  - The RCI “compresses” the equipment intake temperatures into two numbers—the RCI_{HI} and the RCI_{LO}.

$$RCI_{HI} = \left[1 - \frac{\text{Total Over-Temp}}{\text{Max Allowable Over-Temp}}\right]100\%$$

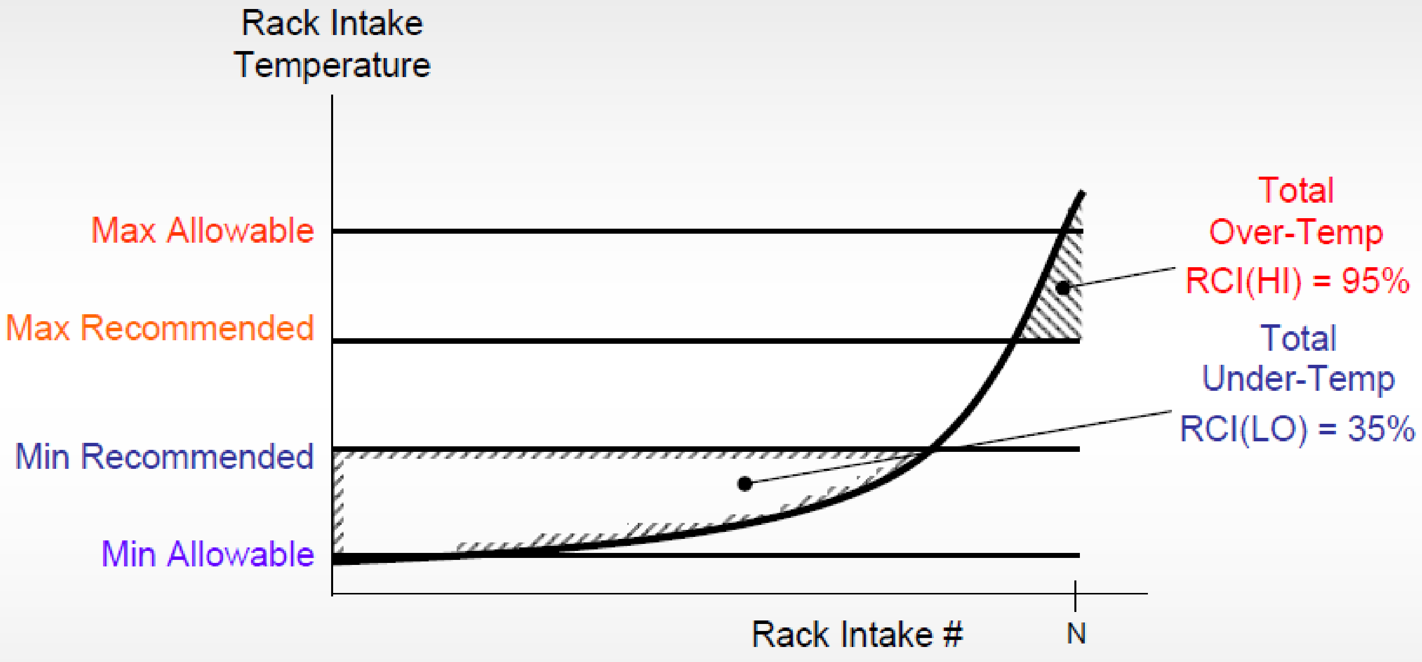

- RCI is a measure of compliance with Temperature Specifications
- 100% mean ideal conditions; no over- or under-temperatures

RTI
(RETURN TEMPERATURE INDEX)

- A measure of the net level of by-pass air or net level of recirculation air in the equipment room. Both effects are detrimental to the overall energy and thermal performance of the space.
- By-pass air does not contribute to the cooling of the electronic equipment, and it depresses the return air temperature.
- Recirculation is one of the main reasons for hot spots or areas significantly hotter than the ambient temperature.

$$RTI = (\Delta T_{AHU}/\Delta T_{Equip})\ 100\% = (V_{Equip}/V_{AHU})\ 100\% \quad (3)$$

where

RTI	=	Return temperature index
ΔT_{AHU}	=	Temperature drop across the air-handler units (airflow weighted average)
ΔT_{Equip}	=	Temperature rise across the IT-equipment (airflow weighted average)
V_{AHU}	=	Total airflow rate through the air-handler units
V_{Equip}	=	Total airflow rate through the IT-equipment

Interpretation of the RTI

Interpretation	RTI
Balanced	100%
Net Recirculation Air	>100%
Net By-Pass Air	<100%

Three key factors in lowering energy usage and maximizing performance

HVAC&R Res. J. 9, No. 2, 137–152 (2003)

- Minimize the infiltration of hot air into the rack inlets.
- Minimize the mixing of hot return air with cold-air streams prior to return to the CRAC units.
- Minimize the short-circuiting of cold air to the CRAC inlet.

- The typical temperature difference in DP (data processing) equipment and CRAC units is 10 °C.
- To maintain the 10 °C air temperature difference as the power dissipation of the DP racks increases, the flow rate through the rack must increase in direct proportion. This is not always the case since many DP racks with higher heat loads are now experiencing temperature differences as high as 15 °C to 20 °C.
- For a rack heat load of 4 kW, the resulting flow rate was 20 m^3/min to maintain the 10 °C temperature difference across the DP rack.

Perforated tile flow rates (m^3/min) to maintain an average inlet air temperature of 10 °C

Rack heat load (kW)	*Perforated-tile flow rate* (m^3/min)		*Increase in flow rate from tile* (%)
	For $\Delta T = 20.8°C$ *across rack*	*For* $\Delta T = 10°C$ *across rack*	
4	12.5	18	44
8	18	20.5	14
12	20	25.5	28

Cooling Technology in Datacom

- Air cooling

Air Cooling – Conditioned air is supplied to the inlets of the rack / cabinet for convection cooling of the heat rejected by the components of the electronic equipment within the rack. It is understood that within the rack, the transport of heat from the actual source component (e.g., CPU) within the rack itself can be either liquid or air based, but the heat rejection media from the rack to the terminal cooling device outside of the rack is air.

- Liquid cooling

Liquid Cooling – Conditioned liquid (e.g., water, etc., and usually above dew point) is channeled to the actual heat-producing electronic equipment components and used to transport heat from that component where it is rejected via a heat exchanger (air to liquid or liquid to liquid) or extended to the cooling terminal device outside of the rack.

Liquid Cooling Overview
– The coolants

- Fluorinerts^{TM} (fluorocarbon liquids)
- Water (or water/glycol mixtures)
- Refrigerants (pump & vapor compression systems)

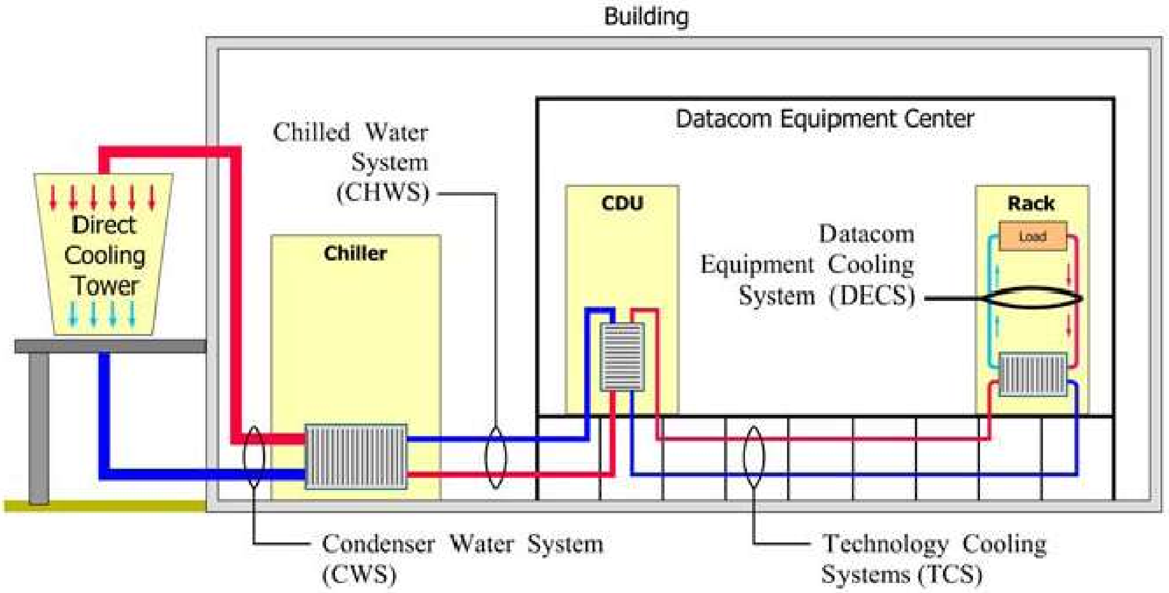

Liquid cooling systems/loops within a data center.

Comparison of Cooling Fluids Based on Cooling Solution Requirements

Cooling Solution Requirement	Refrigerant Technology	Water Based Technology
Capacity to Cool High Heat Densities	★★★ Phase changing of the fluid in the system yields higher capacities in limited space.	★★ One-phase fluid in the system can limit capacity.
Flexibility to Equipment Reconfiguration and Changed Room Layout	★★ Pre-piped room and quick connect couplings can allow flexibility to reconfigure.	★ Pre-piped room and quick connect couplings can allow flexibility to reconfigure. However, reconfiguration cannot be done without introducing water-related risks to the data center.
Energy Efficiency	★★★ Phase changing of the fluid in the circuit yields very good energy efficiency due to smaller pumps and less pressure drop in the heat exchangers located close to the heat source.	★★ Pumping water to the heat exchangers, located close to the heat source, yields good energy efficiency.
Provide Thermal Ride Through in Case of a Failure	★★★ Due to the phase changing of the fluid contained in the piping circuit, thermal ride through time can be achieved.	★★ The water (one-phase fluid) contained in the piping circuit, can yield some thermal ride through time.
Floor Space Efficiency	★★★ Refrigerant technology enables floor space-saving overhead solutions.	★★ With water based technology, non-overhead solutions are typically used because of water related risks.
Low Complexity of Cooling Redundancy	★ Heat exchangers close to the heat source increase complexity of cooling redundancy.	★ Heat exchangers close to the heat source increases complexity of cooling redundancy.
Avoid Possibility for Water Leaks in the Data Center	★★★ No water introduced in the middle of the data center.	★ Requires careful piping layout, piping containment/trays, detection and isolation to minimize the possibility of a water leak.
Possibility to Implement as Retrofit	★ Requires space for distribution piping (and heat exchangers) to implement.	★ Requires space for distribution piping (and heat exchangers) to implement.
Known and Comfortable Technology	★★ Direct expansion refrigerant technology is very well known since many years. Pumped refrigerant technology is known but in a relatively new application when used for data center high heat density cooling.	★★ Water based cooling was more common 20 years ago. The technology is slowly becoming used again because of increasing heat densities.

★ Fair ★★ Good ★★★ Excellent

ASHRAE Transactions 2010, Vol. 116, Part 1, pp. 24-31

Economizer (free cooling)

Economizer mode

- In an economizer mode, the compressor function is fully or partially bypassed, eliminating or reducing its energy use. The compressor is used to move heat from within the data center to the outdoor environment when the outdoor temperature is greater than the data center temperature.
- Economizer modes are sometimes referred to as "free cooling". Most systems using economizer modes spend most of the time in a partial bypass mode, so part of the cooling energy is saved, but the cooling is not "free".

- The operation of data center at partial loads increases the benefit of economizer modes, and more designers recognize that data centers spend a considerable fraction of their life at light load
- The trend toward operation of data centers at higher IT air return temperatures has a dramatic effect on the percent of time economizer mode operation is possible, especially in warmer climates.
- Most new implementations of economizer modes can now operate in a "partial" economizer mode, which greatly increases the amount of energy saved in almost all cases.
- The tools available for quantifying the energy saved by implementing economizer modes are now improved and frequently predict significant savings possibilities with excellent ROI.
- Real-world experience with economizer modes and improvement of controls and monitoring systems, have increased confidence that these modes do not adversely affect the reliability of data centers.

Series / Parallel Operation for Economizer

- In a series configuration, the component that bypasses the compressor (e.g. plate-and-frame heat exchanger) is installed in series with the compressor. This configuration allows for partial economizer mode where the heat exchanger “pre-cools” the air or water. This reduces the total heat energy that the compressor must reject, saving a significant amount of energy.
- In a parallel configuration, the component that bypasses the heat pump is installed in parallel with the heat pump. This configuration prevents the ability to operate in partial economizer mode. This “all or nothing” approach fails to capitalize on the significant energy savings available by operating in partial economizer mode.

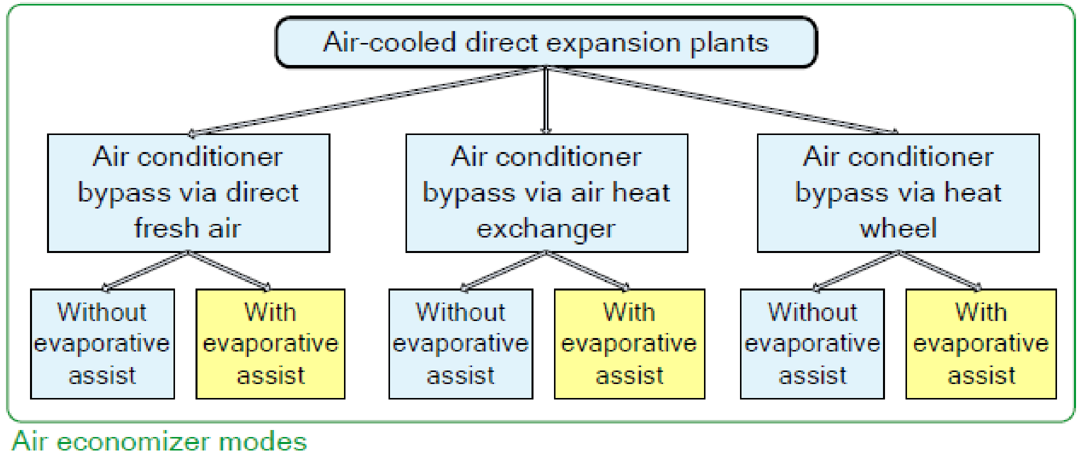

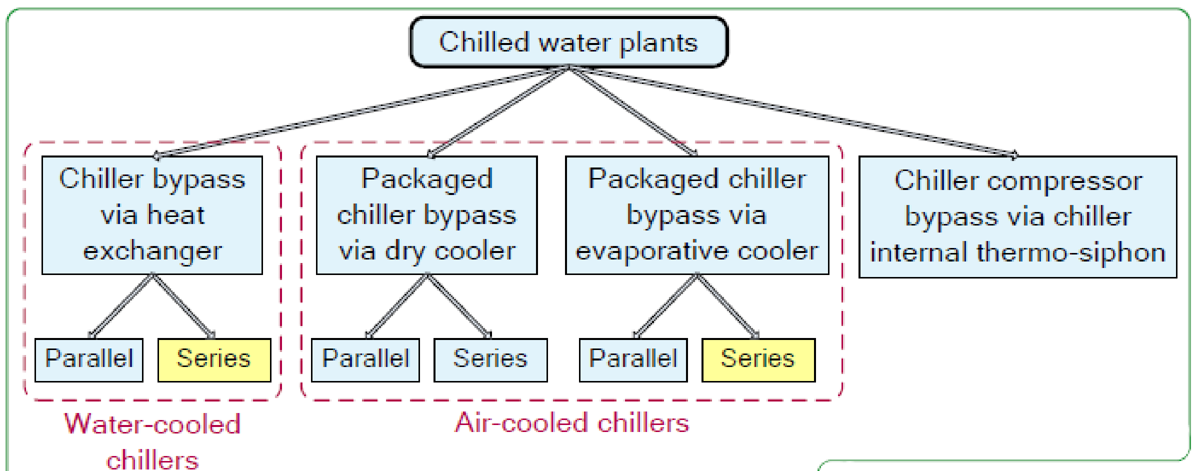

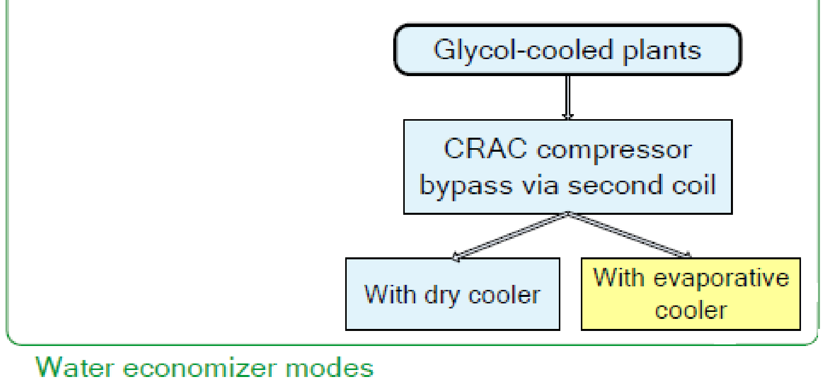

The modes highlighted in yellow are analyzed later in this paper since they allow partial economizer mode operation and evaporative cooling.

Typical Airside economizer

The concept of an airside economizer is to capture outside air with low heat content to replace internal heat gain from occupants, lighting and equipment when outdoor weather condition are favorable.

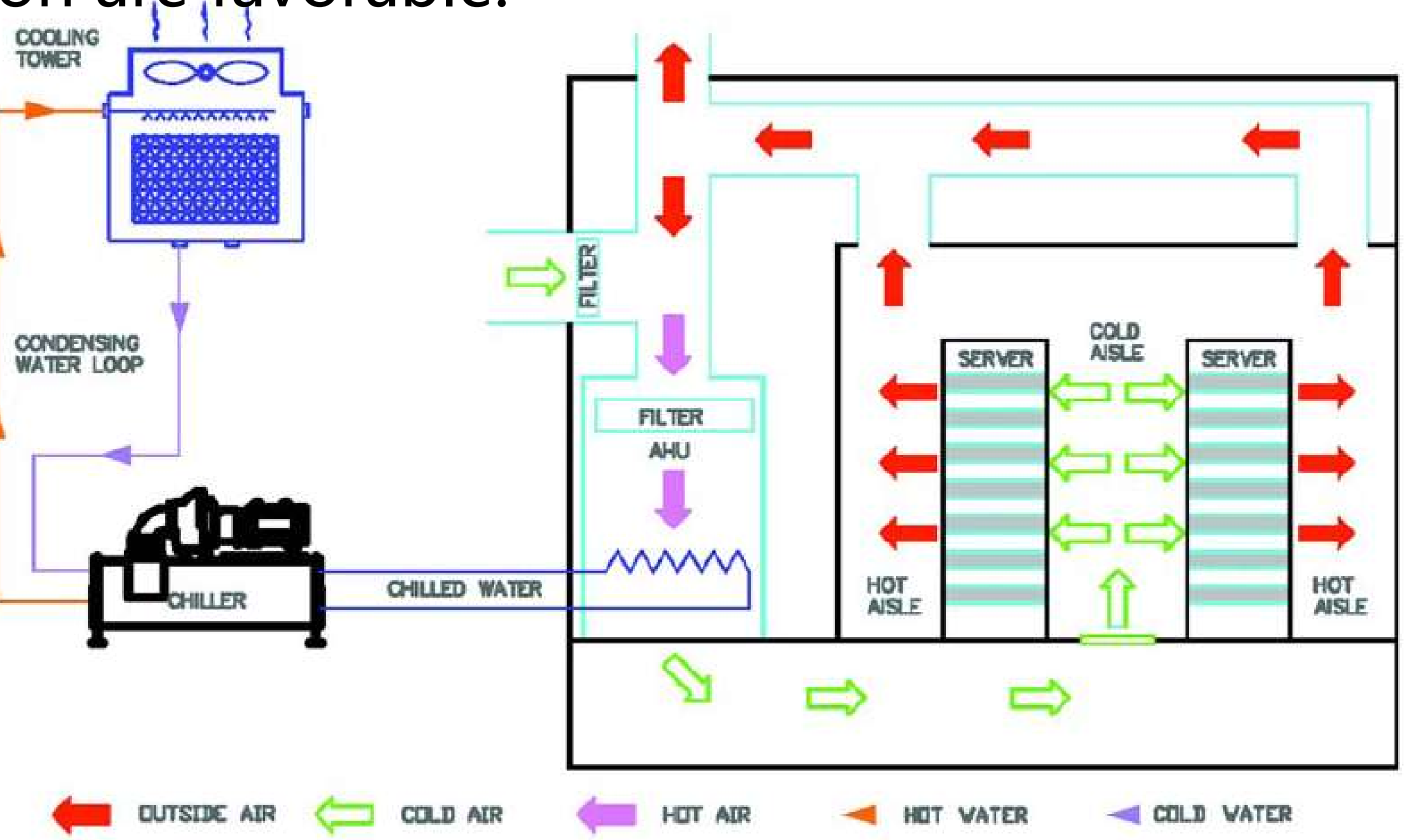

Air-side economizer

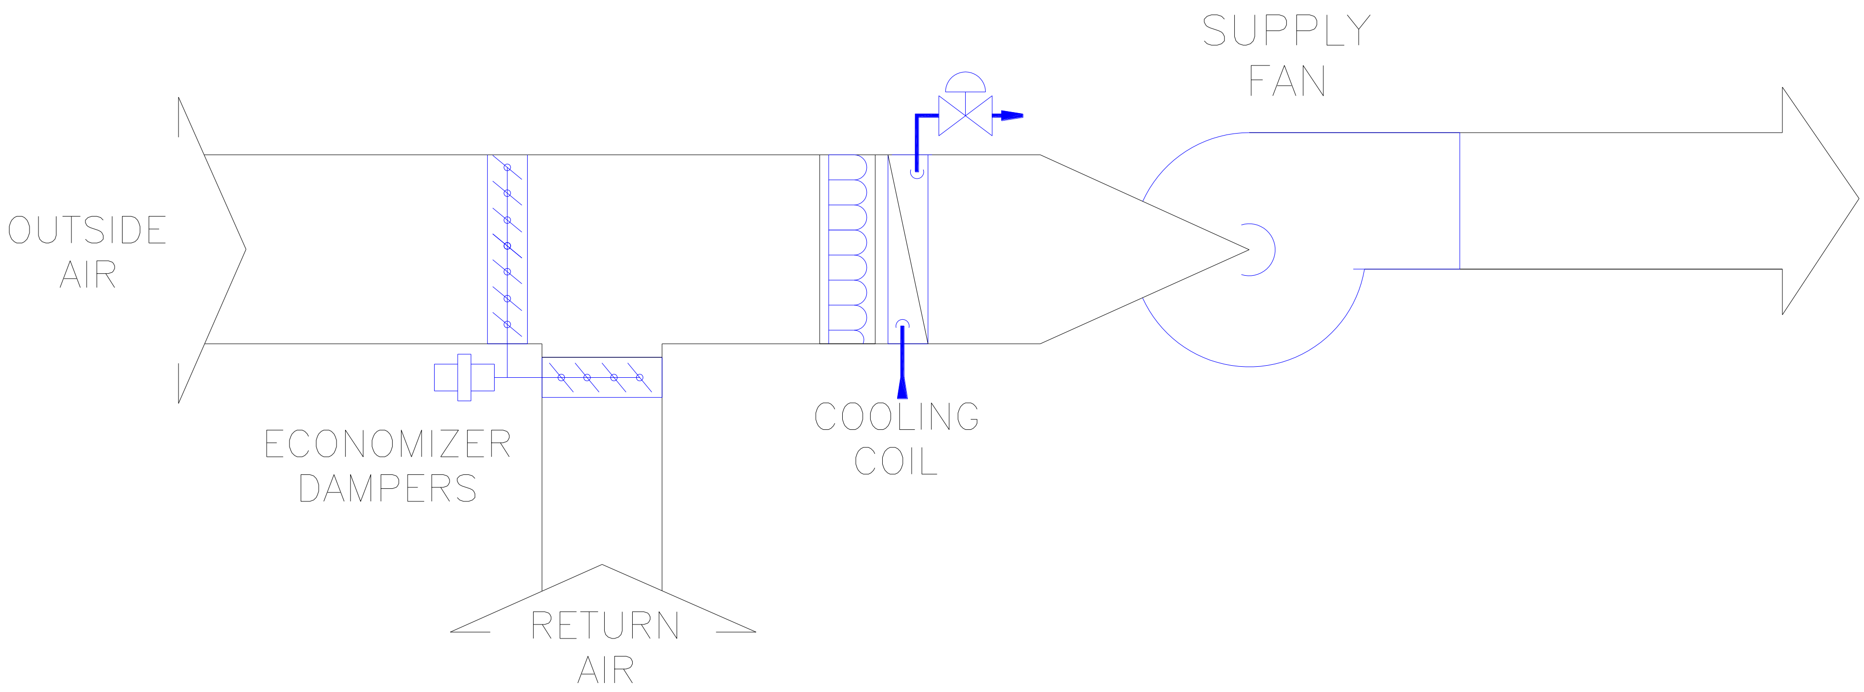

Air-Side Economizer Design conditions

Table 2.1 Class 1, Class 2 and NEBS Design Conditions				
Condition	**Class 1 / Class 2**		**NEBS**	
	Allowable Level	**Recommended Level**	**Allowable Level**	**Recommended Level**
Temperature control range	59°F – 90°F[a,f] (Class 1) 50°F – 95°F[a,f] (Class 2)	68°F – 77°F[a]	41°F – 104°F[c,f]	65°F – 80°F[d]
Maximum temperature rate of change	9°F. per hour[a]		2.9°F/min.[d]	
Relative humidity control range	20% - 80% 63°F. Max Dewpoint[a] (Class 1) 70°F. Max Dewpoint[a] (Class 2)	40% - 55%[a]	5% to 85% 82°F Max Dewpoint[c]	Max 55%[e]
Filtration quality	65%, min. 30%[b] (MERV 11, min. MERV 8)[b]			

[a]These conditions are inlet conditions recommended in the ASHRAE Publication *Thermal Guidelines for Data Processing Environments* (ASHRAE, 2004).
[b]Percentage values per ASHRAE *Standard* 52.1 dust-spot efficiency test. MERV values per ASHRAE Standard 52.2. Refer to Table 8.4 of this publication for the correspondence between MERV, ASHRAE 52.1 & ASHRAE 52.2 Filtration Standards.
[c]Telecordia 2002 GR-63-CORE
[d]Telecordia 2001 GR-3028-CORE
[e]Generally accepted telecom practice. Telecom central offices are not generally humidified, but grounding of personnel is common practice to reduce ESD.
[f]Refer to Figure 2.2 for temperature derating with altitude

Air conditioner bypass via direct fresh air

- A fresh air economizer mode (sometimes referred to as direct air) uses fans and louvers to draw a certain amount of cold outside air through filters and then directly into the data center when the outside air conditions are within specified set points.

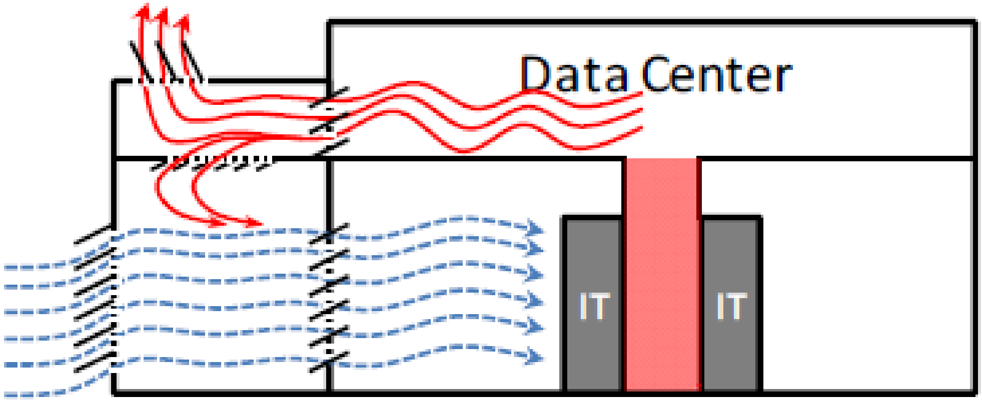

Air conditioner bypass via air heat exchanger

- An air conditioner bypass via air heat exchanger mode (sometimes referred to as indirect air) uses outdoor air to *indirectly* cool data center air when the outside air conditions are within specified set points.

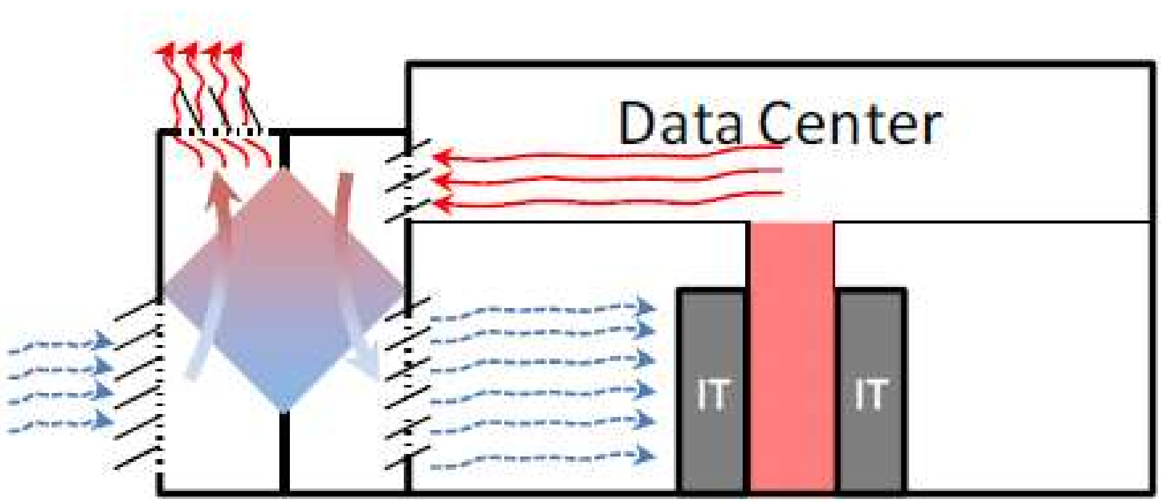

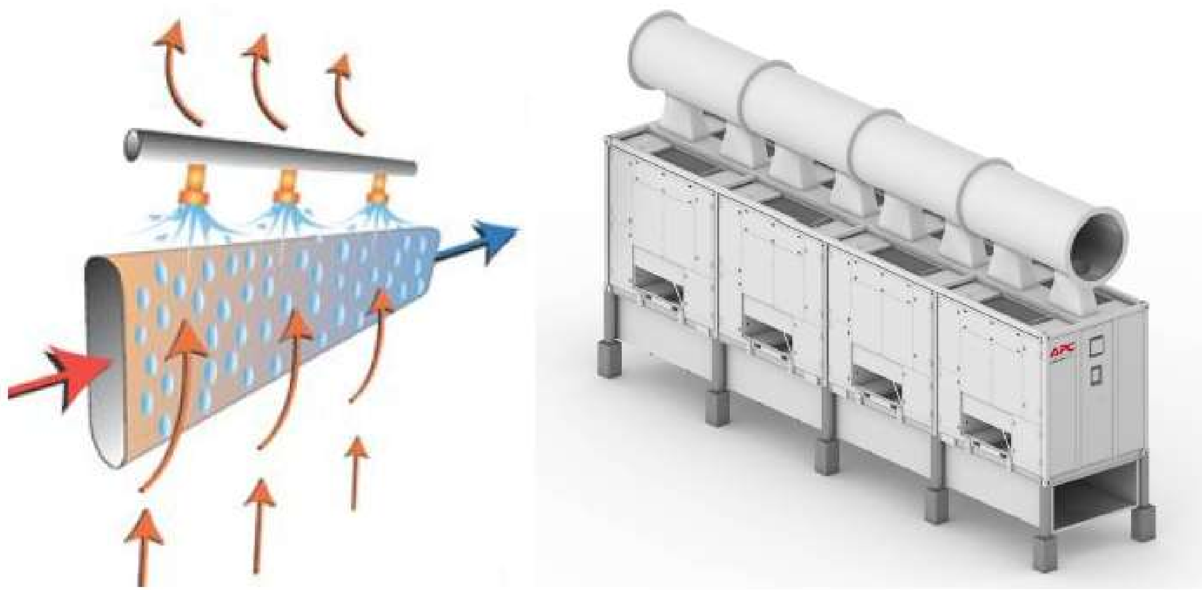

Figure 3b

Illustration of an air-to-air heat exchanger with evaporative assist (left) and an example of a complete cooling system with an integrated air conditioner bypass via air heat exchanger mode (right)

Air conditioner bypass via heat wheel

- An air conditioner bypass via heat wheel mode uses fans to blow the cold outside air through a rotating heat exchanger which preserves the dryer air conditions of the data center space

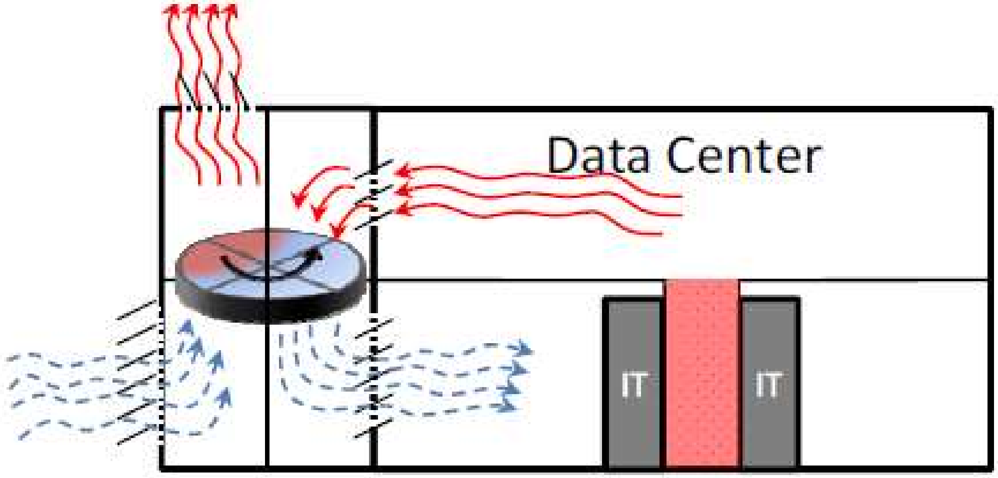

Chiller bypass via heat exchanger

- A chiller bypass via heat exchanger economizer mode uses the condenser water to indirectly cool the data center chilled water when the outside air conditions are within specified set points.

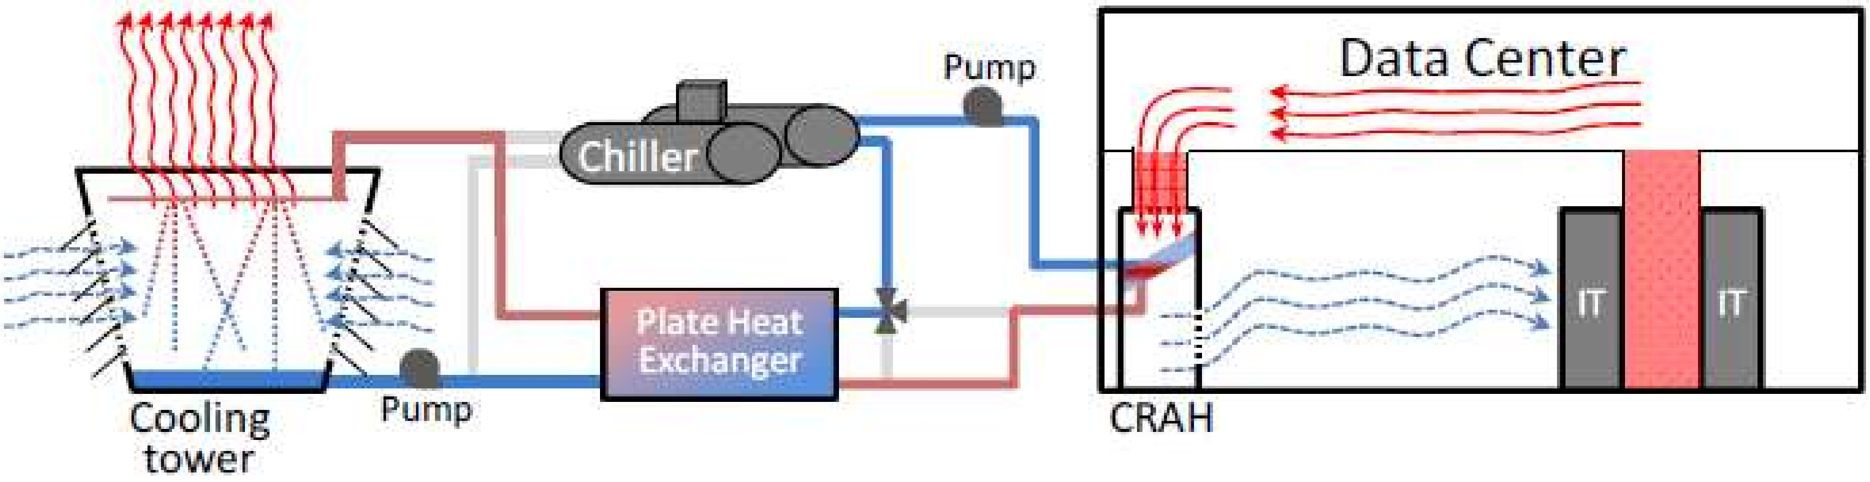

Chiller compressor bypass via chiller internal thermo-siphon

- Some chillers offer a thermo-siphon economizer mode option that allows the compressor to be turned off when the outside air conditions are within specified set points. In this mode, the chiller acts like a simple heat exchanger. The principle of thermo-siphon causes the hot refrigerant to naturally move toward the cold condenser coil where is it cooled. The cold refrigerant then relies on gravity or a pump to travel back to the evaporator coil where it cools the data center chilled water.

Packaged chiller bypass via dry cooler (or via evaporative cooler)

- A packaged chiller bypass via dry cooler economizer mode uses a heat exchanger known as a dry cooler to directly cool the data center chilled water when the outside air conditions are within specified set points.

Figure 6a

Packaged chiller bypass via dry cooler mode

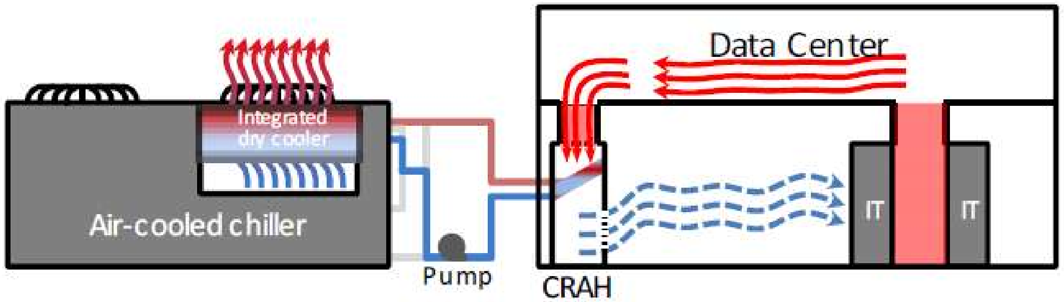

Figure 6b

Example of a packaged chiller with integrated dry cooler

CRAC compressor bypass via second coil

- In this type of economizer mode, the direct expansion (DX) CRAC includes an independent second coil that uses the condenser water during economizer mode operation.

Figure 7

CRAC compressor bypass via second coil mode

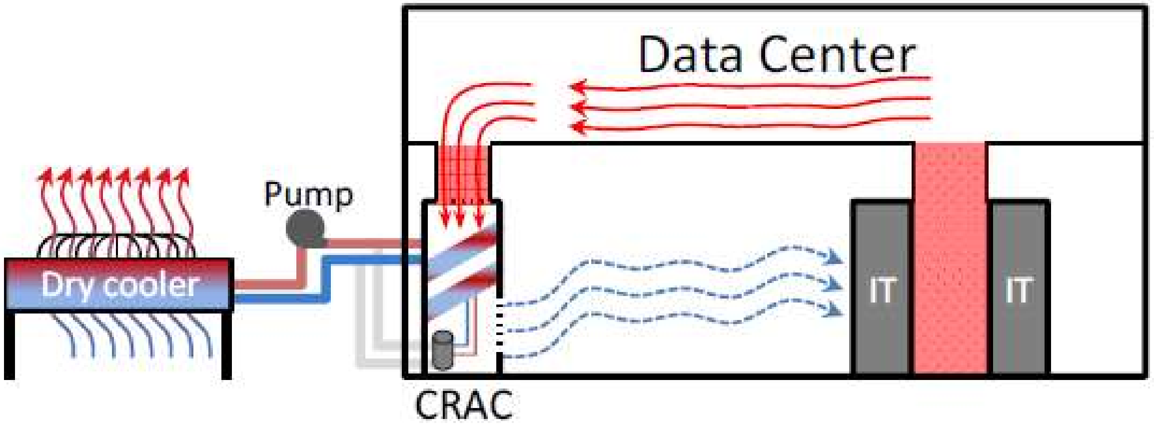

Comparison of the different economizer modes

Table 1

Qualitative comparison between types of economizer modes (blue cells indicate best performer for that attribute)

Economizer mode attribute	Air economizer modes: Air conditioner bypass via direct fresh air (w/ evap assist)	Air economizer modes: Air conditioner bypass via air heat exchanger (w/ evap assist)	Air economizer modes: Air conditioner bypass via air heat wheel (w/ evap assist)	Water economizer modes: Chiller bypass via heat exchanger [3]	Water economizer modes: Packaged chiller bypass via evaporative cooler [3]	Water economizer modes: CRAC compressor bypass via second coil (w/ evap assist)
Building shell compatibility	May require building shell modification	May require building shell modification	May require building shell modification	No issue with building shell	No issue with building shell	No issue with building shell
Ability to retrofit	Not logical to retrofit into existing system	Not logical to retrofit into existing system	Not logical to retrofit into existing system	Practical if space available	Practical if space available	Requires swapping out CRAC unit
Complexity of controls	Fewer devices to control	Fewer devices to control	Fewer devices to control	Most devices to control	Moderate number devices to control	Moderate number devices to control
Data center humidity control	Dependent on outdoor humidity	Independent of outdoor humidity	Independent of outdoor humidity	Independent of outdoor humidity	Independent of outdoor humidity	Independent of outdoor humidity
Life expectancy	20-40 years on heat exchanger	20-40 years on heat exchanger	20-40 years on heat exchanger	10-15 yrs on plate heat exchanger	10-20 years on evaporative cooler	10-20 years on cooling unit
Availability risks -loss of cooling water - poor air quality - fire suppression	Highly susceptible to outdoor air quality Shutdown required with clean agent suppression	Low downtime risk due to water loss. No risk due to poor air quality, or fire suppression	Low downtime risk due to water loss. No risk due to poor air quality, or fire suppression	Downtime due to loss of make-up water for cooling tower	No downtime due to water loss, poor air quality, or fire suppression	No downtime due to water loss, poor air quality, or fire suppression
Footprint	0.41 ft^2 / kW 0.038 m^2 / kW	0.788 ft^2 / kW 0.073 m^2 / kW	1.72 ft^2 / kW 0.16 m^2 / kW	1.94 ft^2 / kW 0.18 m^2 / kW	3.34 ft^2 / kW 0.31 m^2 / kW	2.02 ft^2 / kW 0.19 m^2 / kW
Need for backup refrigerant mode	Fully sized backup in case of poor outdoor air quality	Partially sized for extreme climates	Partially sized for extreme climates	Partially sized for extreme climates	Partially sized for extreme climates	Partially sized for extreme climates

Table 2

Quantitative comparison between types of economizer modes

Economizer mode attribute	Air economizer modes: Air conditioner bypass via direct fresh air (w/ evap assist)	Air conditioner bypass via air heat exchanger (w/ evap assist)	Air conditioner bypass via air heat wheel (w/ evap assist)	Water economizer modes: Chiller bypass via heat exchanger [5]	Packaged chiller bypass via evaporative cooler [5]	CRAC compressor bypass via second coil (w/ evap assist)
The following attributes are based on a 1MW data center at 50% IT load, located in St. Louis, MO, U.S. See side bar for all assumptions.						
Annual water consumption	100 gal 379 L	1,262,000 gal 4,777,000 L	257,000 gal 973,000 L	7,000,000 gal 26,000,000 L [6]	128,000 gal 485,000 L	128,000 gal 485,000 L
Capital cost of entire cooling system	$2.2 / watt	$2.4 / watt	$2.8 / watt	$3.0 / watt	$2.3/ watt	$2.0 / watt
Annual maintenance cost of entire system [7]	75%	75%	83%	100%	100%	92%
Total cooling energy	737,506	340,365	377,625	589,221	736,954	960,974
Annual hours - full economizer mode	5,723	7,074	5,990	4,705	5,301	4,918
Annual hours - partial economizer mode	0	1,686	2,770	3,604	1,773	3,800
Est. annual PUE	1.34	1.25	1.26	1.31	1.34	1.39

Economics of evaporative assist

- The cost of evaporative coolers and evaporative assist in general include the material cost, water usage, and water treatment.
- These costs must be considered when deciding upon a data center cooling system.
- Evaporative assist is most effective in dry climates such as Las Vegas and Dubai. The cost of an evaporative cooler must be balanced against its effectiveness in climates that are more humid.
- It is possible to spend more on evaporative cooling than is saved on cooling system energy.

Water side economizer

The basic principle of waterside economizer is to pre-cool some or all of the return water in a chilled water loop with the cooling tower, substantially reducing or even eliminating the need of mechanical cooling. Through the use of plate and frame heat exchangers, building heat is transferred from the chilled water loop into the cooling tower loop and eventually dissipated into the atmosphere.

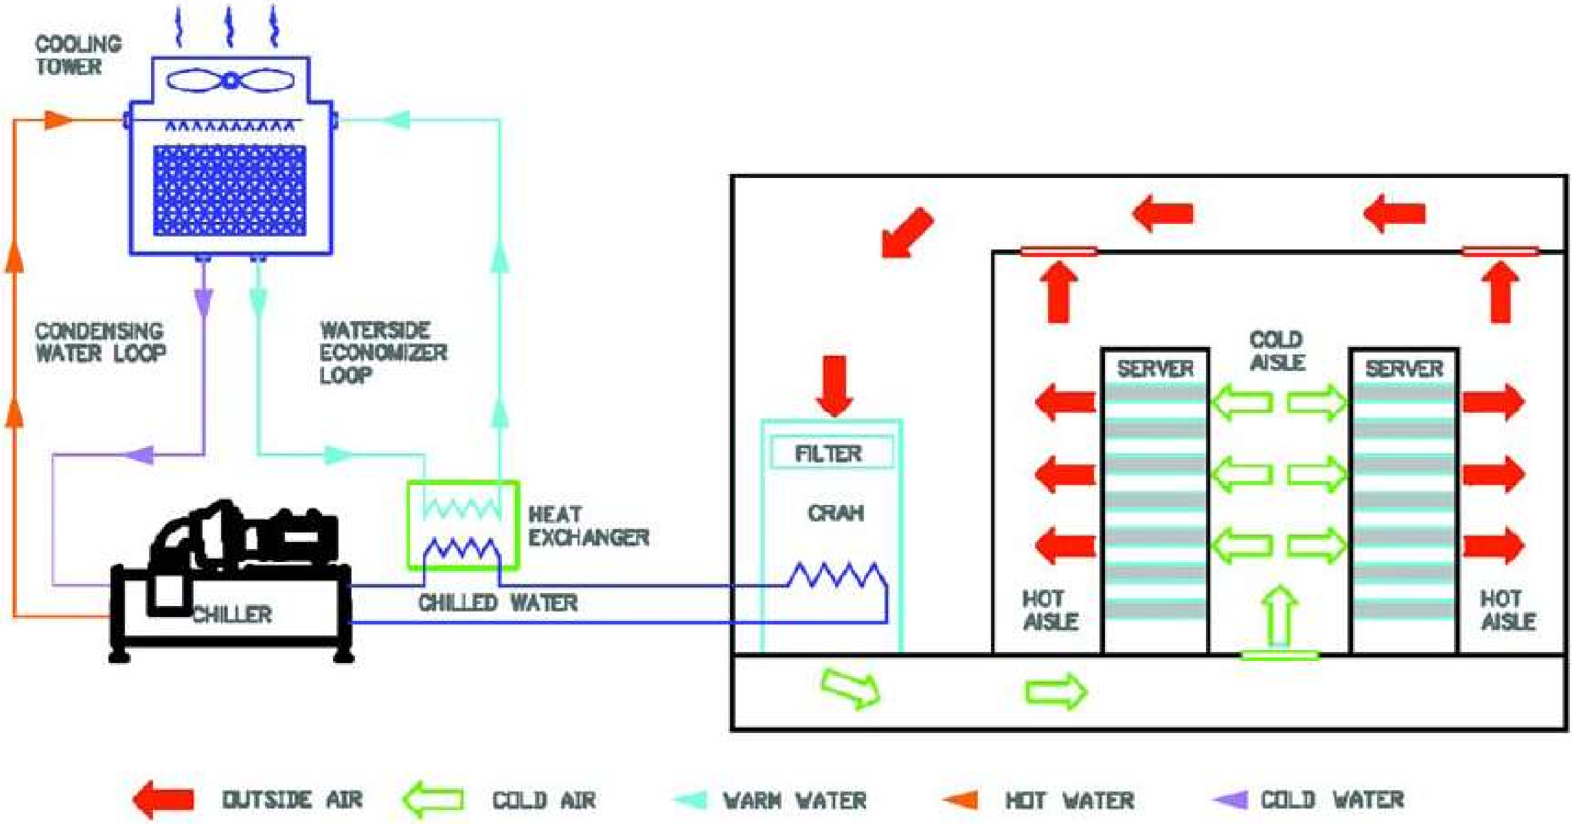

Water-Side Economizer

Integrated

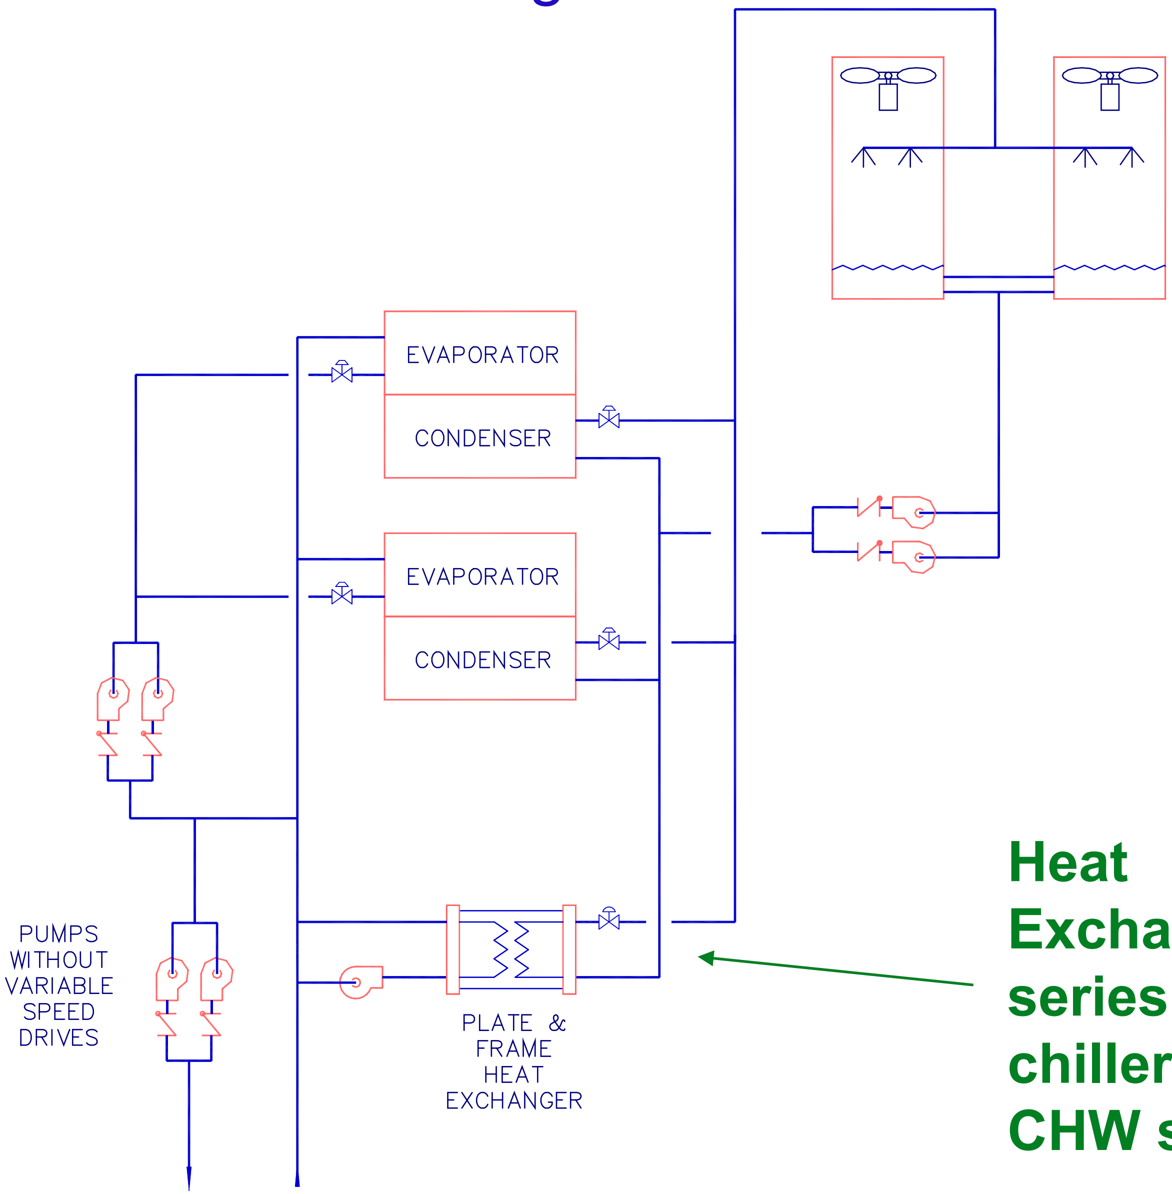

Heat Exchanger in series with chillers on CHW side

Typical parallel waterside economizer (PWSE) flow diagram in normal operation (left) and economizer operation (right).

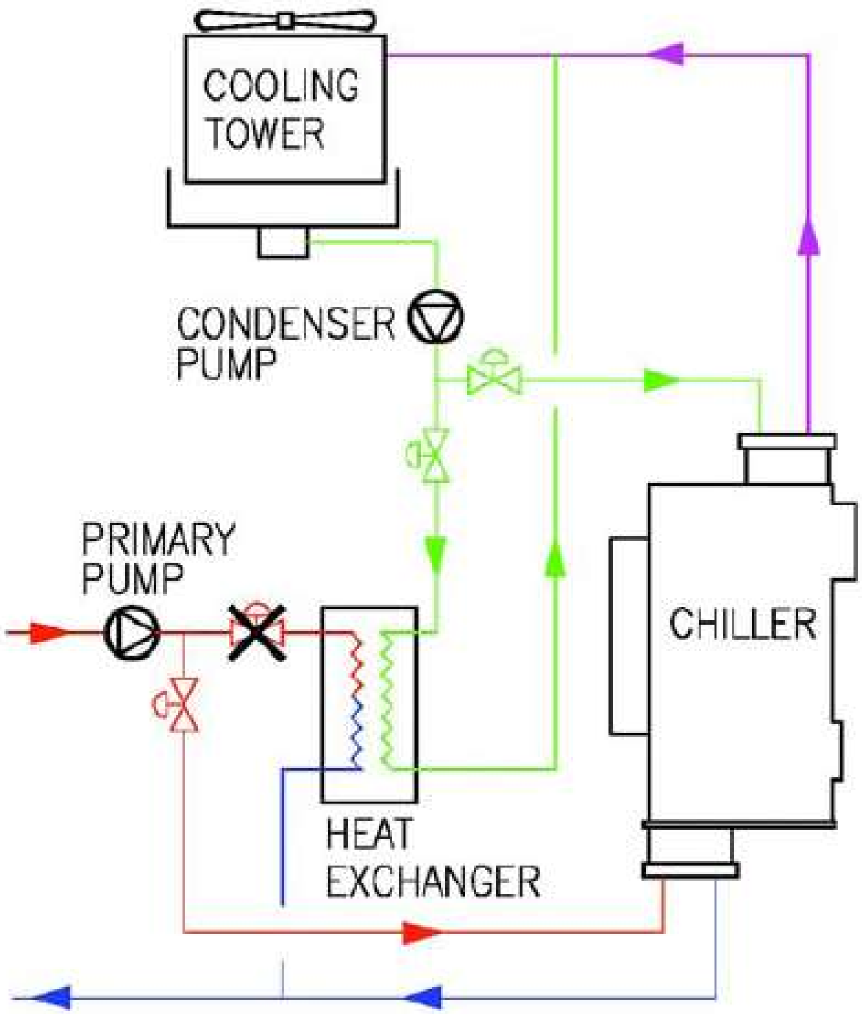

Typical series waterside economizer (SWSE) flow diagram in normal operation (left) and economizer operation (right)

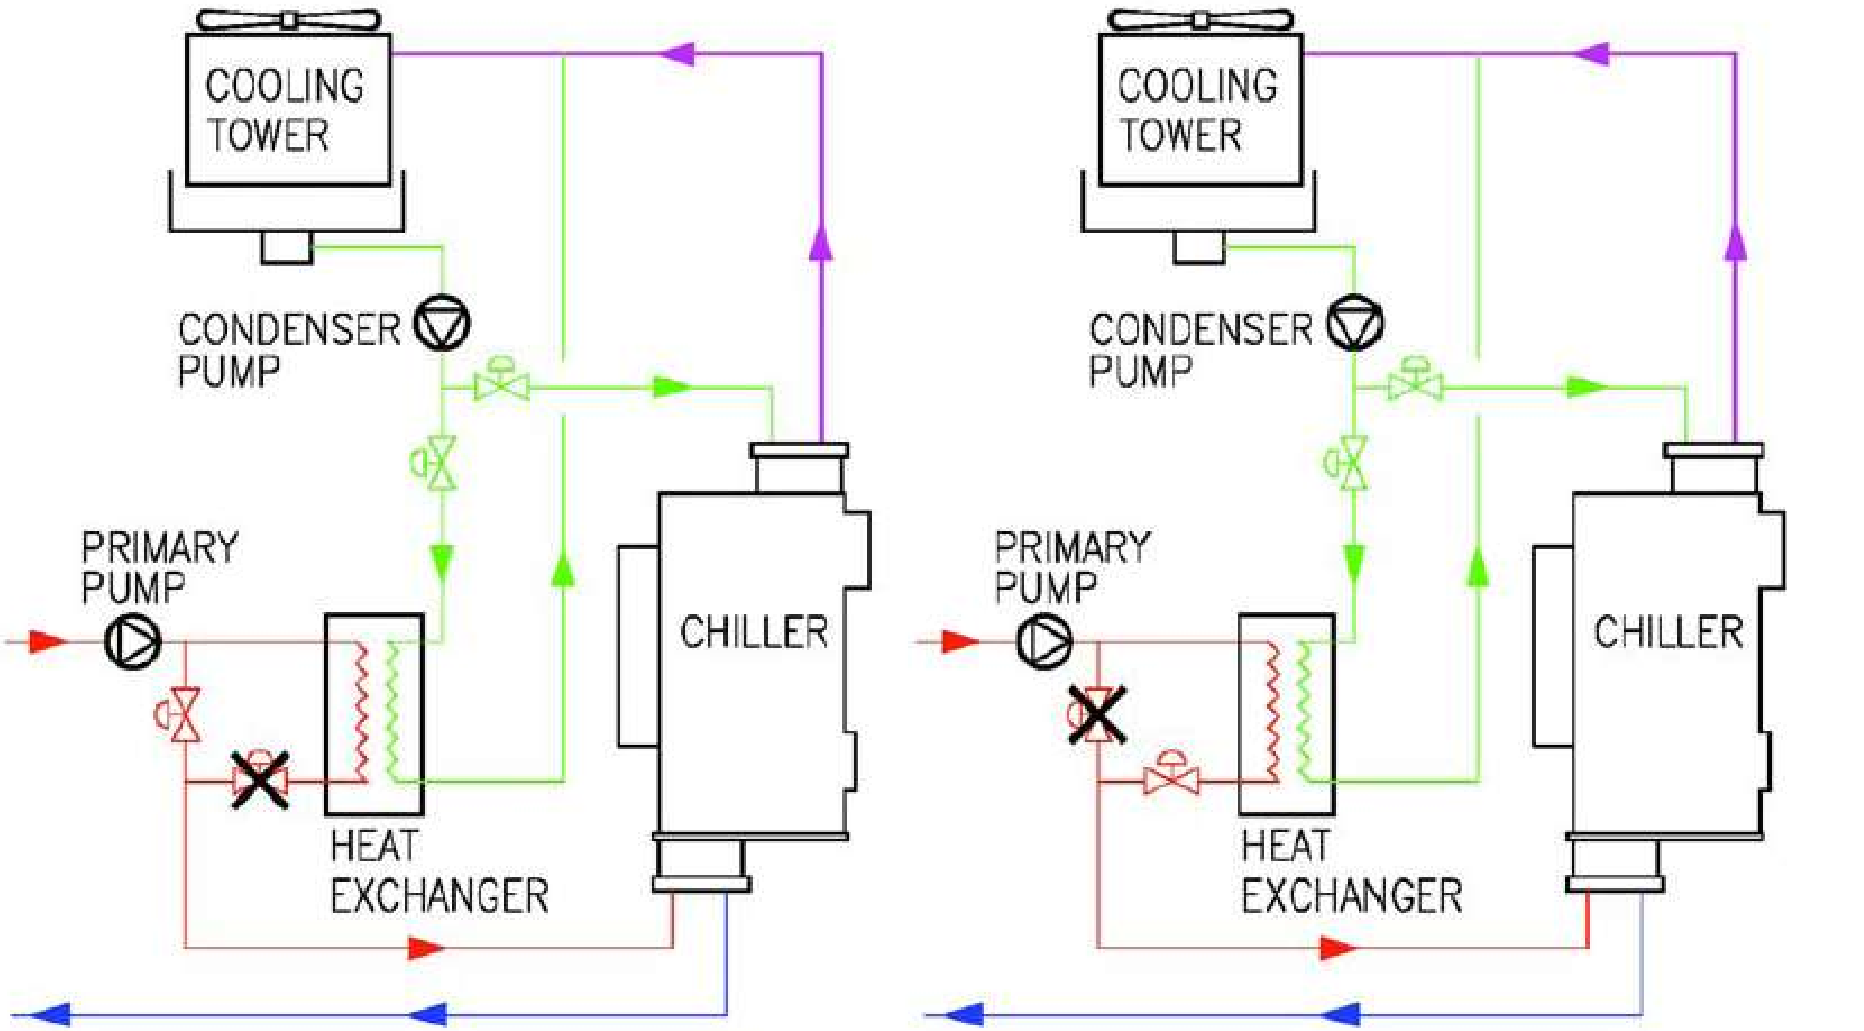

Economizer Summary

Air-Side Economizers

- Provides free cooling when dry-bulb temperatures are below 78°F-80°F.
- May increase particulates (LBNL research indicates this is of little concern).
- Should be integrated to be most effective.
- Improves plant redundancy!
- Can work in conjunction with water-side economizers on data centers!
- Need to incorporate relief.

Water-Side Economizers

- Provides low energy cooling when wet-bulb temperatures are below 55°F-60°F.
- Avoids increased particulates (and low humidity if that concerns you).
- Should be integrated to be most effective (see previous slide).
- Improves plant redundancy!
- Can work in conjunction with air-side economizers on data centers!

Both are proven technologies on data centers!

Study on Free Cooling Systems for Data Centers in Japan

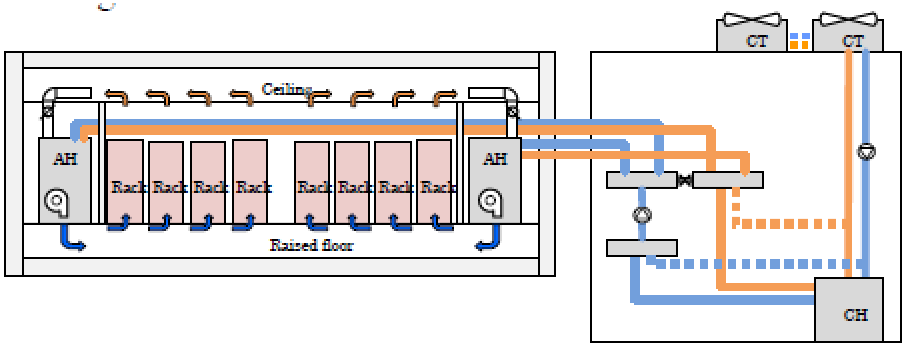

Figure 7: Indirect free cooling system
(water-side economizer)

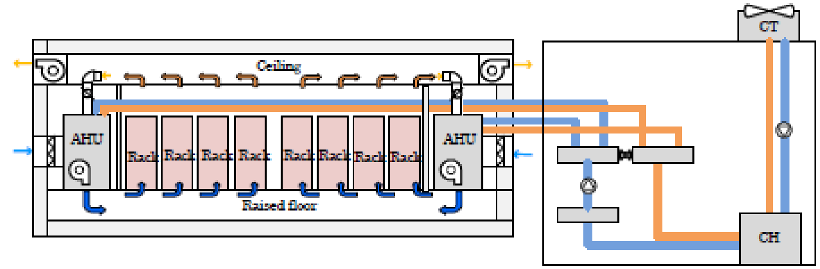

Figure 5: Direct free cooling system
(air-side economizer)

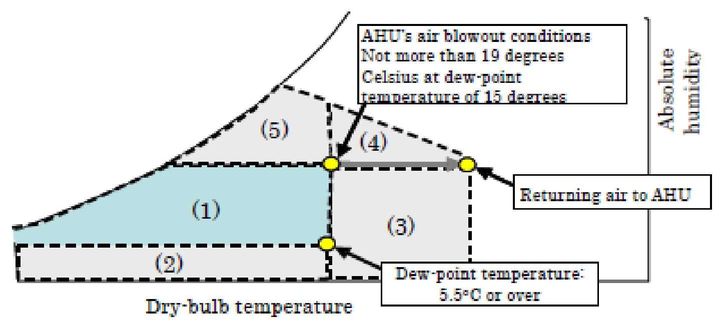

(1) Non-humidifying and free cooling zone
(2) Humidifying and free cooling zone
(3) Heat source load reducing zone
(4) Dehumidifying and heat source load reducing zone
(5) Dehumidifying and free cooling zone

Figure 6: Scope of the air-side economizer

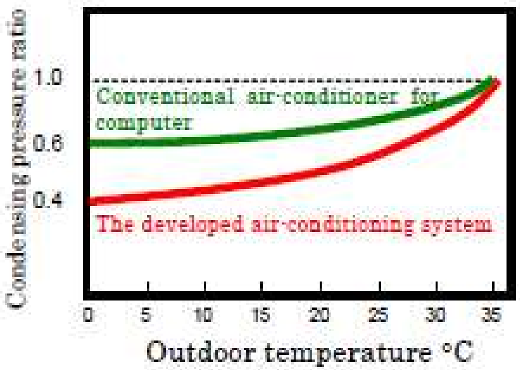

Figure 8: Relation between condensing pressure and outdoor temperature

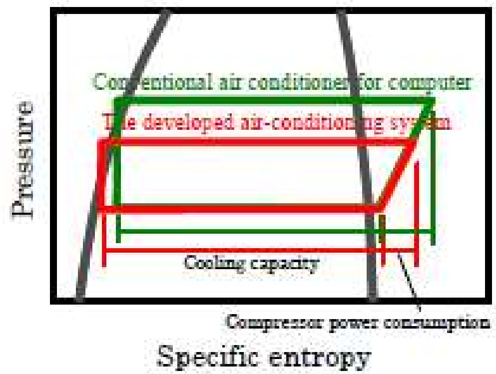

Figure 9: Mollier diagram

Study on Free Cooling Systems for Data Centers in Japan

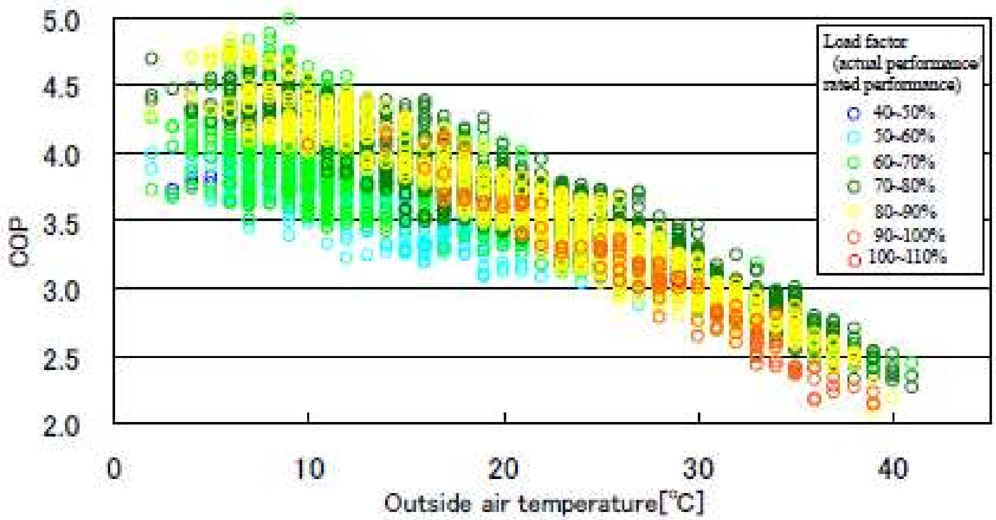

Figure 10: Outside air temperature vs. air conditioner's coefficient of performance (field data)

Energy Efficient Free Cooling System for Data Centers

2011 Third IEEE International Conference on Cloud Computing Technology and Science

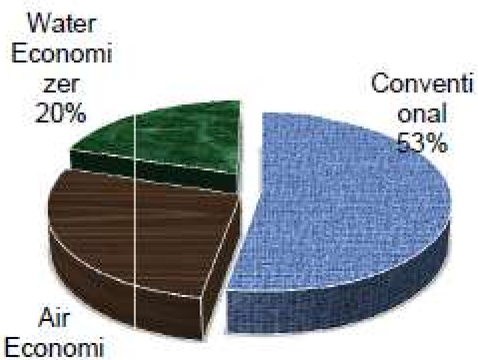

Figure 5: Annual Energy Cost - CHICAGO

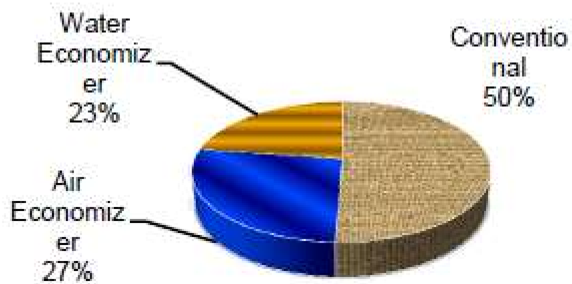

Figure 7 :Annual Energy Cost - ATLANDA

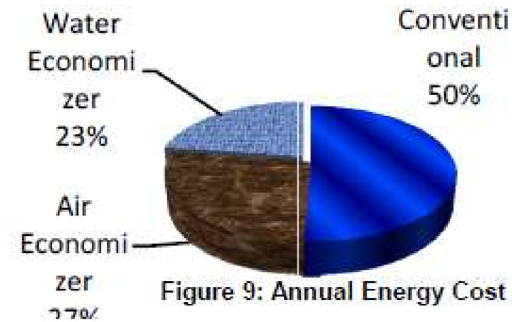

Figure 9: Annual Energy Cost -PHOENIX

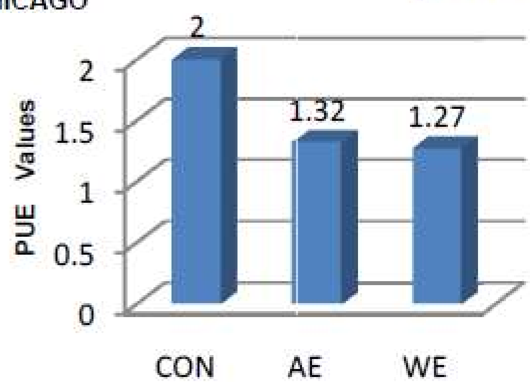

Figure 10: PUE Values for the various scenarios

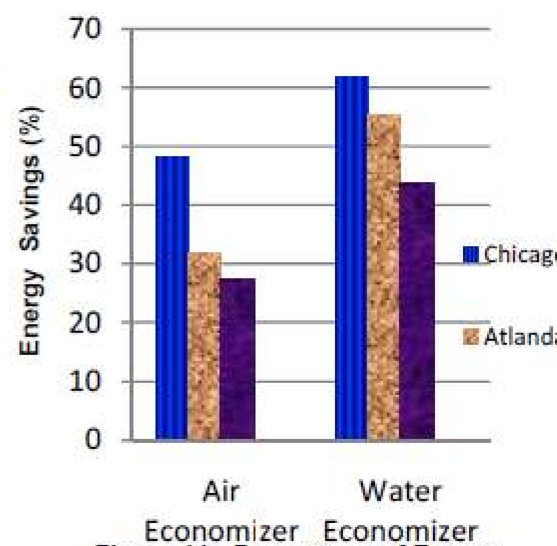

Figure 11: Percentage of Energy Savings in all the three Locations

Liquid Cooling

Liquid Cooling

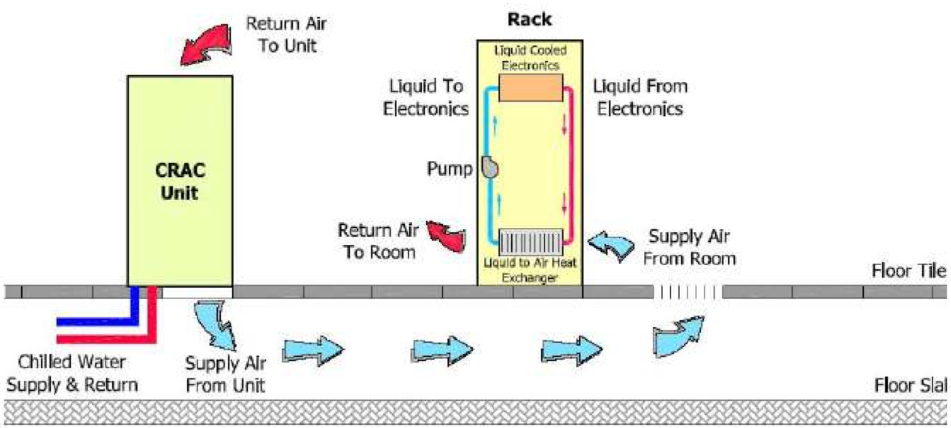

Internal liquid cooling loop restricted within rack extent.

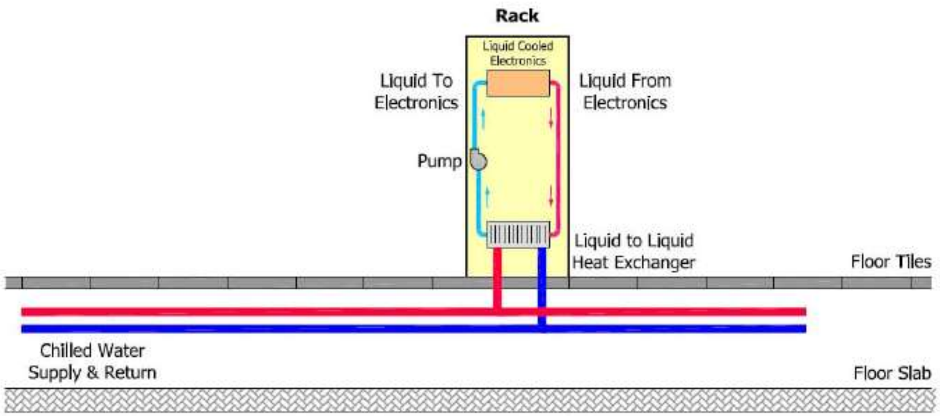

Internal liquid cooling loop with rack extents and liquid cooling loop external to racks.

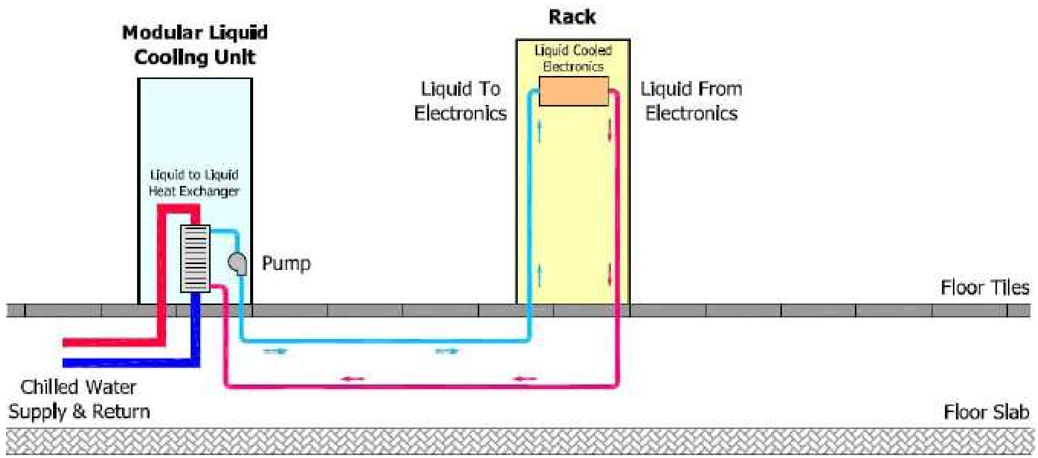

Internal liquid cooling loop extended to liquid-cooled external modular cooling unit.

Basic Piping Architecture

- Direct return
  - The most basic type featuring reduced number of connection pt.

Advantages

1. Least expensive to construct, uses a minimal amount of pipe, valves, and fittings.
2. Simplest to operate and understand.

Disadvantages

1. Least reliable since only one source of cooling exists.
2. No redundancy in piping to the load. Any pipe failure or leak or future addition could jeopardize system availability.
3. May require additional balancing valves.

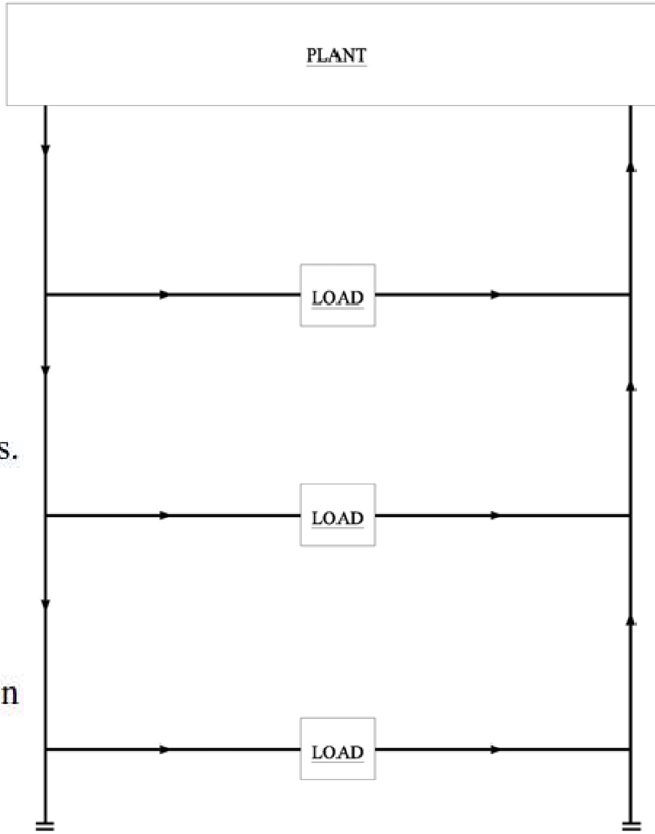

Example of direct return flow principle.

- Reverse return
  - Create a network with self-balancing

Advantages

1. Simple to operate and understand.
2. Self-balancing.

Disadvantages

1. Less reliable; again, only one source of cooling.
2. No redundancy in pipe or chilled-water routes. Routine maintenance or system expansion could require complete system shutdown.
3. A little more expensive to install than direct return (i.e., more piping required).

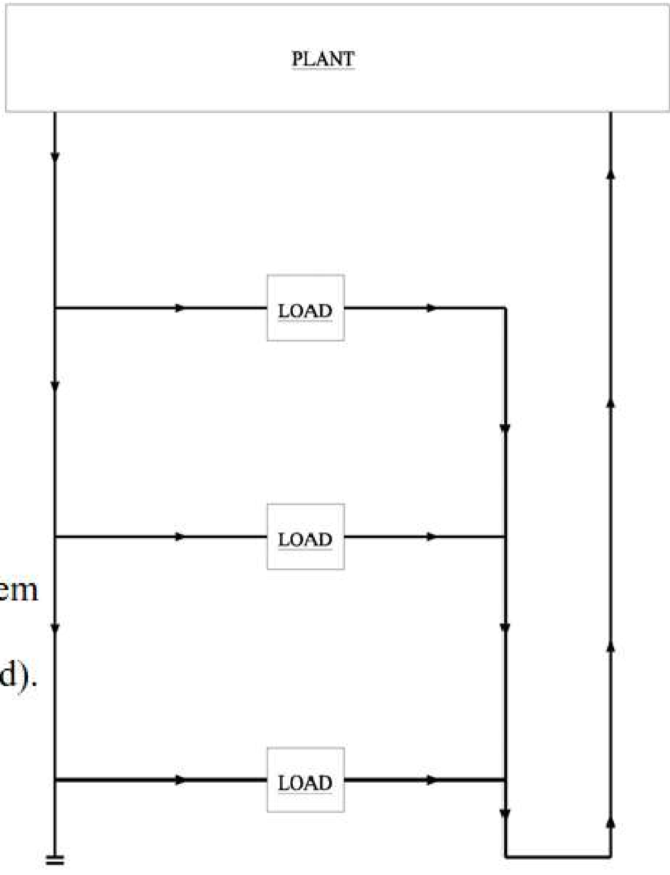

Example of reverse return flow principle.

Single ended loop with direct feed

- Loop piping main allow a section in the main to be maintained and repaired

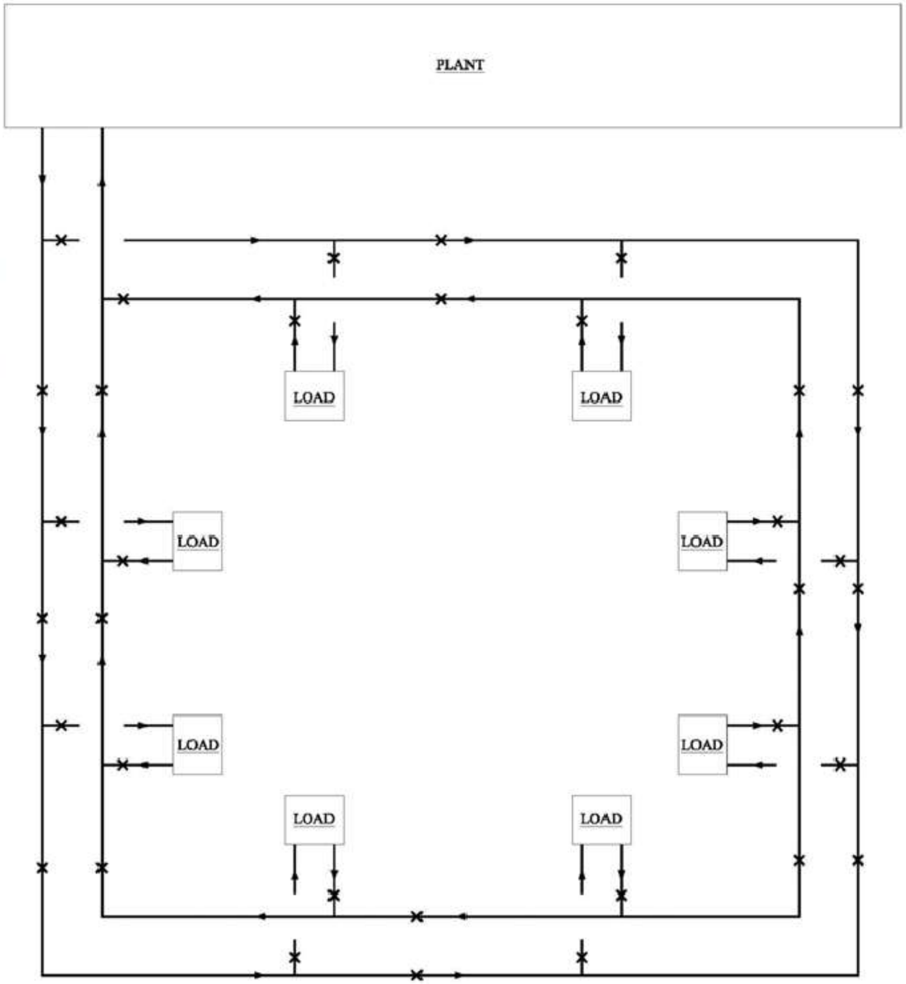

Advantages

1. Self-balancing.
2. Increased reliability over direct and reverse returns systems with two piping routes to the load.
3. Individual pipe sections and future equipment installations are serviceable without system shutdown.

Disadvantages

1. Increased complexity and understanding.
2. Increased installation costs.

Single ended with common cross branch

- Allow bi-direction flow in the main and cross branch

Advantages

1. Increased reliability with multiple piping routes to load.
2. Self-balancing.
3. Individual pipe sections and future equipment installations are serviceable without system shutdown.

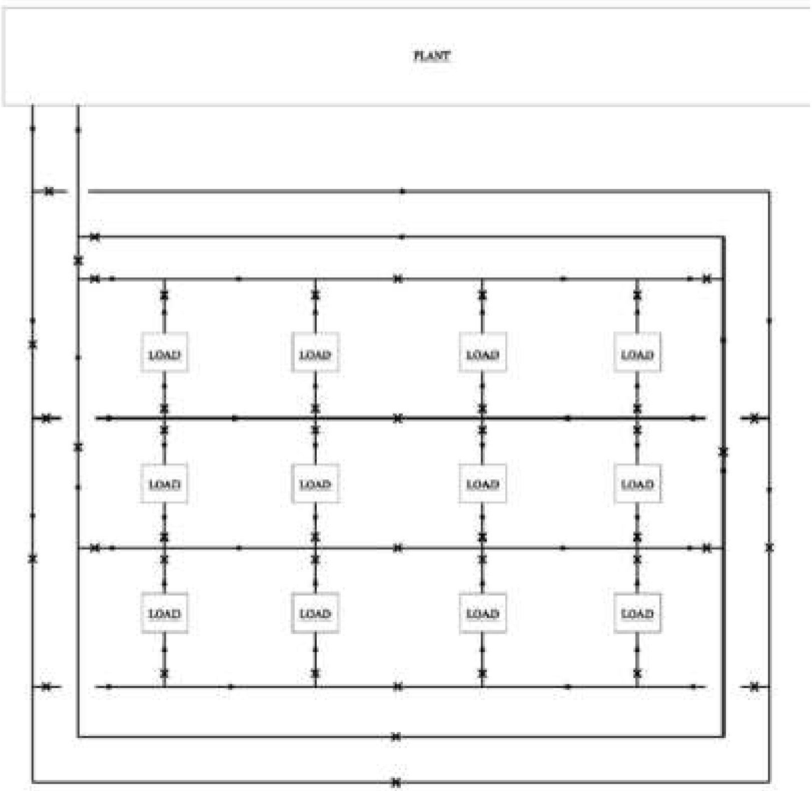

Disadvantages

1. Increased installation costs.
2. Increased operational complexity.

Double ended with direct feed

- Two connections to the plant, avoid the single point failure

Advantages

1. High reliability.
2. Redundant piping routes to load and a second cooling supply and return mains from the plant.
3. Redundant cooling supply and return piping from a second central plant.
4. Individual pipe sections and future equipment installations are serviceable without system shutdown.
5. Self-balancing.

Disadvantages

1. Increased installation costs.
2. Increased operational complexity.

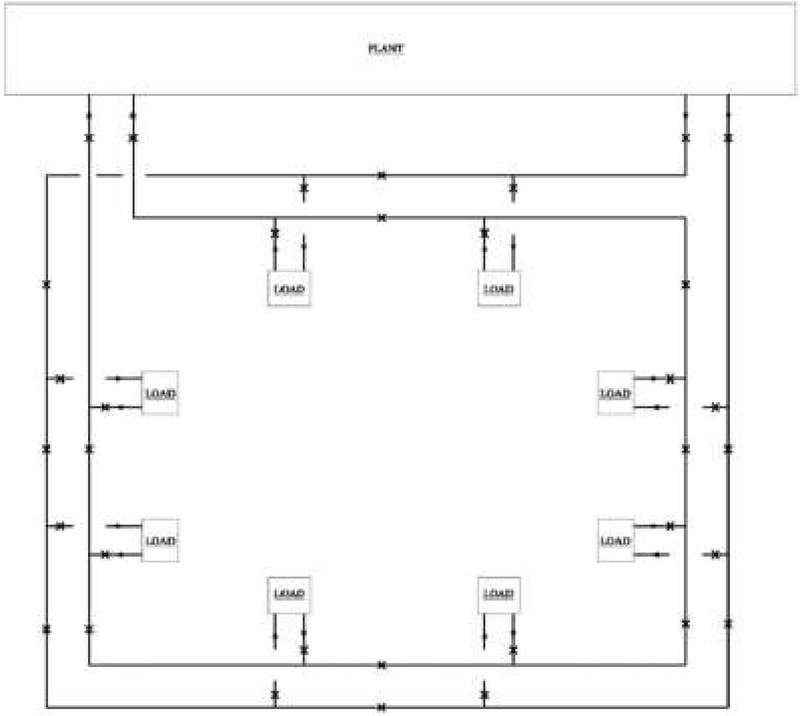

Double ended with cross branch

Advantages

1. High reliability.
2. Redundant piping routes to load and a second cooling supply and return mains from the plant.
3. Redundant cooling supply and return piping from a second central plant.
4. Individual pipe sections and future equipment installations are serviceable without system shutdown.
5. Self-balancing.

Disadvantages

1. Increased installation costs.
2. Increased operational complexity.

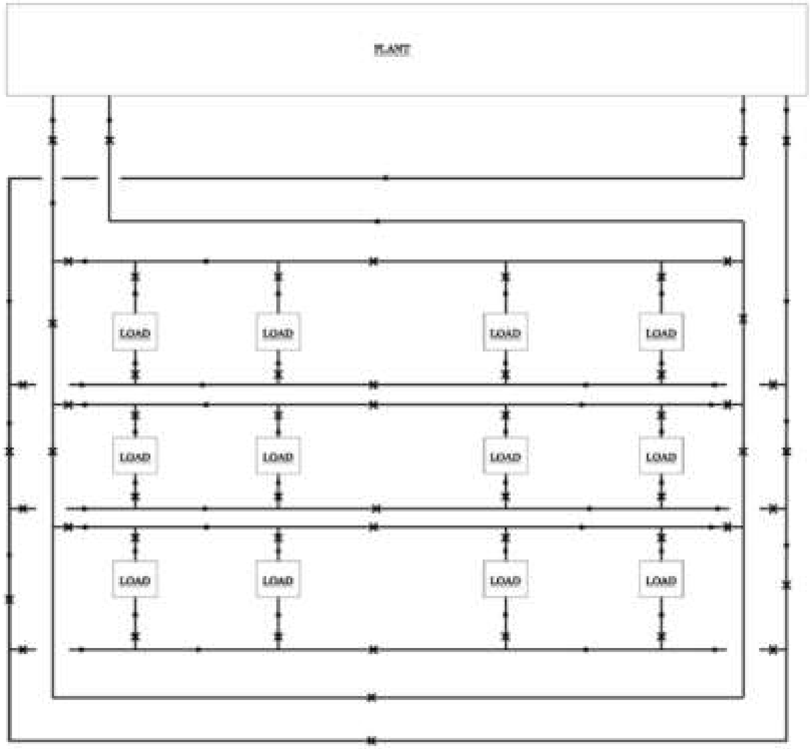

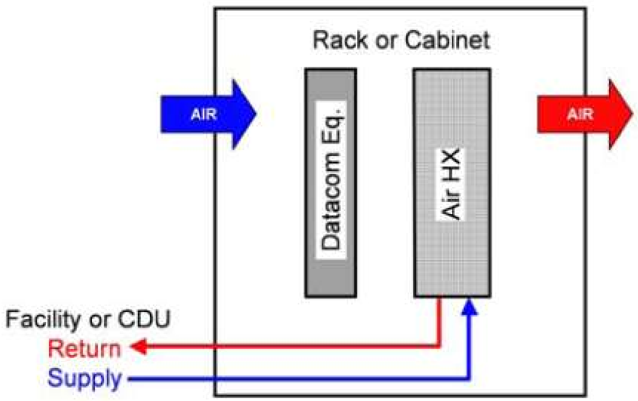

Open air-cooled datacom equipment in an air/liquid-cooled rack.

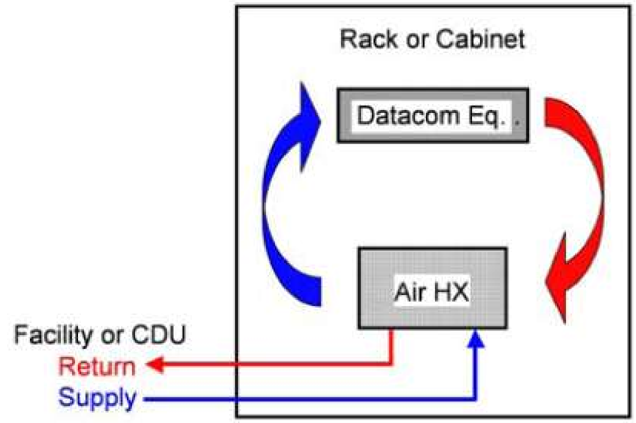

Closed air-cooled datacom equipment in a liquid-cooled cabinet.

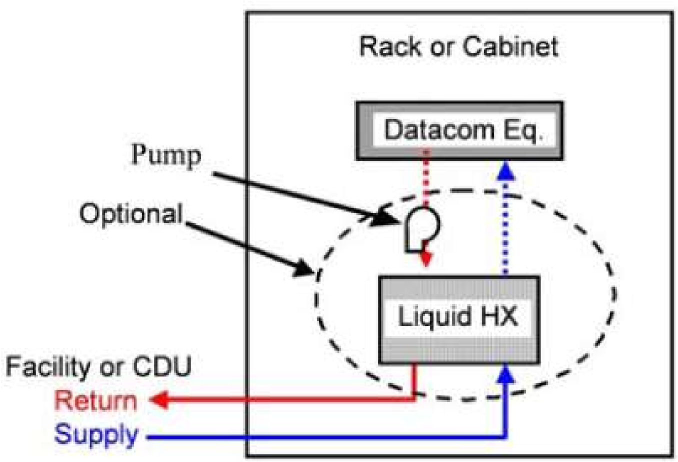

Liquid-cooled datacom equipment in a liquid-cooled rack.

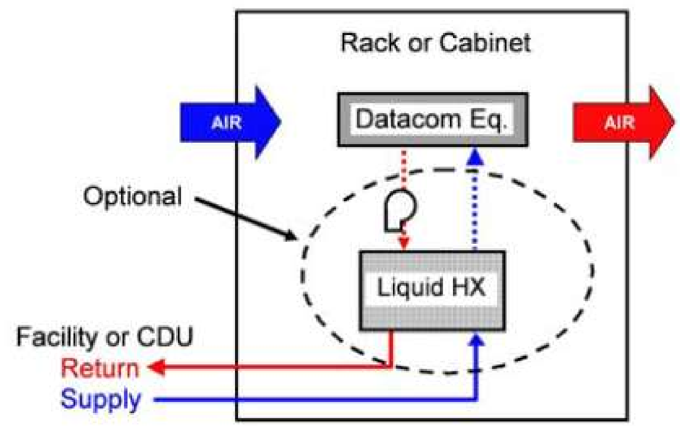

Open air- and liquid-cooled datacom equipment in an air/liquid-cooled rack.

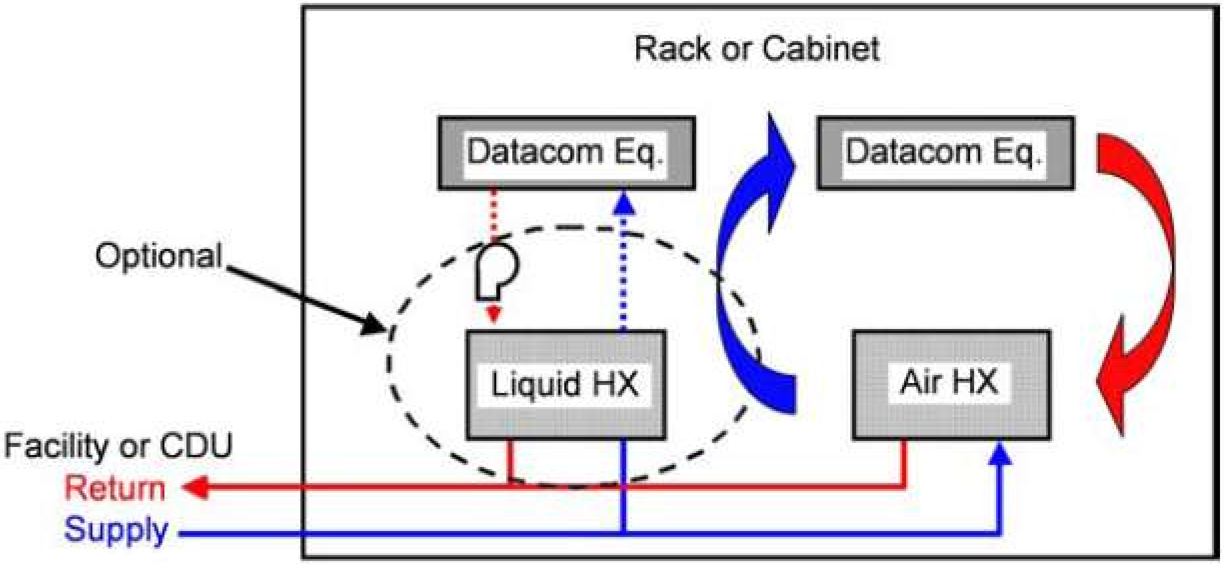

Closed air- and liquid-cooled datacom equipment in a liquid-cooled rack.

Overview of coolant distribution unit (CDU)

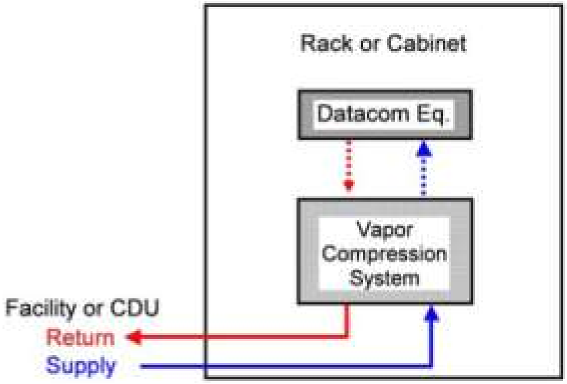

Liquid-cooled datacom equipment in a liquid-cooled rack using a vapor-compression system.

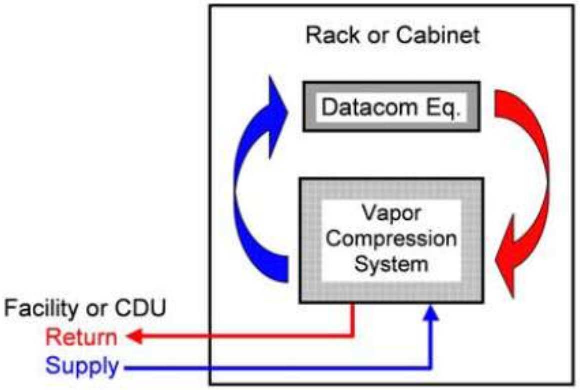

Air-cooled datacom equipment in a liquid-cooled cabinet using a vapor-compression cycle.

The CDU provides a number of important benefits, as follows:

- **Prevention of condensation:** It provides an opportunity for temperature conditioning, which could allow the coolant to be delivered to the electronics above the dew point.
- **Isolation:** It may allow the electronics to be isolated from the harsher facility water in the CHWS or CWS loop. The loop supplied by the CDU also utilizes a lower volume of coolant, so a coolant leak is less catastrophic.
- **Coolant flexibility:** The separate loop associated with the CDU allows users the flexibility to use any number of coolants.
- **Temperature control:** The separate loop associated with the CDU allows users the flexibility of running the electronics at a desired temperature. This temperature can be above, at, or below the dew point.

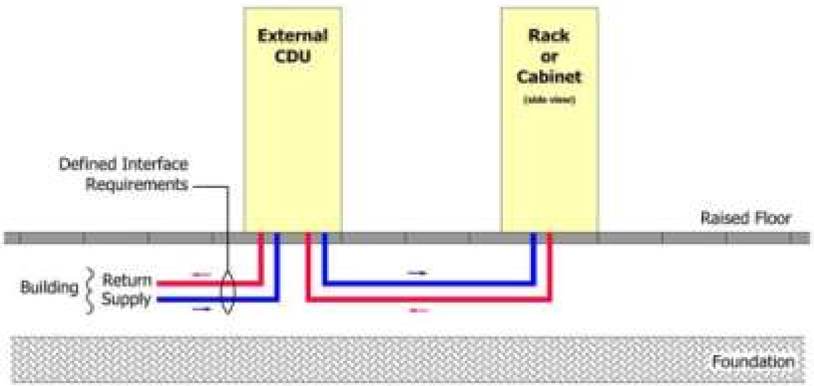
External CDU
Rack or Cabinet
Defined Interface Requirements
Raised Floor
Building
Return
Supply
Foundation

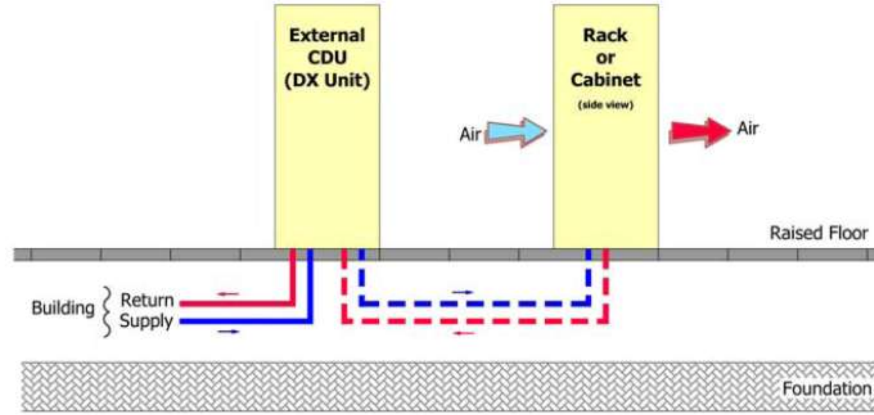
External CDU (DX Unit)
Rack or Cabinet
(side view)
Air
Air
Raised Floor
Building
Return
Supply
Foundation

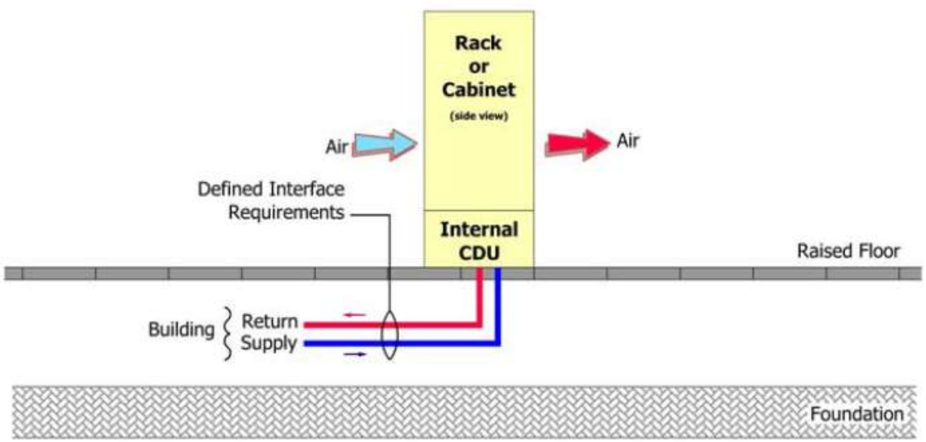
Rack or Cabinet
(side view)
Air
Air
Defined Interface Requirements
Internal CDU
Raised Floor
Building
Return
Supply
Foundation

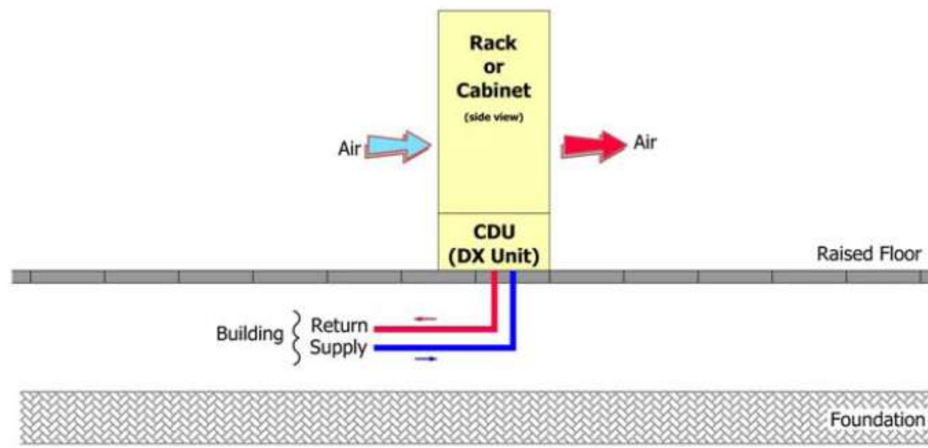
Rack or Cabinet
(side view)
Air
Air
CDU (DX Unit)
Raised Floor
Building
Return
Supply
Foundation

Characteristics of traditional hard piping methods

APC white paper #131

- **Carbon steel pipe schedule 40 and hard copper pipe type L or M are most commonly used**. Hard piping requires the use of threaded, grooved, welded or brazed fittings at every turn, at every valve, at every branch to multiple air conditioners and at every 1.8 or 6 meters (6 or 20 feet), depending on the available length of the pipe run. It is common to have multiple fittings in one pipe run from the chilled water source to the air conditioner.

Failure modes of hard piping

- Leak
- Galvanic Corrosion
- Condensation
- Mineral build-up (scaling)

Underfloor hard piping installation

***Figure 1** – Traditional underfloor chilled water piping installation with branches to different air conditioners using multiple fittings*

Overhead hard piping installation

Figure 2 – Overhead piping with drain pan above racks.

Flexible piping methodology

- The flexible piping is a multi-layered composite tubing consisting of an aluminum tubing sandwiched between inner and outer layers of cross-linked polyethylene. This gives the piping flexibility to be routed through the data center with the rigidity to stay in place. The cross-linked polyethylene or PEX also offers excellent protection against corrosion and the smooth interior walls and chemical properties make it resistant to mineral buildup with hard or soft water eliminating the risk of pinholes.

Overhead flexible piping installation

Figure 4 – Layout drawing of data center with flexible piping overhead

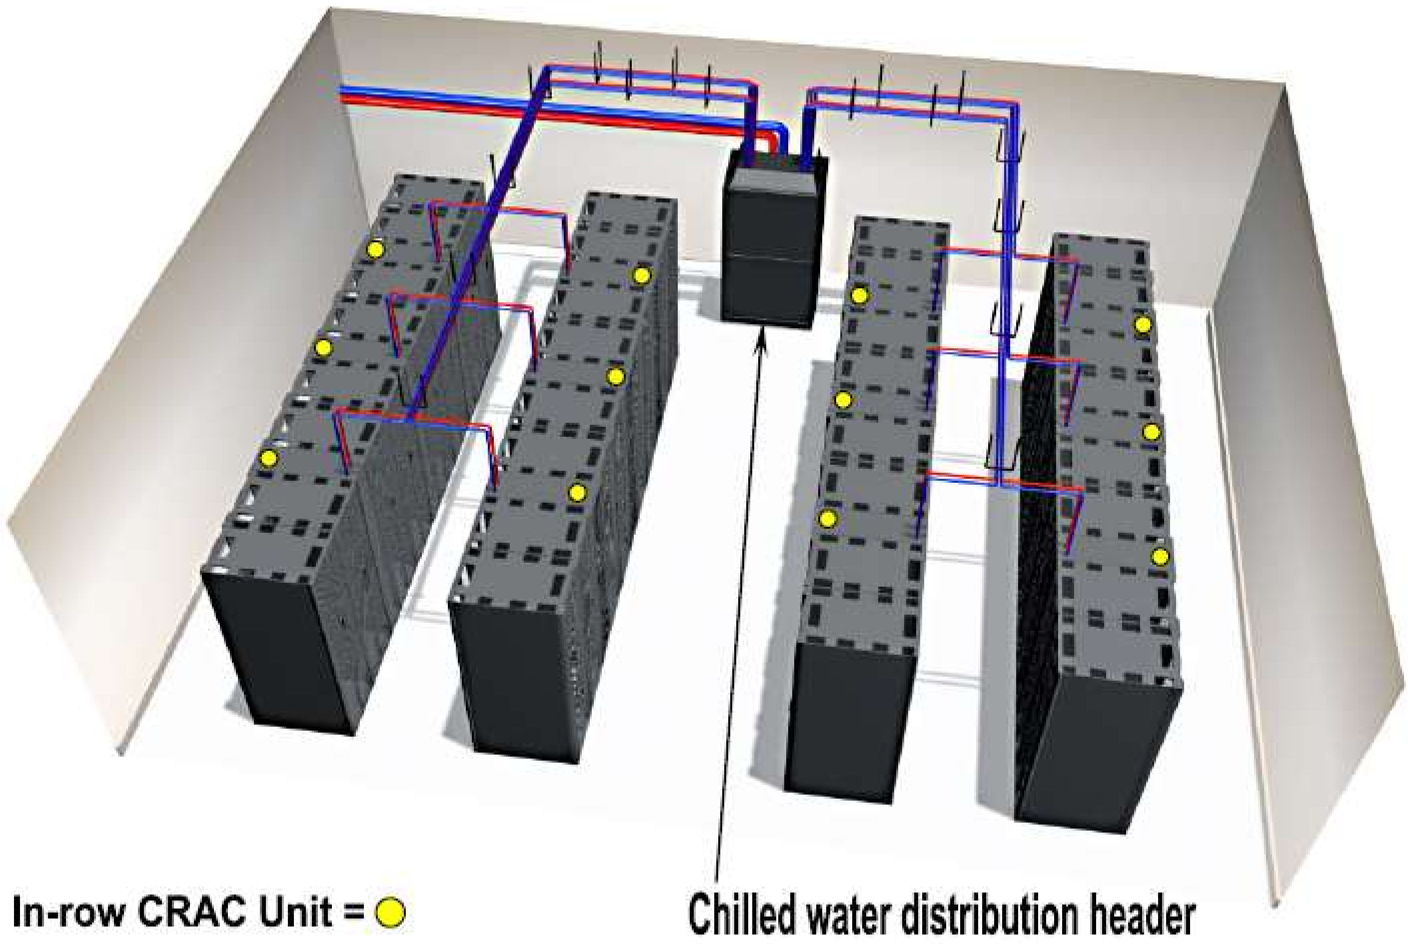

Comparison Between Hard Piping and Flexible Piping

Table 1

Physical attributes of hard and flexible piping

Physical attributes	Carbon steel schedule 40	Hard copper piping type "L"	Flexible piping PEX
Pipe weight in kg per linear meter (2.54 cm nominal size pipe without water)	2.49	0.975	0.324
Pipe weight in pounds per linear foot (1" nominal size pipe without water)	1.67	0.655	0.218
Temperature rating	Up to 399°C (750°F)	Up to 204°C (400°F)	Up to 93°C (200°F)
Rated internal working pressure in megapascal	19.7 MPa @ 38°C 19.7 MPa @ 93°C	3.41 MPa @ 38°C 2.79 MPa @ 93°C	1.38 MPa @ 23°C 0.689 MPa @ 93°C
Rated internal working pressure in psi	2857 psi @ 100°F 2857 psi @ 200°F	494 psi @ 100°F 404 psi @ 200°F	200 psi @ 73°F 100 psi @ 200°F
Type of fittings	Welded, brazed, grooved or threaded fittings	Soldered, brazed, grooved or threaded fittings	Multipress threaded or compression fittings
Size range	3.2 to 660 mm (1/8" to 26")	6.4 to 305 mm (¼" to 12")	12.7 to 5.08 mm (½" to 2") in North America 12.7 to 609 mm (½" to 24") in Europe[3]
Termination connection	Welded, brazed or threaded	Soldered, brazed or threaded	Multipress threaded or compression
Corrosion resistance	Limited, depends on the relative humidity of the environment and PH of water	Very good	Excellent
Thermal conductivity	High	High	Medium to Low

	Hard piping	Flexible piping
Agility	Slow speed of deployment due to multiple brazed joints required. Balancing of system is not easily accessible either under the raised floor or above the ceiling tiles. Non-scalable expansions or relocations require one time engineering and downtime for other units.	Increased speed of deployment by 40%. Balancing of the water system is located in a centralized accessible location. Scalable, allows for moves, adds, changes, and future expansions without disturbing other units.
Availability	Leak potentials at every fitting and joint decreasing reliability.	Increased reliability by eliminating intermediate joints drastically reducing leak potential.
MTTR	If leakage occurs on the main, repair may take from hours to days depending on the leak. If leakage occurs on a distribution branch in the data center, repair may take several hours, causing shutdown for several units.	If leakage occurs from the chiller to the centralized distribution header, repair may take from hours to days depending on the leak. If leakage occurs on a flexible branch in the data center, new flexible piping can be routed and repair may take up several hours causing shutdown on one unit only.
Installation	Higher installation costs. System balancing requires more time adding cost to start-up. Brazed, threaded, or mechanical joints and fittings are used, and intermediate isolation and balancing valves are required.	Lower installation cost. System start-up and balancing is less complex with the centralized distribution system. No brazed joints, intermediate fittings, or valves are required.
Turning radius	Allows a shorter turning radius using elbow fittings.	Minimum bending radius is 5 to 7 times the outside diameter of the tube.
Maintainability	Visual checks for leaks at each joint and valve, visual check for condensation at fittings and valves and visual check at corrosion points. Water and glycol concentration measured and validated.	Less time spent in visual checking for leaks and condensation formation on valves at the centralized distribution header (all valves are in one location). Water and glycol concentration measured and validated, routine maintenance
Pressure drop	The use of elbows for turns and mineral buildup causes additional pressure drop	Smooth interior and larger radius turns without fittings reduce the pressure drop for typical piping runs
White space	Piping is run underfloor or overhead, no white space is occupied by the piping system	White space is required for the centralized distribution header in the room.
Distances	Long pipe distances can be performed with hard pipe since several pieces of pipe are joined through fittings.	Maximum distance recommended is 46 meters (150 ft) from the distribution header to the air conditioners due to the complexity that longer distances would create for the installer.
Upfront cost (installation and material)	Hard pipe cost is lower but the overall installation cost is higher due to the increased labor required for brazing and threads and system balancing requires more time adding cost to start-up.	PEX piping has a higher cost, however the overall installation may be lower due to the elimination of brazing or threaded fittings and the system start-up and balancing is less complex with the centralized distribution system.
Pipe location	Can be installed outdoors or exposed to sunlight.	PEX must not be stored or installed in areas where it is exposed to sunlight, either direct or indirect.

Table 3

Failure mode comparison of hard and flexible piping

	Hard piping	Flexible piping
Punctures	Less susceptible to leakage due to puncture by a sharp object.	More susceptible to leakage due to puncture by a sharp object.
Single point failures	Failure in a branching pipe causes loss of cooling in all CRACs connected to the branch.	Failure in a line causes loss of cooling in only one CRAC.
Joint leaks	Multiple joints and fittings in the pipe increase leak potential due to possible galvanic corrosion, failure of thread sealant over time, poor machining of the threads, gasket deterioration in grooved connections or poor quality of the threaded fittings.	Reduced amount of joints - two per line per CRAC. Multipress threaded fittings crimp the PEX-AL-PEX tube making a stronger connection than a threaded or gasketed fitting.
Earthquake / vibration	Vibration or earthquake movement can cause leakage at joints and fittings.	Less susceptible to break or leak in vibration or earthquake conditions.
Stepping on	May damage brazed or threaded fittings which can produce a leak.	Less susceptible to damage due to the flexibility of the pipe.
Insulation dripping from condensation in the data center.	More potential for condensation due to difficulty to insulate multiple valves, strainers, and fittings. Small cracks or spaces left without insulation may cause condensation.	Less potential for condensation due to the elimination of intermediate valves or fitting between the distribution system and the CRACs.
Abrasions / cuts	Resistant to exterior abrasions or cuts	Less resistant to exterior abrasions. Cut can damage the PEX piping exterior.
Pinholes and mineral buildup	Susceptible to pinholes and leakage due to mineral buildup if water is not treated periodically.	Very resistant to mineral buildup due to smooth interior walls and chemical properties.

Note: shading indicates best performance for the characteristics

- The concern of water in the data center is also reduced with a flexible piping system for three reasons:
  - The overall piping system failure rate is greatly decreased due to the dramatic reduction in joints
  - The fundamental reliability of the base piping itself is higher
  - The potential for condensation is reduced by not having intermediate fittings or valves to insulate, which are the main points of condensation formation in a chilled water system.

Water Quality Specifications—TCS Cooling Loop

Parameter	Recommended Limits
pH	7 to 9
Corrosion inhibitor	Required
Biocide	Required
Sulfides	<1 ppm
Sulfate	<10 ppm
Chloride	<5 ppm
Bacteria	<100 CFU/mL
Total hardness (as $CaCO_3$)	<20 ppm
Conductivity	0.20 to 20 micromho/cm
Total suspended solids	<3 ppm
Residue after evaporation	<50 ppm
Turbidity	<20 NTU (nephelometric)

Air Cooling & Airflow Distribution

Basic airflow requirements
(APC white paper #49)

- The airflow in and around the rack cabinet is critical to cooling performance. **The key to understanding rack airflow is to recognize the fundamental principle that IT equipment cares about two things**:
  - **That appropriate conditioned air is presented at the equipment air intake**
  - **That the airflow in and out of the equipment is not restricted**.
- The two key problems that routinely occur and prevent the ideal situation are
  - The CRAC air becomes mixed with hot exhaust air before it gets to the equipment air intake
  - The equipment airflow is blocked by obstructions.

Air system design overview

- Data center layout
- Airflow configurations
  - Distribution: overhead or underfloor
  - Control: constant or variable volume
- Airflow issues
- Economizers
- Humidity control issues

Air cooling issues

- **Limitations on the data densities served (~200 W/ft^2)**
  - Air delivery limitations
  - Real estate
- Working conditions
  - Hot aisles are approaching OSHA (Occupational Safety & Health Administration) limits
- Costly infrastructure
- High energy costs
- Management over time
- Reliability
  - Loss of power recovery
  - Particulates

Wet Bulb Globe Temperature (WBGT)

- A composite temperature used to estimate the effect of temperature, humidity, wind speed (wind chill) and solar radiation on humans. It is used by industrial hygienists, athletes, and the military to determine appropriate exposure levels to high temperatures. It is derived from the following formula:

$$WBGT = 0.7T_w + 0.2T_g + 0.1T_d$$

- Where
- T_w = Natural wet-bulb temperature (combined with dry-bulb temperature indicates humidity)
- T_g = Globe thermometer temperature (measured with a globe thermometer, also known as a black globe thermometer, to measure solar radiation)
- T_d = Dry-bulb temperature (actual air temperature)
- Temperatures may be in either Celsius or Fahrenheit
- Indoors, or when solar radiation is negligible, the following formula is used:

$$WBGT = 0.7T_w + 0.3T_d$$

APC white paper #123

OSHA guidelines (US only)

Work / rest regimen (hourly duty cycles)	Light work load Max permissible WGBT	Moderate work load Max permissible WGBT	Heavy work load Mass permissible WGBT
Continuous	86F (30C)	80F (26.7C)	77F (25C)
75% work / 25% rest	87F (30.6C)	82F (27.8C)	78F (25.6C)
50% work / 50% rest	89F (31.7C)	85F (29.4C)	82F (27.8C)
25% work / 75% rest	90F (32.2C)	88F (31.1C)	86F (30C)

Airflow with constant volume systems

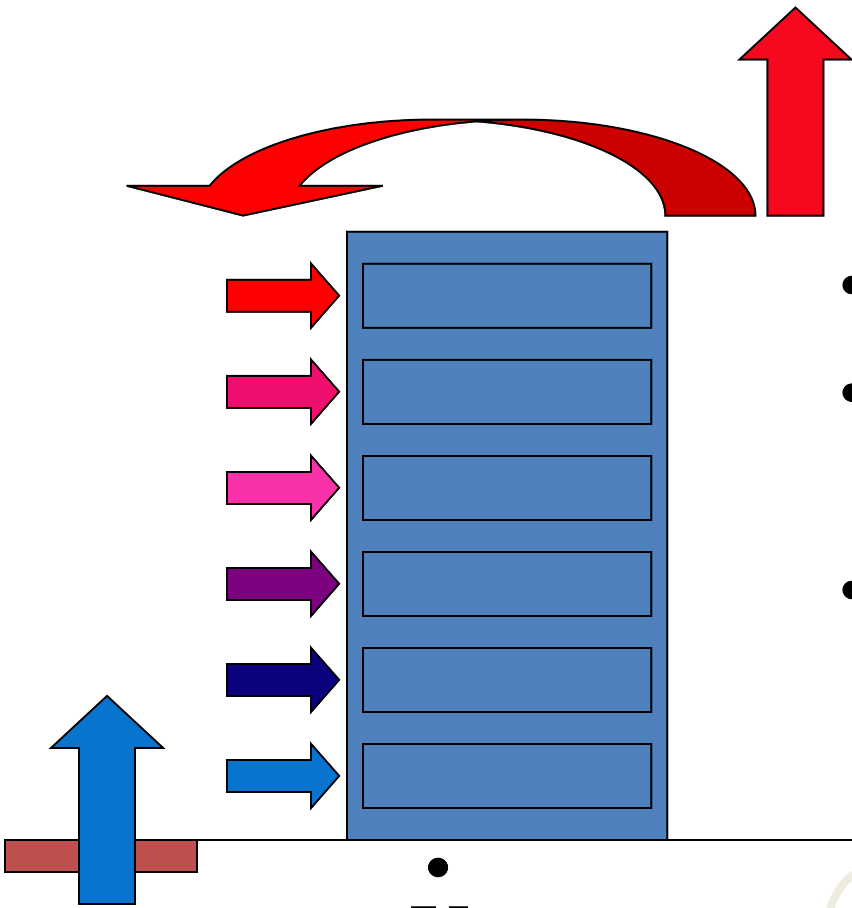

- Hot spots
- Higher hot aisle temperature
- Possible equipment failure or degradation

$$\dot{V}_{HVAC\_Supply} < \sum \dot{V}_{Servers}$$

Airflow with constant volume systems

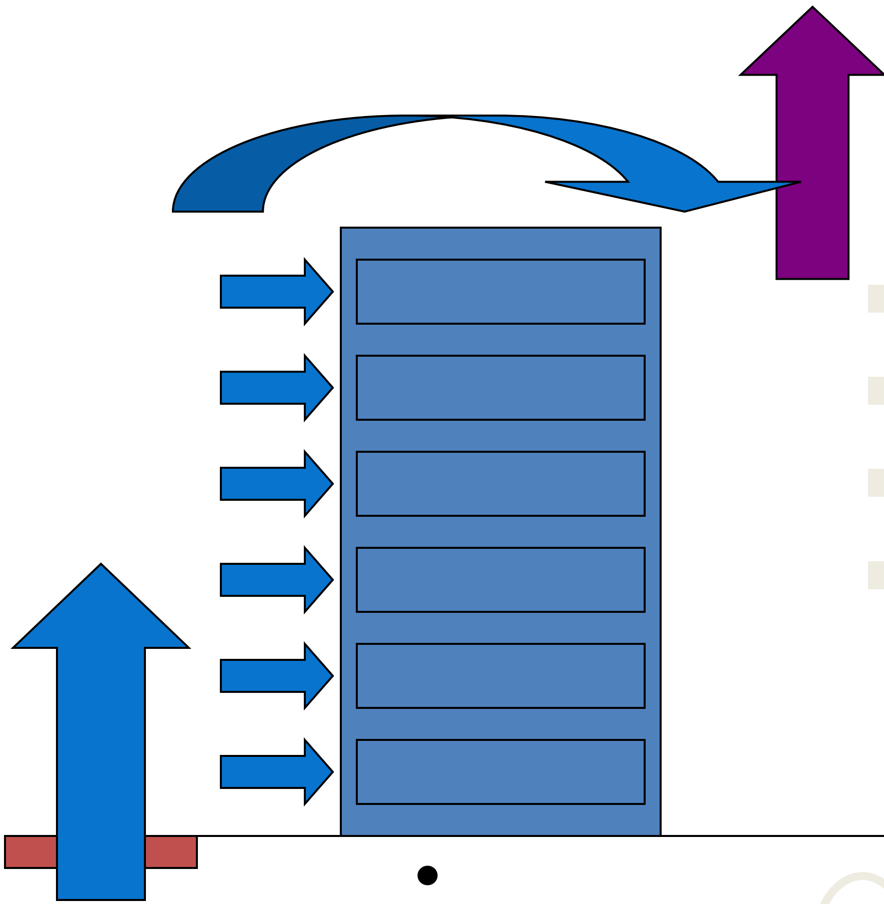

- Least hot spots
- Higher air velocities
- Higher fan energy
- Reduced economizer effectiveness (due to lower return temperatures)

$$\dot{V}_{HVAC\_Supply} > \sum \dot{V}_{Servers}$$

Airflow with constant volume systems

- Note most of these observations apply to overhead and underfloor distribution
- With constant volume fans on the servers you can only be right at one condition of server loading!
- The solution is to employ variable speed server and distribution fans…

Airflow with variable volume systems

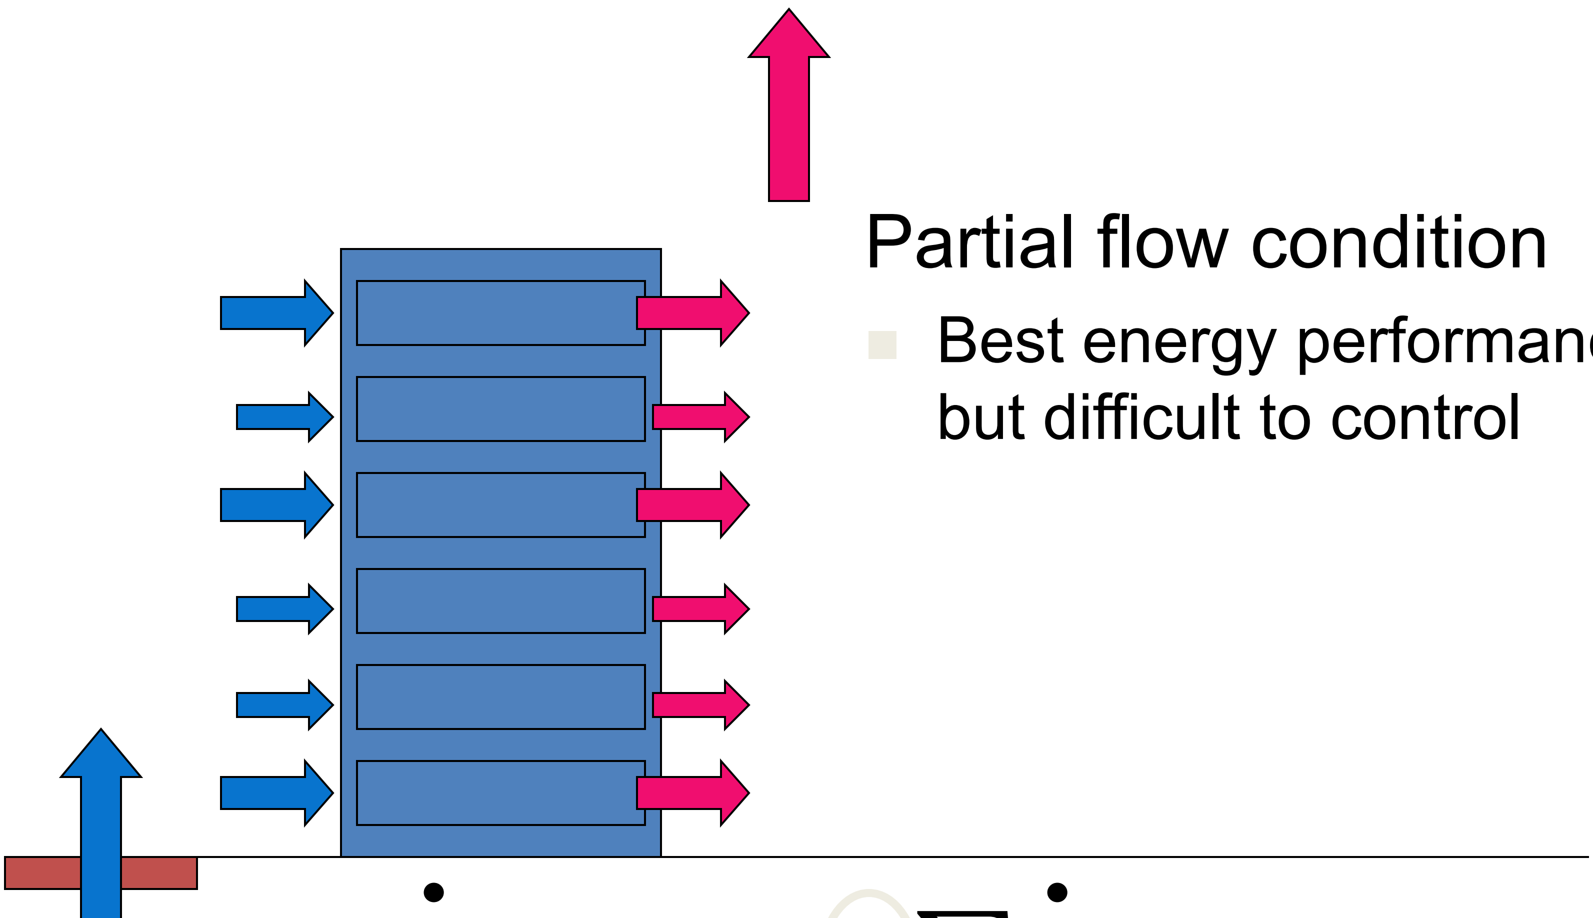

Partial flow condition

- Best energy performance but difficult to control

$$\dot{V}_{HVAC\_Supply} > \sum \dot{V}_{Servers}$$

How Do You Balance Airflow?

- Spreadsheet
- CFD
- Monitoring/Site Measurements

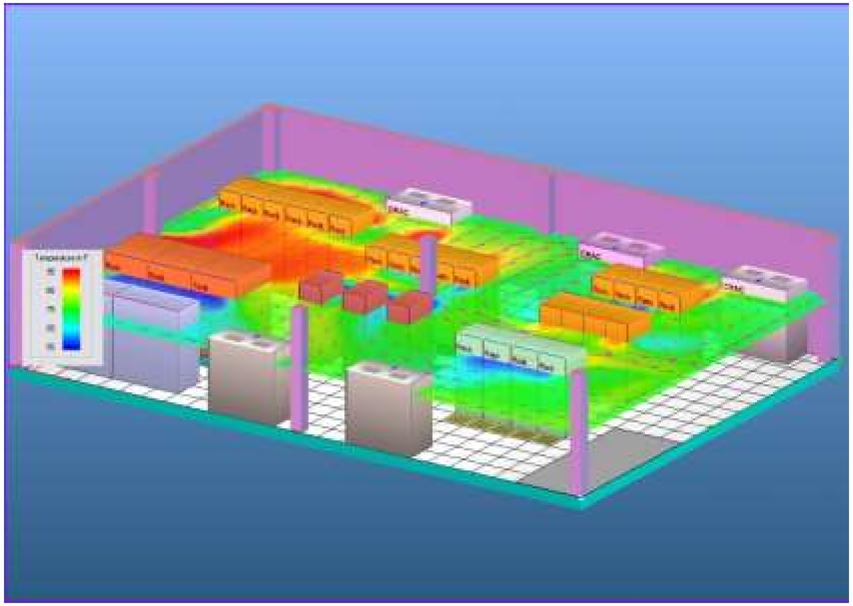

Image from TileFlow http://www.inres.com/Products/TileFlow/tileflow.html, Used with permission from Innovative Research, Inc.

What's the server airflow?

	SUN	SUN	DELL	DELL
	V490	V240	2850	6850
num fans	9	3	n/a	n/a
total CFM (max)	150	55.65	42	185
total CFM (min)			27	126
fan speed	single speed	variable	2 speed	2 speed
fan control	n/a	inlet temp.	77F inlet	77F inlet
Form Factor (in U's)	5	2	2	4
heat min config (btuh)		798		454
heat max config (btuh)	5,459	1,639	2,222	4,236
heat max (watts)	1,599	480	651	1,241
dT min config	-	13	-	3
dT max config	33	27	48	21
servers per rack	8	21	21	10
CFM/rack (hi inlet temp)	1,200	1,169	882	1,850
CFM/rack (low inlet temp)	1,200		567	1,260
max load / rack (kW)	13	10	14	12

Best air delivery practices (short summary)

- Arrange racks in hot aisle/cold aisle configuration
- Try to match or exceed server airflow by aisle
  - Get thermal report data from IT if possible
  - Plan for worst case
- Get variable speed or two speed fans on servers if possible
- Provide variable airflow fans for AC unit supply
  - Also consider using air handlers rather than CRACs for improved performance
- Use overhead supply where possible
- Provide aisle capping (preferably cold aisles, refer to LBNL presentation for more details)
- Plug floor leaks and provide blank off plates in racks
- Draw return from as high as possible
- Use CFD to inform design and operation

Short stops..

- Use air- or water-side economizers where possible
- Consider personal grounding in lieu of humidification
- Consider AHUs as an alternative to CRACs
- Consider VSDs on fans, pumps, chillers and towers
- Refer to ASHRAE, LBNL and Uptime Institute for more recommendations

Some data center airflow configurations

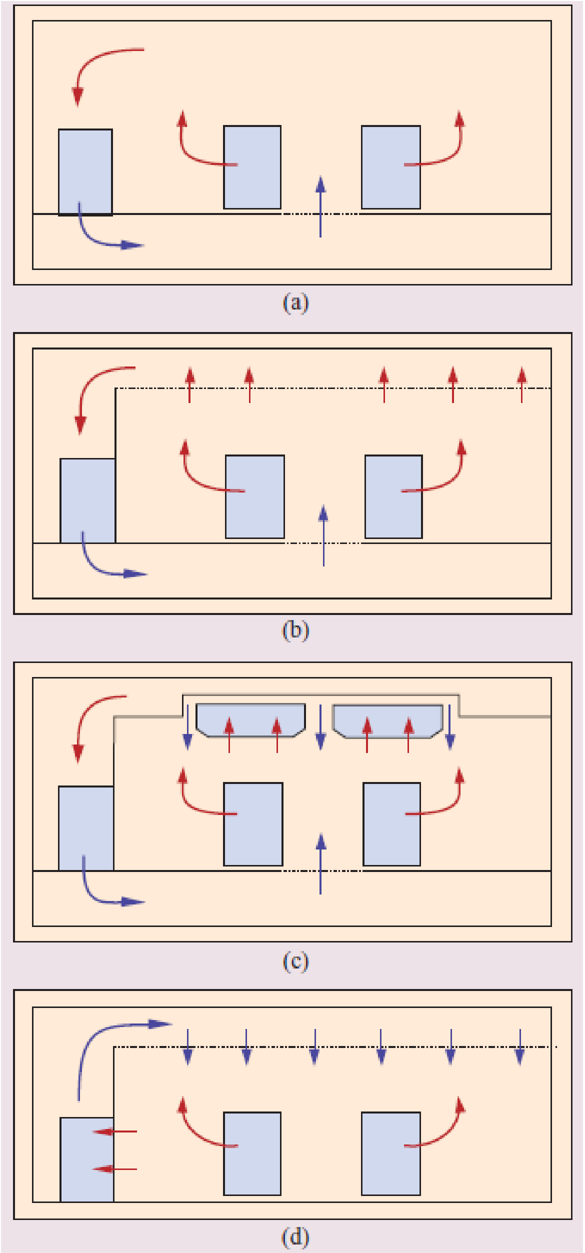

(a) raised-floor supply;

(b) raised-floor supply/ceiling return;

(c) raised-floor supply/ceiling supply;

(d) non-raised floor/ceiling supply

Data center layout

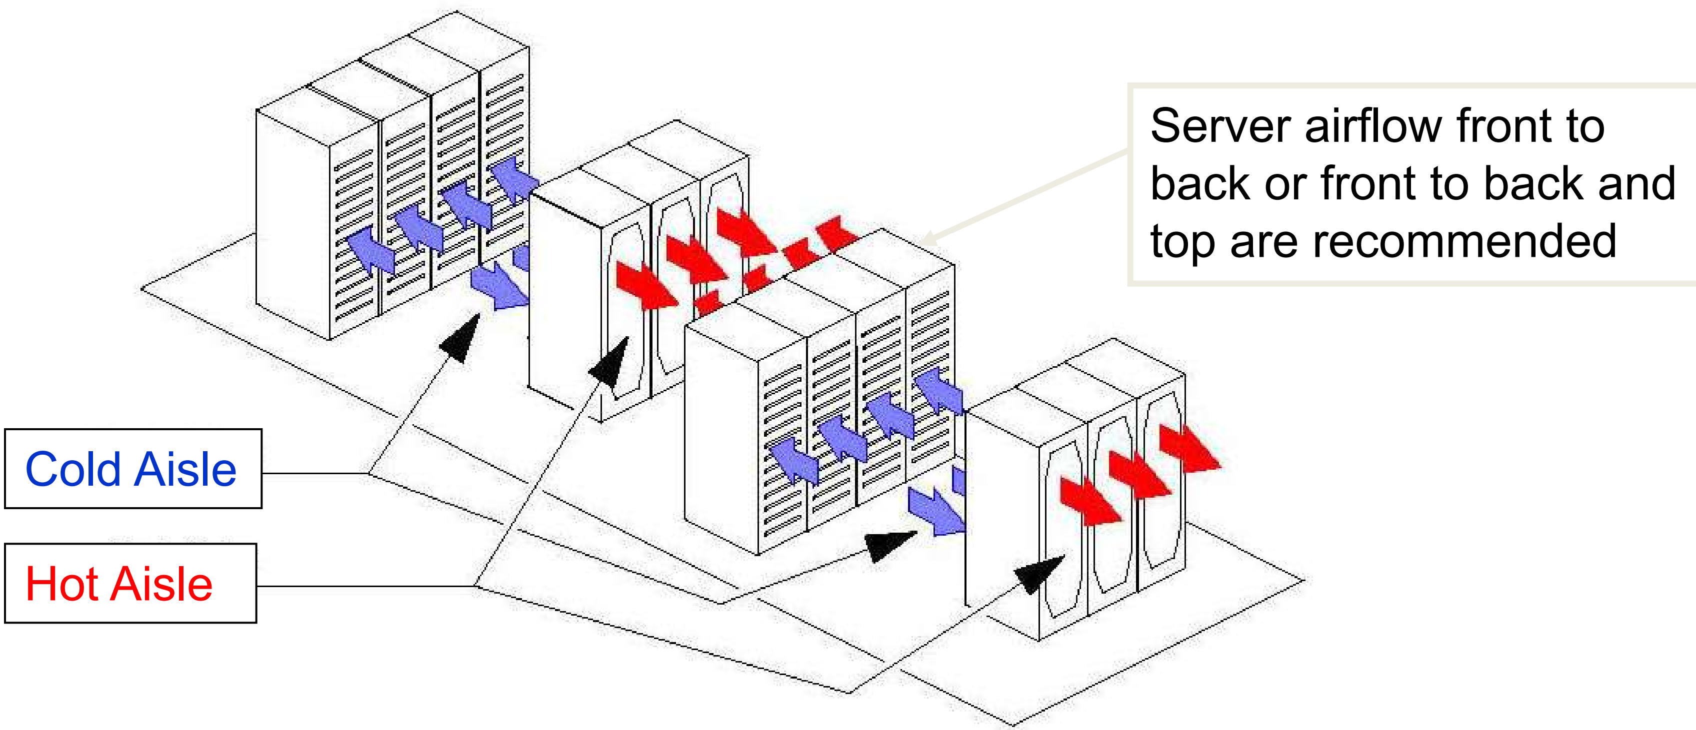

Data center layout

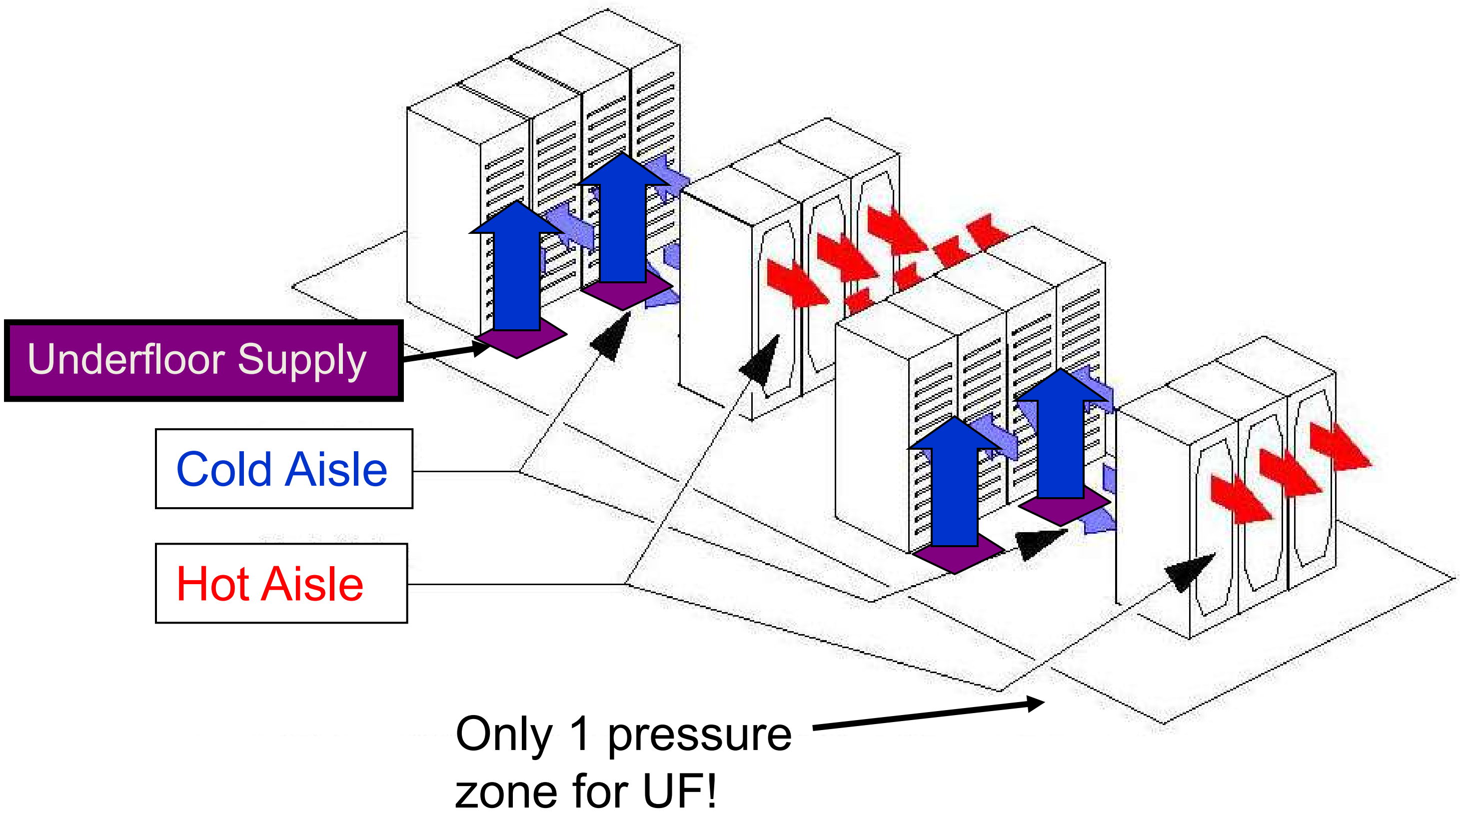

Data center layout

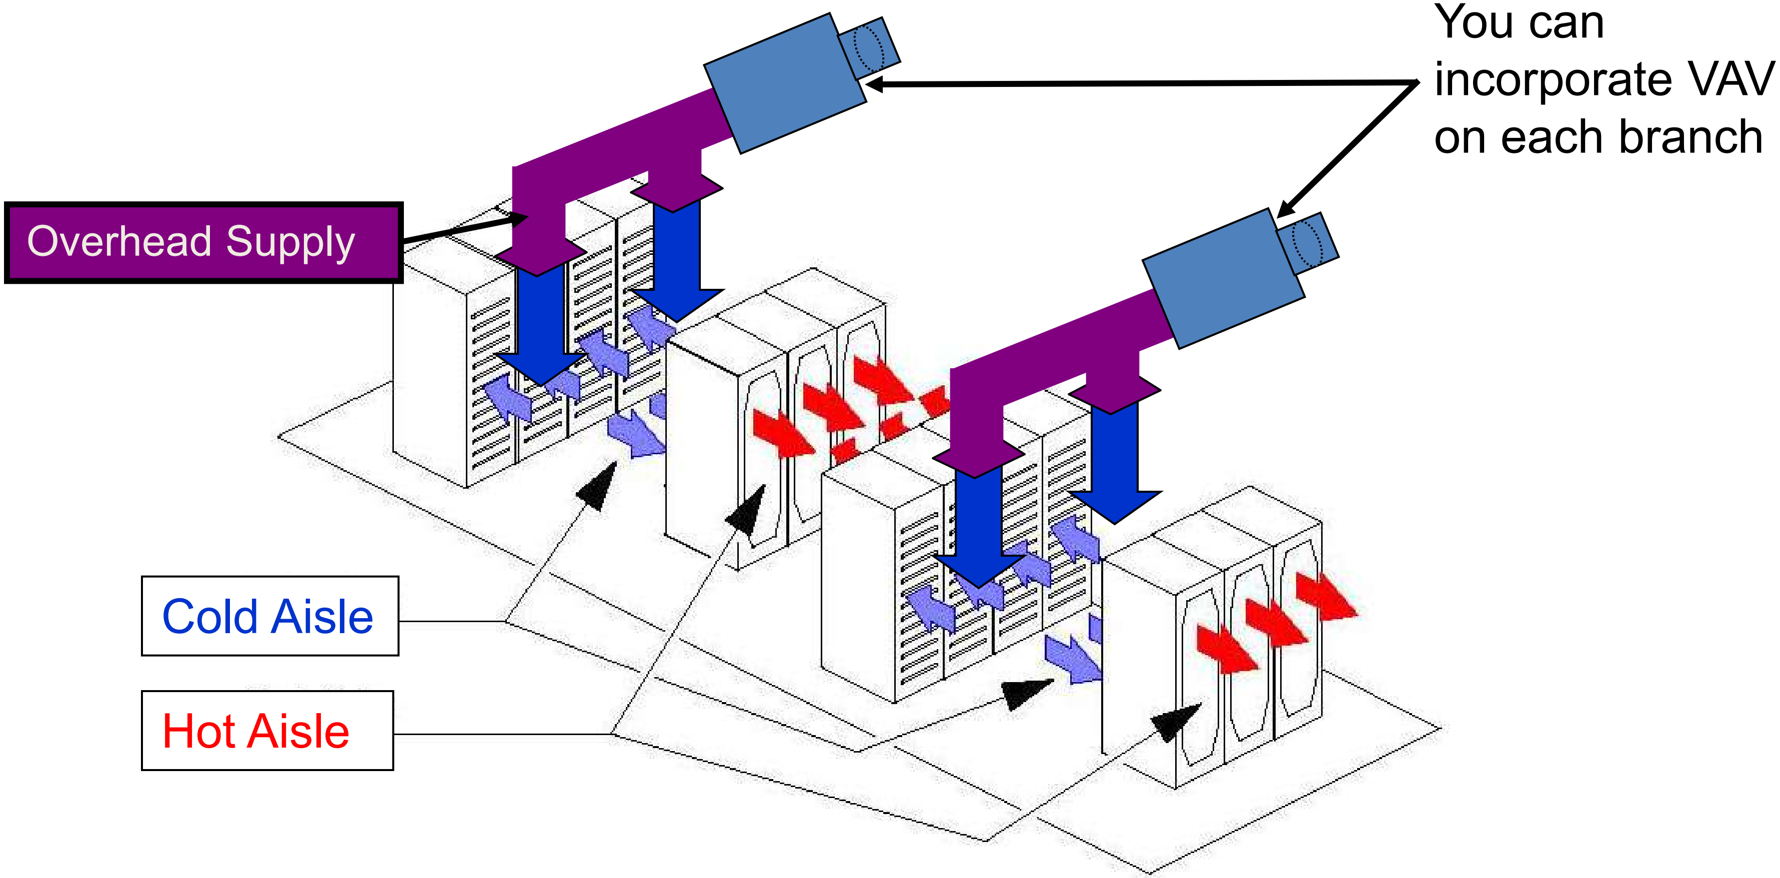

Typical temperature profile with UF supply

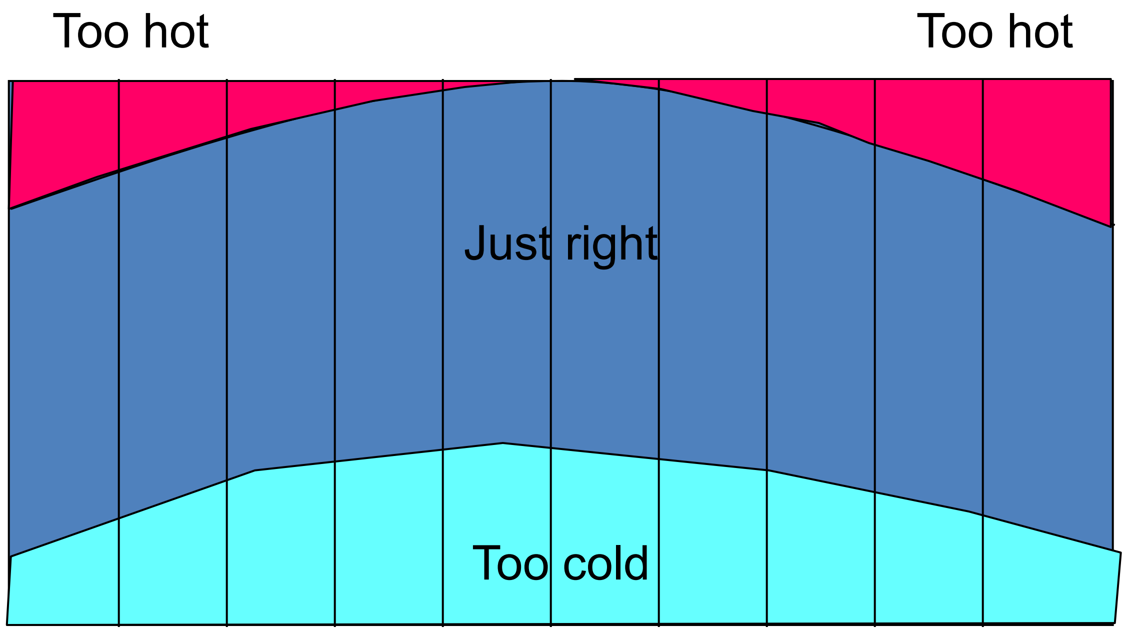

Elevation at a cold aisle looking at racks

There are numerous references in ASHRAE. See for example V. Sorell et al; "Comparison of Overhead and Underfloor Air Delivery Systems in a Data Center Environment Using CFD Modeling"; ASHRAE Symposium Paper DE-05-11-5; 2005

Typical temperature profile with OH supply

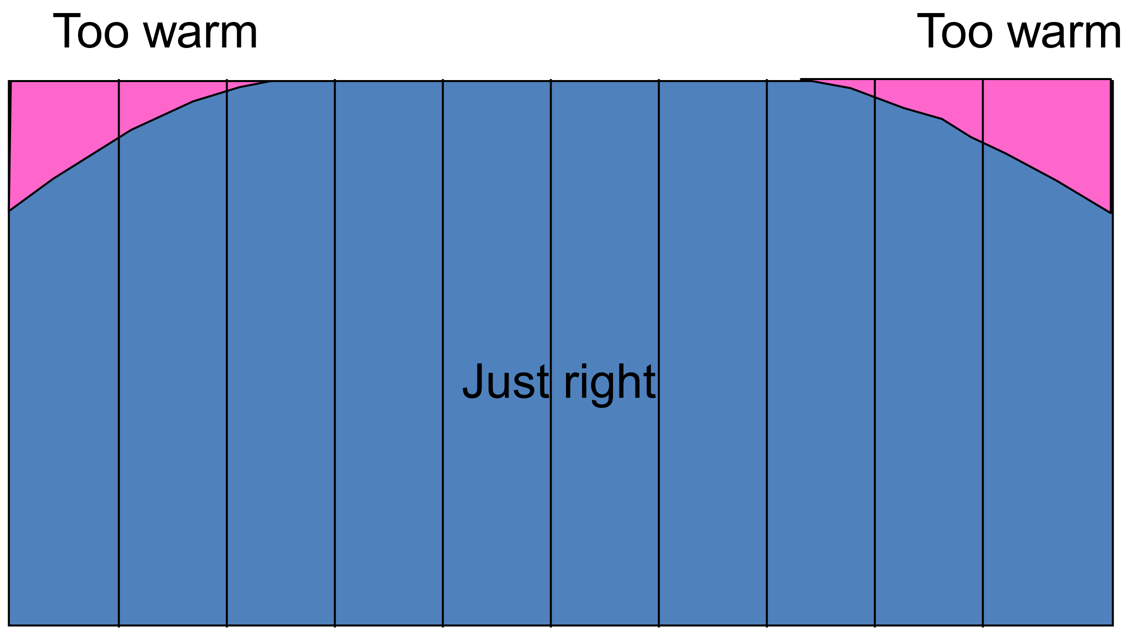

Elevation at a cold aisle looking at racks

國立交通大學 National Chiao Tung University

Aisle capping

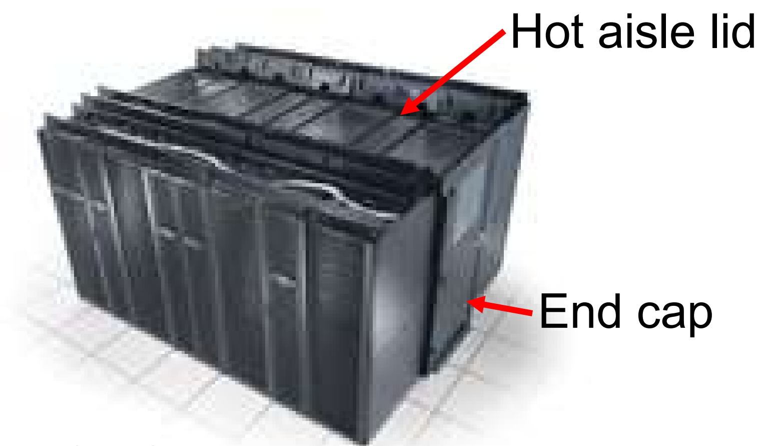

© APC reprinted with permission

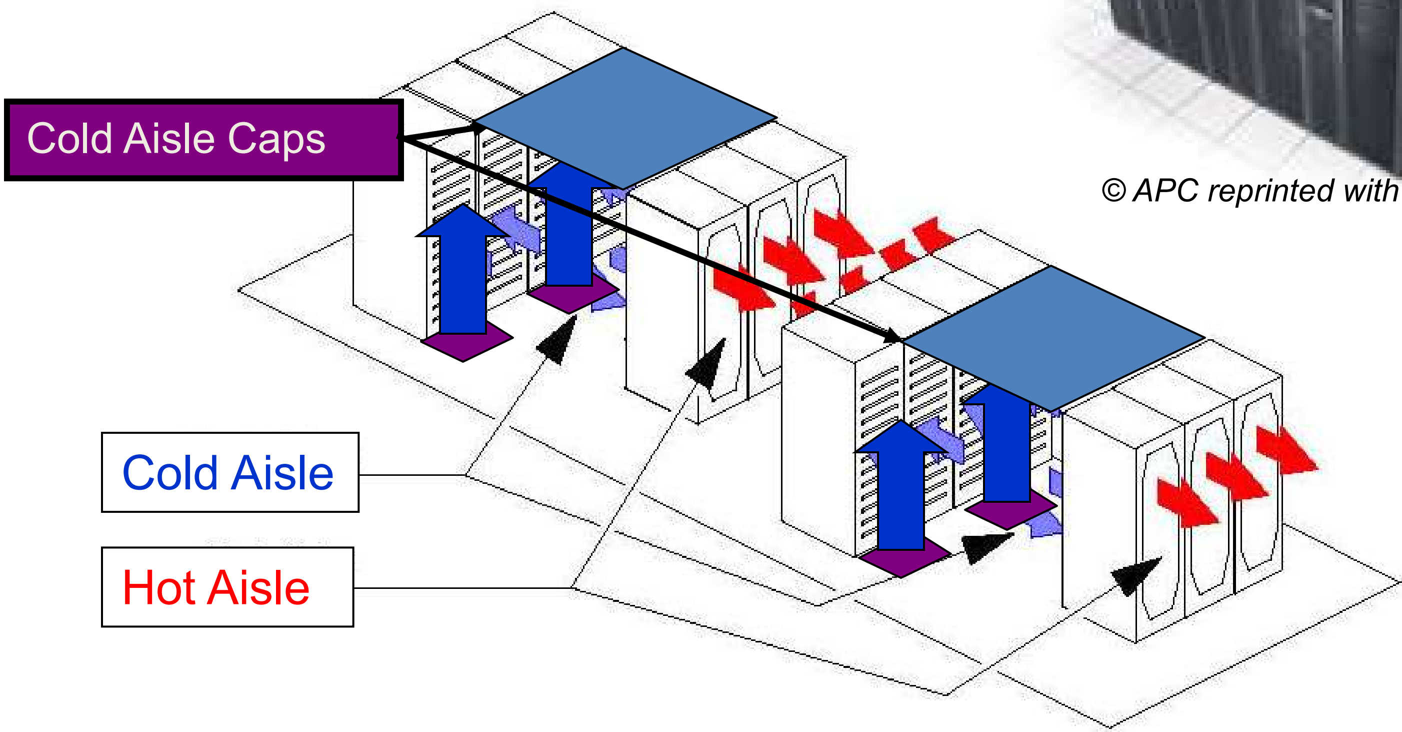

Aisle capping

LBNL has recently performed research on aisle capping

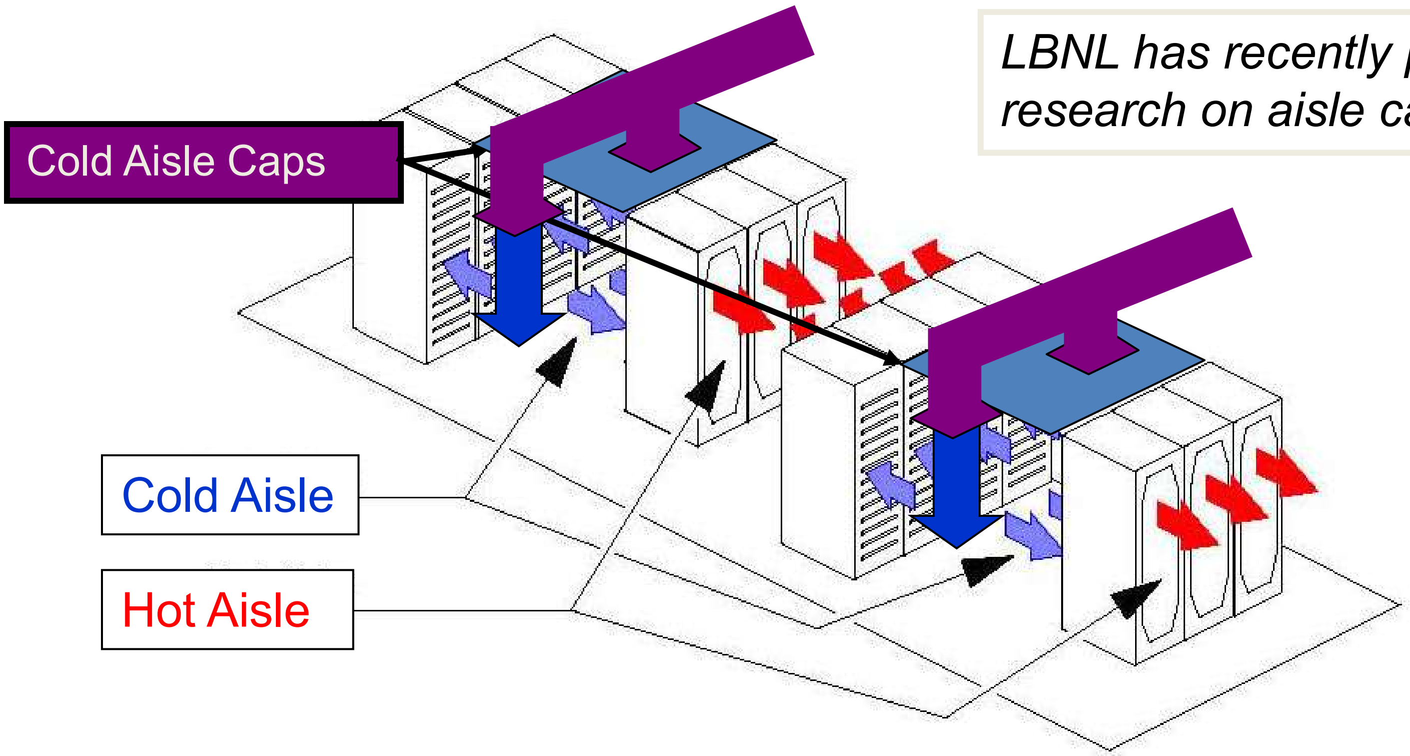

Overhead (OH) vs. Underfloor (UF)

Issue	**Overhead (OH) Supply**	**Underfloor (UF) Supply**
Capacity	Limited by space and aisle velocity.	Limited by free area of floor tiles.
Balancing	Continuous on both outlet and branch.	Usually limited to incremental changes by diffuser type. Some tiles have balancing dampers. Also underfloor velocities can starve floor grilles!
Control	Up to one pressure zone by branch.	Only one pressure zone per floor, can provide multiple temperature zones.
Temperature Control	Most uniform.	Commonly cold at bottom and hot at top.
First Cost	Best (if you eliminate the floor).	Generally worse.
Energy Cost	Best.	Worst.
Aisle Capping	Hot or cold aisle possible.	Hot or cold aisle possible.

- Prevention of hot and cold air mixing is a key to all efficient data center cooling strategies.
- Both HACS and CACS offer improved power density and efficiency when compared with traditional cooling approaches. A hot-aisle containment system (HACS) is a more efficient approach than a cold-aisle containment system (CACS) because it allows higher hot aisle temperatures and increased chilled water temperatures which results in increased economizer mode hours and significant electrical cost savings. Cooling set points can be set higher while still maintaining a comfortable temperature in the uncontained area of the data center.

APC white paper #46

Figure 2

Frequency distribution of actual rack power consumption, showing relation to maximum possible rack configuration

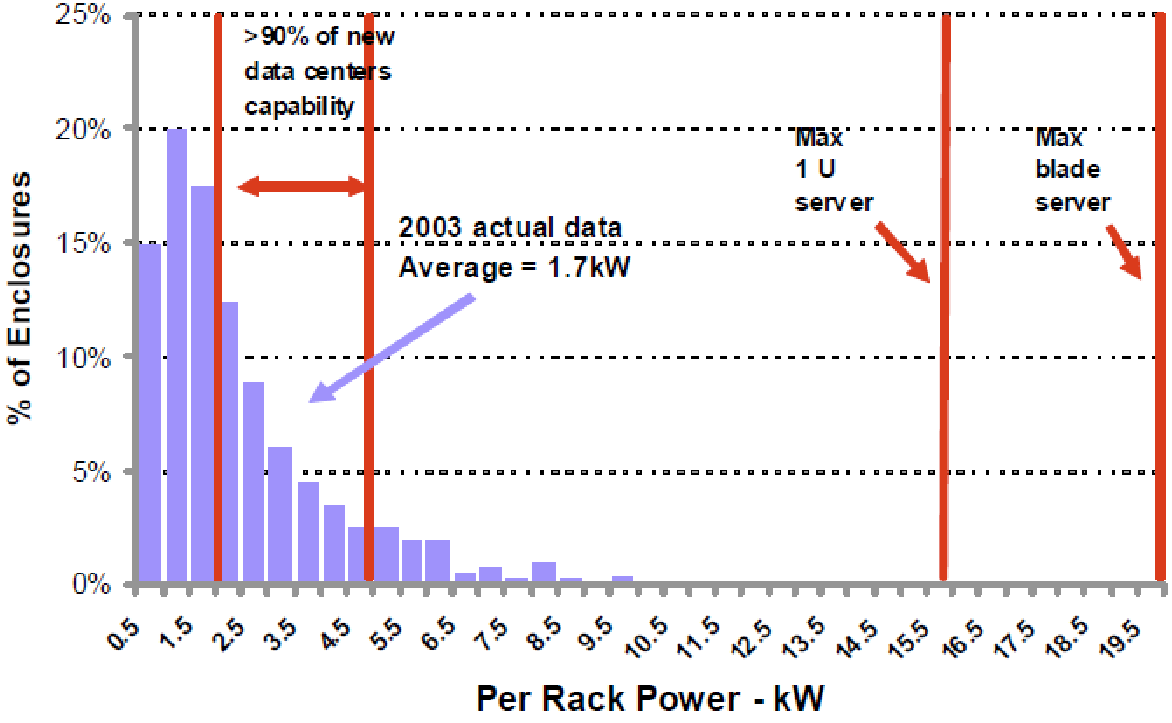

Room, row, and rack based cooling architectures

APC White paper #130

***Figure 1** – Floor plans showing the basic concept of room, row, and rack-oriented cooling architecture. Blue arrows indicate the relation of the primary cooling supply paths to the room.*

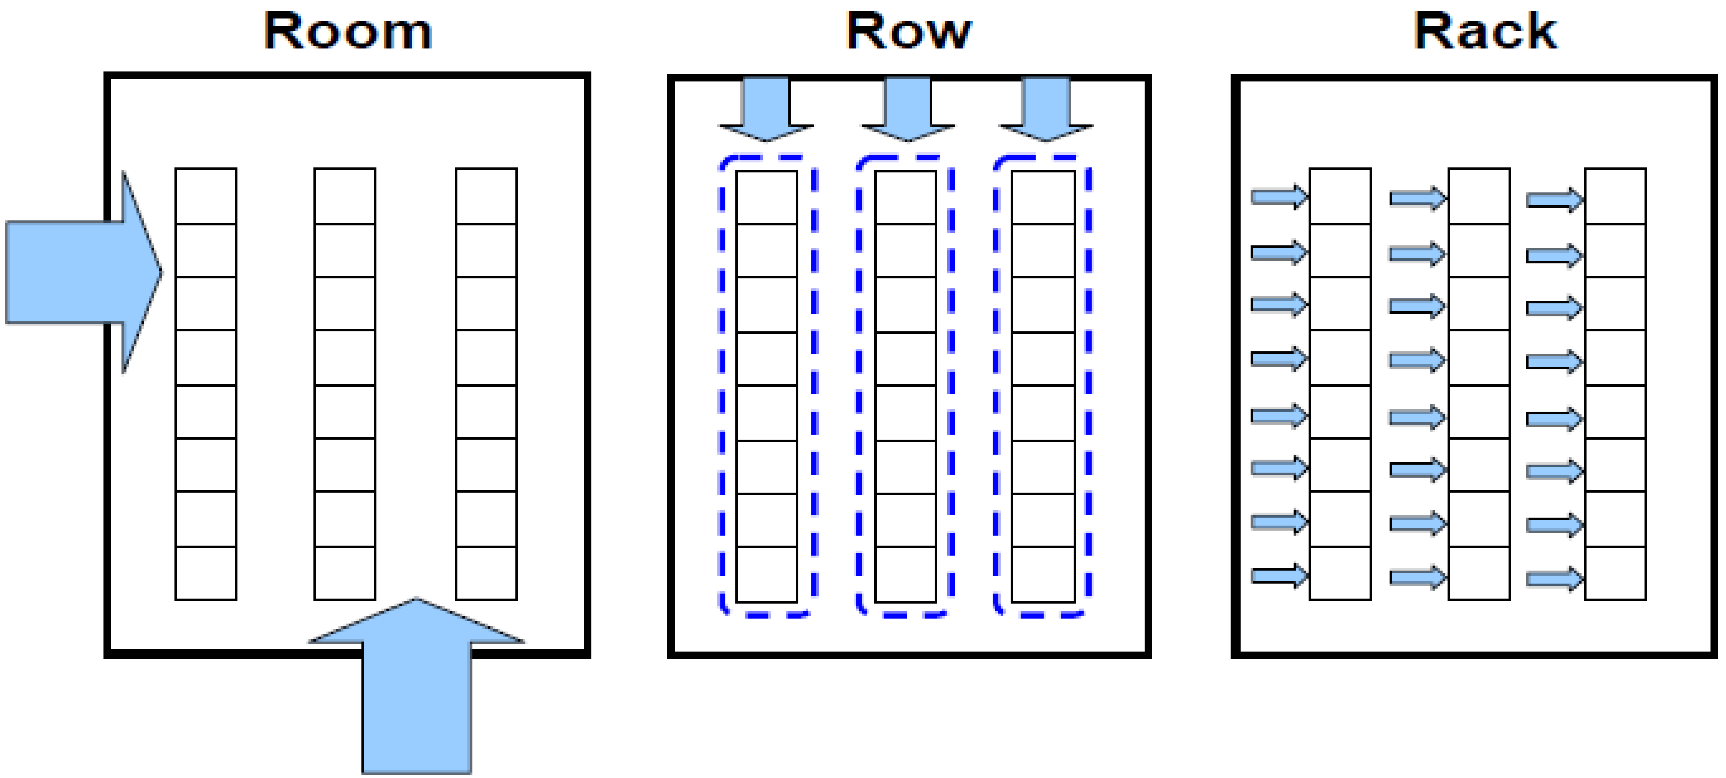

- Room-oriented: Supplying a room but primarily serving a low density area of mixed equipment such as communication equipment, low density servers, and storage. Target: 1-3 kW per rack, 323-861 W/m^2 (30-80 W/ft^2)
- Row-oriented: Supplying a high density or ultra-high density area with blade servers or 1U servers.
- Rack-oriented: Supplying isolated high density racks, or ultra-high density racks.

Room-oriented architecture

- In room-oriented architecture, the CRAC units are associated with the room and operate concurrently to address the total heat load of the room. A room-oriented architecture may consist of one or more air conditioners supplying cool air completely unrestricted by ducts, dampers, vents, etc. or the supply and/or return may be partially constrained by a raised floor system or overhead return plenum.

- The room-oriented design is **heavily affected by the unique constraints of the room**, including the **ceiling height, the room shape, obstructions** above and under the floor, rack layout, CRAC location, the distribution of power among the IT loads, etc. The result is that performance prediction and performance uniformity are poor, particularly as power density is increased.
- Computational fluid dynamics (CFD) may be required to help understand the design performance of specific installations.

- Another significant shortcoming of room-oriented architecture is that in many cases **the full rated capacity of the CRAC cannot be utilized**. This condition is a result of room design and occurs when **a significant fraction of the air distribution pathways from the CRAC units bypass the IT loads and return directly to the CRAC**. This bypass air represents CRAC airflow that is not assisting with cooling of the loads; in essence a decrease in overall cooling capacity. The result is that cooling equipments of the IT layout can exceed the cooling capacity of the CRAC.

Row-oriented architecture

- With a row-oriented architecture, the CRAC units are associated with a row and are assumed to be dedicated to a row for design purposes. The CRAC units may be mounted among the IT racks, they may be mounted overhead, or they may be mounted under the floor.

Figure 2a – *In-row cooling solution*

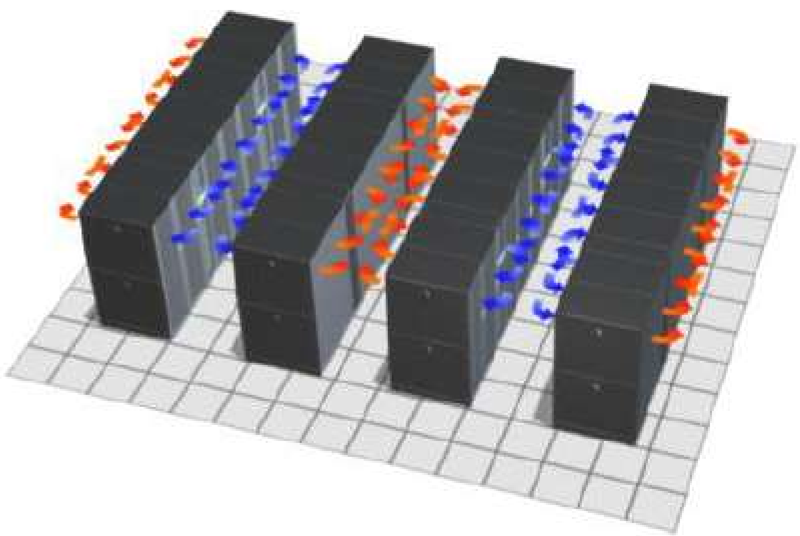

Figure 2b – *Overhead cooling solution*

Rack-oriented architecture

- The CRAC units are associated with a rack and are assumed to be dedicated to a rack for design purposes. The CRAC units are directly mounted to or within the IT racks. Compared with the room-oriented or row-oriented architecture, **the rack-oriented airflow paths are even shorter and exactly defined, so that airflows are totally immune to any installation variation or room constraints**. All of the rated capacity of the CRAC can be utilized, and the highest power density (**up to 50 kW per rack**) can be achieved.

- The rack-oriented architecture has other unique characteristics in addition to extreme density capability. The reduction in the airflow path length reduces the CRAC fan power required, increasing efficiency.

Figure 3 – *Rack cooling solution with cooling completely internal to rack*

Mixed architecture

Figure 4 – *Floor layout of a system utilizing room, row, and rack-oriented architectures simultaneously*

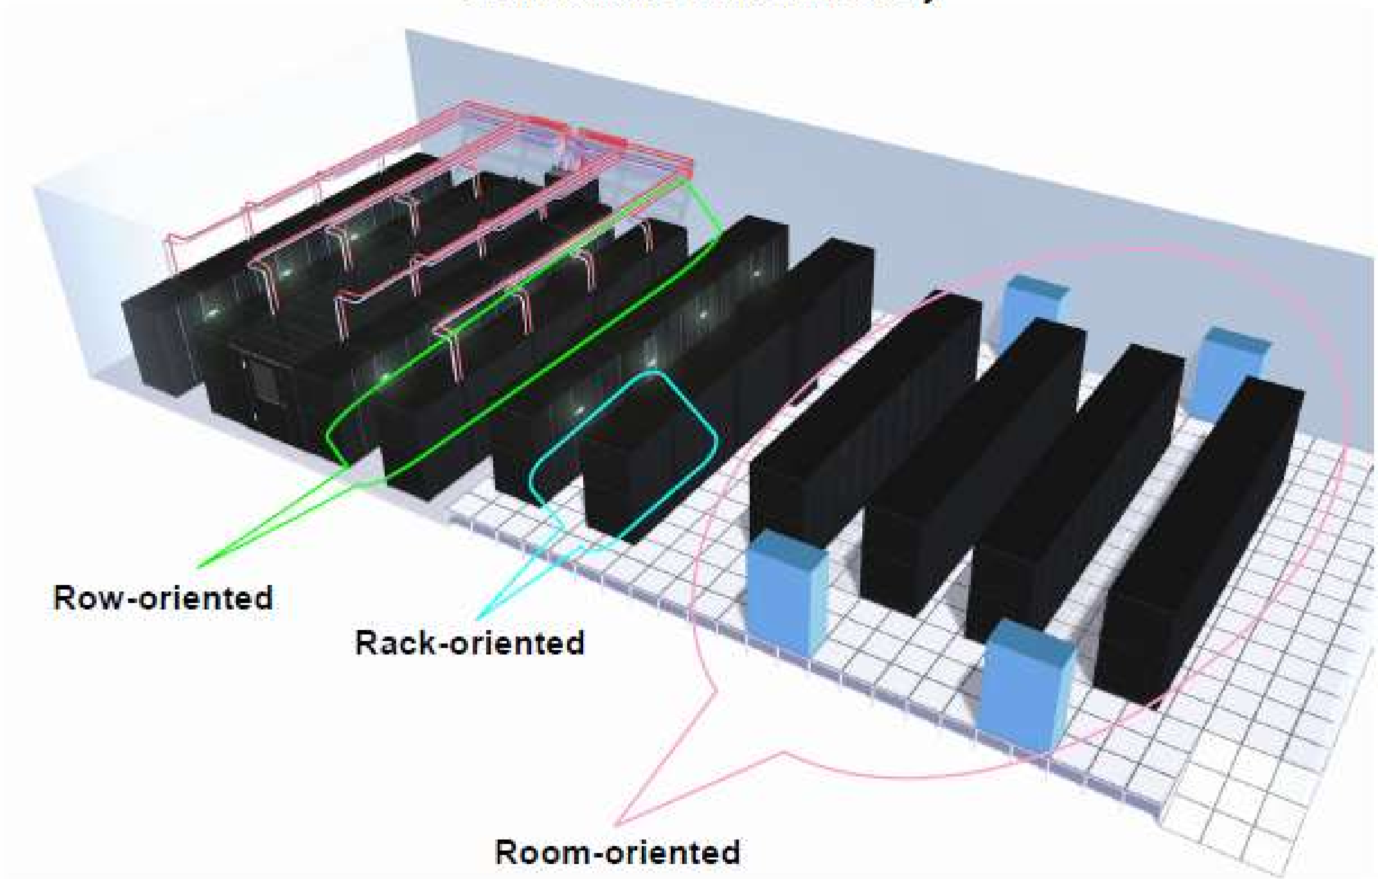

Benefit comparison of cooling architectures

- Agility
- System availability
- Lifecycle costs (TCO)
- Serviceability
- Manageability

Table 1 – Effectiveness of the room, row, and rack-oriented cooling architectures in addressing agility challenges. Best performance highlighted in blue.

Agility Challenges			
Challenge	**Rack**	**Row**	**Room**
Plan for a power density that is increasing and unpredictable	Modular; deployable at rack level increments targeted at specific density	Modular; deployable at row level increments targeted at specific density	Complex to upgrade or adapt; typically built out in advance of requirement
Reduce the extensive engineering required for custom installations	Immune to room effects; rack layout may be completely arbitrary	Immune to room effects when rows laid out according to standard designs; configure with simple tools	Complex CFD analysis required which is different for every room
Adapt to ever-changing requirements or any power density	Rack cooling capacity that is not used cannot be used by other racks	Cooling capacity is well defined and can be shared across a group of racks	Any change may result in overheating; complex analysis required to assure redundancy and density are achieved
Allow for cooling capacity to be added to an existing operating space	New loads may be added that are completely isolated from the existing cooling system; limited to rack cooling capacity	New loads may be added that are completely isolated from the existing cooling system; each additional cooling system increases density for entire row	May require shutdown of existing cooling system; requires extensive engineering
Provide a highly flexible cooling deployment with minimal reconfiguration	Racks may need to be retrofit or IT equipment moved to accommodate new architecture	Requires the rack rows to be spaced to accommodate or changes to overhead infrastructure for new architecture	Floor tiles can be reconfigured quickly to change cooling distribution pattern for power densities <3 kW

Table 2 – Effectiveness of the room, row, and rack-oriented cooling architectures in addressing availability challenges. Best performance highlighted in blue.

Availability Challenges			
Challenge	**Rack**	**Row**	**Room**
Eliminate hot spots	Closely couples heat removal with the heat generation to eliminate mixing The airflow is completely contained in the rack	Closely couples heat removal with the heat generation to minimize mixing	Supply and return paths promote mixing; engineered ductwork required to separate air streams
Assure redundancy when required	2N cooling capacity required for each rack; many rack cooling systems are not redundant capable	Utilizes shared N+1 capacity across common air return	Complex CFD analysis required to model failure modes; requires localized redundancy
Eliminate vertical temperature gradients at the face of the rack	Heat captured at the rear of the rack before mixing with cold supply air	Heat captured at the rear of the rack before mixing with cold supply air	Warm air may recirculate to front of rack as a result of insufficient heat removal or supply
Minimize the possibility of liquid leaks in the mission critical installation	Operates at warmer return temperatures to reduce or eliminate moisture removal and make-up sources. Rack targeted cooling requires additional piping and leakage points	Operates at warmer return temperatures to reduce or eliminate moisture removal and make-up sources	Mixed air return promotes the production of condensate and increases requirement for humidification
Minimize human error	Standardized solutions are well documented and can be operated by any user	Standardized solutions are well documented and can be operated by any user	Uniquely engineered system requires a highly trained and specialized operator

Table 3 – Effectiveness of the room, row, and rack-oriented cooling architectures in addressing lifecycle cost challenges. Best performance highlighted in blue.

Lifecycle Cost Challenges			
Challenge	**Rack**	**Row**	**Room**
Optimize capital investment and available space	Dedicated system for each rack may result in oversizing and wasted capacity	Ability to match the cooling requirements to a much higher percentage of installed capacity	System performance is difficult to predict, resulting in frequent oversizing
Accelerate speed of deployment	Pre-engineered system that eliminates or reduces planning and engineering	Pre-engineered system that eliminates or reduces planning and engineering	Requires unique engineering that may exceed the organizational demand
Lower the cost of service contracts	Standardized components reduce service time and facilitate the ability for user serviceability. Likely higher number of units with 1:1 ratio to IT rack enclosures.	Standardized components reduce service time and facilitates the ability for user serviceability	Specialized service contracts required for custom components
Quantify the return on investment for cooling system improvements	Standardized components for accurate measurement of system performance	Standardized components for accurate measurement of system performance	Customer engineered solutions makes system performance difficult to predict
Maximize the operational efficiency by matching capacity to load	Cooling system will likely be oversized and full potential not realized.	Right-sized cooling capacity to the cooling load matching heat load to installed capacity	Air delivery dictates oversized capacity; pressure requirements for under floor delivery are a function of the room size and floor depth.

Table 4 – Effectiveness of the room, row, and rack-oriented cooling architectures in addressing serviceability challenges. Best performance highlighted in blue.

Serviceability Challenges			
Challenge	**Rack**	**Row**	**Room**
Decrease Mean-Time-To-Recover (includes repair time plus technician arrival, diagnosis, and parts arrival times)	Modular components reduces downtime; 2N redundancy required for system repair and maintenance	Modular components reduces downtime; N+1 or excess capacity allows for repair without interruption to system performance	Custom spare parts are not readily available and require trained technician extending recovery time
Simplify the complexity of the system	Standardized components reduce the technical expertise required for routine service and maintenance	Standardized components reduce the technical expertise required for routine service and maintenance	Operation and repair of the system requires trained experts.
Simpler service procedures	In-house staff can perform routine service procedures. Modular subsystems with interfaces that mistake-proof service procedures.	In-house staff can perform routine service procedures. Modular subsystems with interfaces that mistake-proof service procedures.	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.
Minimize vendor interfaces	Modular units designed to integrate with a small set of ancillary systems	Modular units designed to integrate with a small set of ancillary systems	Engineered solution with multi-vendor subsystems
Learn from past problems and share learning across systems	Standardized building block approach with single rack and cooling unit interaction maximizes learning	Standardized building block approach with low interactions increases learning but with fewer systems to learn from	Unique floor layouts all have unique problems, limiting learning

Table 5 – Effectiveness of the room, row, and rack-oriented cooling architectures in addressing manageability challenges. Best performance highlighted in blue.

Manageability Challenges			
Challenge	**Rack**	**Row**	**Room**
System menu must be clear and provide ease of navigation	Low option configuration allows user to navigate through menu interface quickly	Low option configuration allows user to navigate through menu interface quickly	Highly configurable system complicates the menu structure. Requires advanced service training
Provide predictive failure analysis	Ability to provide real-time models of current and future performance.	Ability to provide near real-time models of current or future performance as a result of limited control effects	Virtually impossible to provide real-time models of current or future performance due to room-specific effects
Provide, aggregate, and summarize cooling performance data	Cooling capacity information at the rack level is determined and available in real time	Cooling capacity information at the row level is determined and available in real time. Rack level information can be effectively estimated.	Cooling capacity information is not available at the rack or row level

Row based cooling can improved Cooling efficiency (APC white paper #126)

Figure 2

2a (left)
Row floor mounted cooling

2b (right)
Row overhead cooling

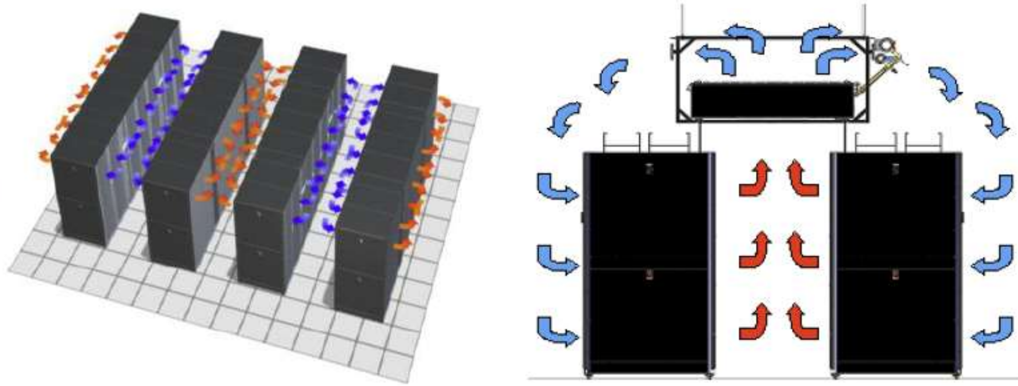

Figure 6

Cooling used row-based CRACs with shorter air flow paths

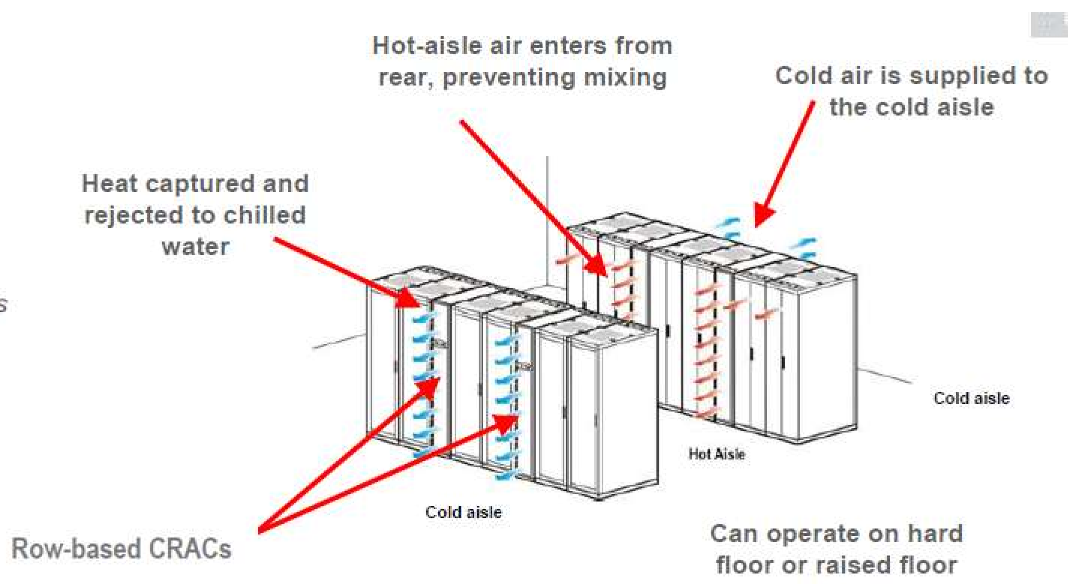

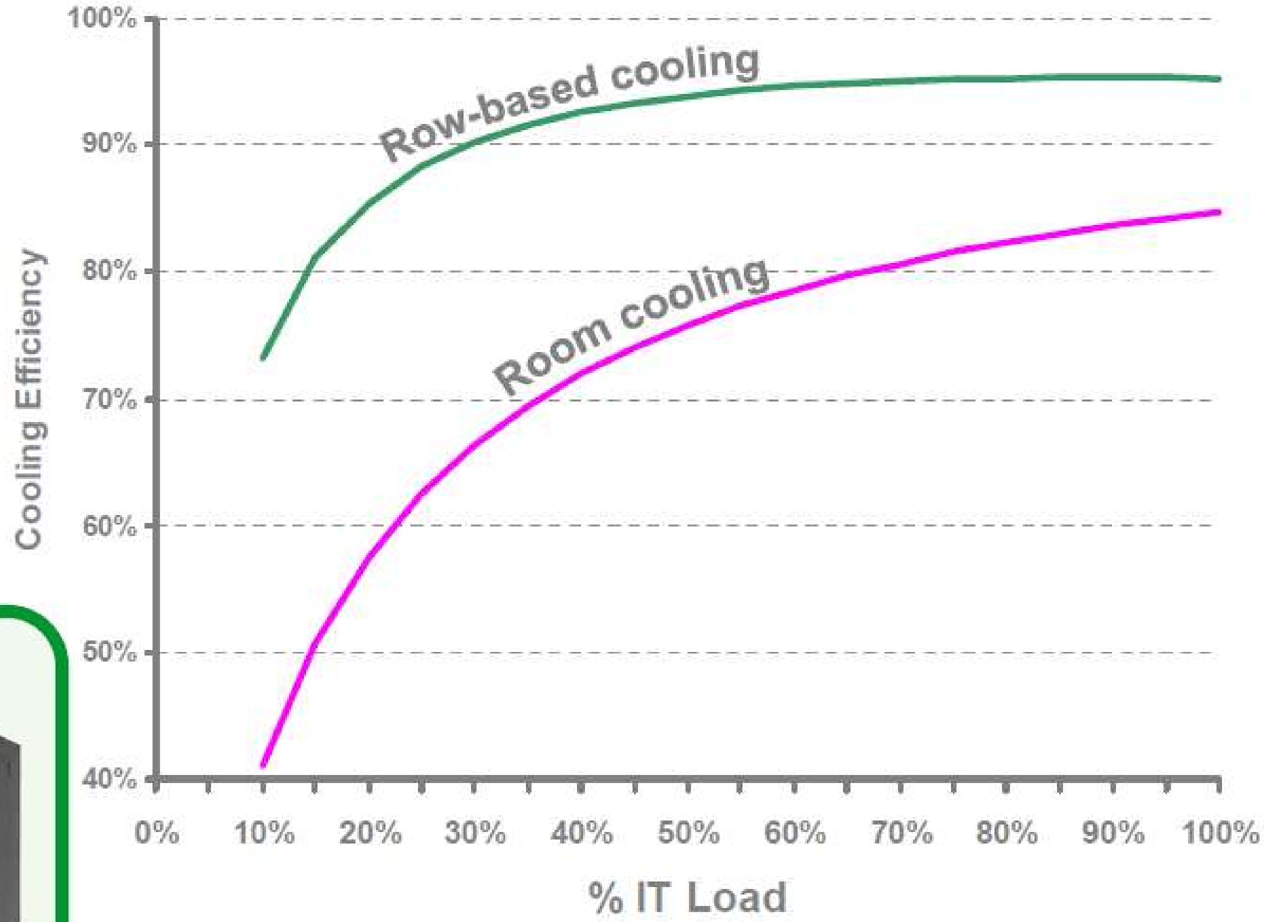

Figure 7

Computer room air conditioner efficiency curves comparing row-based cooling to traditional room cooling

> Row-based cooling units

Compared with the traditional room-oriented approach, the airflow paths of row-based air conditioners are shorter and much more predictable. In addition, all of the rated capacity of the air conditioner can be utilized, and higher power density can be achieved. At the same time, the usable capacity of the perimeter (room-based) cooling system increases and in some cases its cooling redundancy is restored to the original design as IT load is removed from this system and placed into the zone.

Although not discussed in this paper, row-based cooling is also an effective method for entirely cooling small low-density data rooms (1-3 rows of racks).

Row based cooling

Figure 1 – *High-density "hot spots" that vary both in power density AND location can result from dynamic virtualized IT loads*

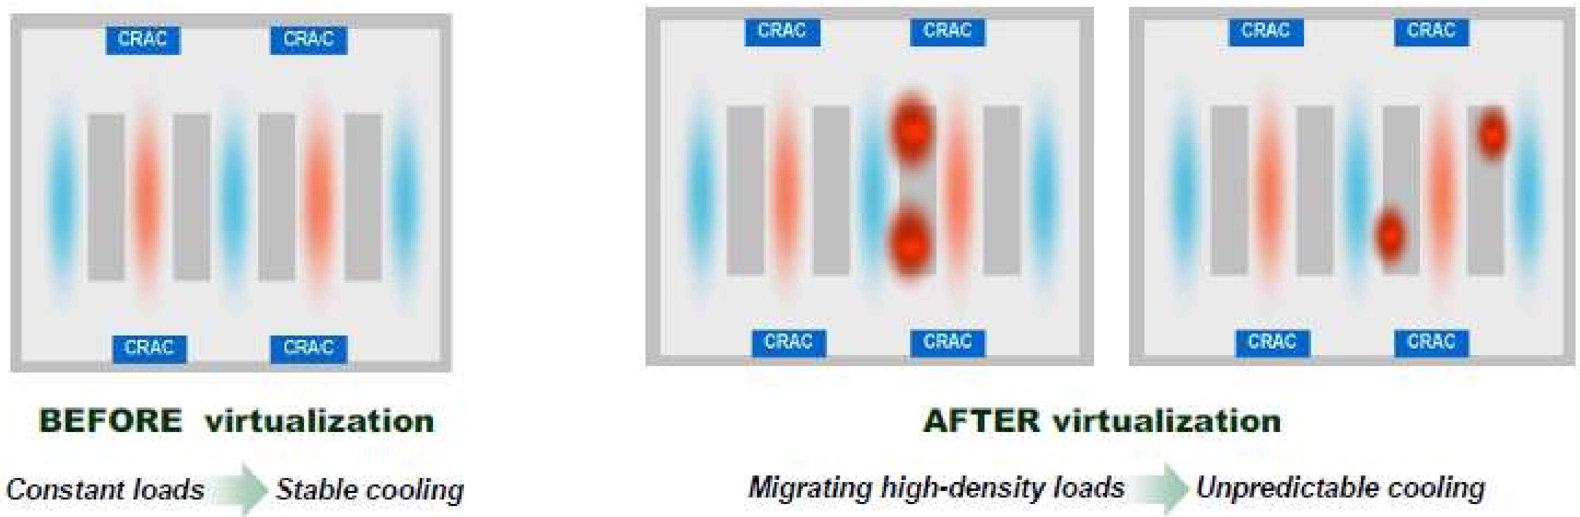

Predictable and efficient cooling calls for a system that comprehends these variations and automatically matches cooling – both in location and in amount – to changing power densities. The key characteristics of such a cooling system are:

- Short air path between cooling and load
- Dynamic response to load changes

Cooling units located within the rows – and instrumented to sense and respond to temperature changes – meet the above two essential criteria. Row-based cooling substantially increases efficiency by providing cooling only *where* needed, only *when* needed, and only in the *amount* needed (**Figure 2**).

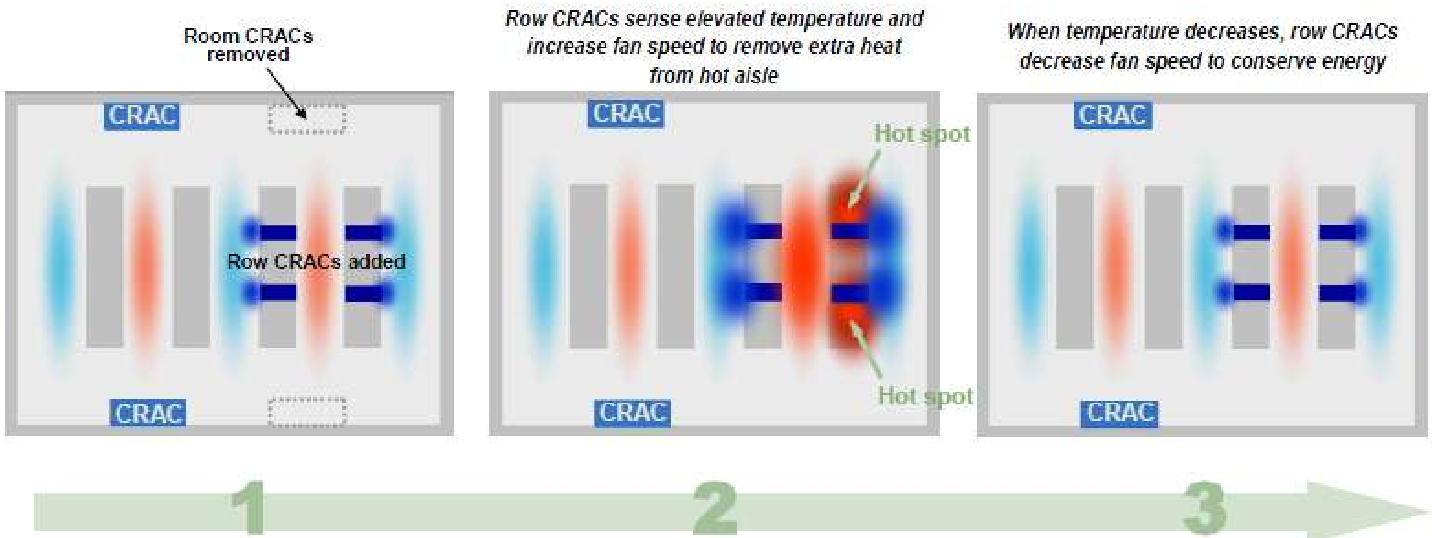

Figure 2

Row-based CRACs work together to remove extra heat from hot aisle

The placement of cooling units close to the servers provides the essential element that is the key to efficient cooling: **short air paths**. A short air path between cooling and the load enables a number efficiency and availability benefits:

- Reduced mixing of cold supply air with hot return air
- Increased return temperature (increases rate of heat transfer to coil)
- Targeted cooling that can respond to localized demand
- Conservation of fan power
- Reduced – often eliminated – need for make-up humidification (to replace condensation formed on a too-cold coil resulting from a too-low set point)

The benefits of variable, short-air-path cooling are only part of the advantage of row-based over room-based cooling. Other significant benefits arise from the fact that row-based cooling is **modular** and **scalable**, which addresses the second challenge when virtualizing: increasing efficiency by deploying correctly sized power and cooling capacity, discussed in the next section.

Redundancy

- Redundancy is necessary in cooling systems to permit maintenance of live systems and to ensure the survival of the data center mission if an air conditioning device fails.
- Power systems often use dual path feeds to IT systems to assure redundancy. This is because the power cords and connections themselves represent a potential single point of failure.
- **In the case of cooling, N+1 designs are common instead of dual path approaches because the common air distribution paths**, being simply open air around the rack, have a very low probability of failure.

Redundancy for rack-oriented architecture

- There is no sharing of cooling between racks, and no common distribution path for air. Therefore, the only way to achieve redundancy is to provide **a full 2N dual path CRAC system for each rack: essentially 2 CRAC systems per rack**.
- This is a severe penalty when compared with the alternative approaches. However, for isolated high density racks this is very effective as the redundancy is completely determined and predictable and independent of any other CRAC systems.

Redundancy for room-oriented architecture

- In principle, redundancy to be provided by introducing a single additional CRAC, independent of the size of the room. This is the case for very low densities, and gives this approach a cost advantage at low densities. However, at higher densities the ability of a particular CRAC to make up for the loss of another is strongly affected by room geometry. For example, the air distribution pattern of a specific CRAC cannot be replaced by a backup CRAC unit that is remotely located from the failed unit.
- The number of additional CRAC units that are required to establish redundancy increases from the single additional unit required at low densities to a doubling of CRAC units at densities greater than 10 kW per rack.

Redundancy for row-oriented architecture

- **Row-oriented architecture provides redundancy at the row level**. **This requires an additional or N+1 CRAC unit for each row**. Even though the row CRAC units are smaller and less expensive than room units, **this is a significant penalty at light loads of 1-2 kW per rack. However, for higher density this penalty is eliminated and the N+1 approach is sustained up to 25 kW per rack**. This is a major advantage when compared with either room or rack-oriented designs, which both trend to 2N at higher densities. The ability to deliver redundancy in high density situations with fewer additional CRAC units is a key benefit of the row-oriented architecture and provides it a significant total cost of ownership (TCO) advantage.

High-Density Zones in a Low-Density Data Center (APC white paper #134)

- High-density equipment such as **blade servers, 1U servers, and multi-core high-end servers** provide more computing per watt compared to previous generation servers. However, when consolidated, this new generation of equipment requires concentrated power and cooling resources. Data center operators and IT executives are often uncertain about the capability of their existing data center and whether a new data center must be built to support higher rack densities. Fortunately, a simple solution exists that allows for the rapid deployment of high-density racks within a traditional low-density data center. A high-density zone, as illustrated in the Figure, allows data center managers to support a mixed-density data center environment for a fraction of the cost of building an entire new data center.

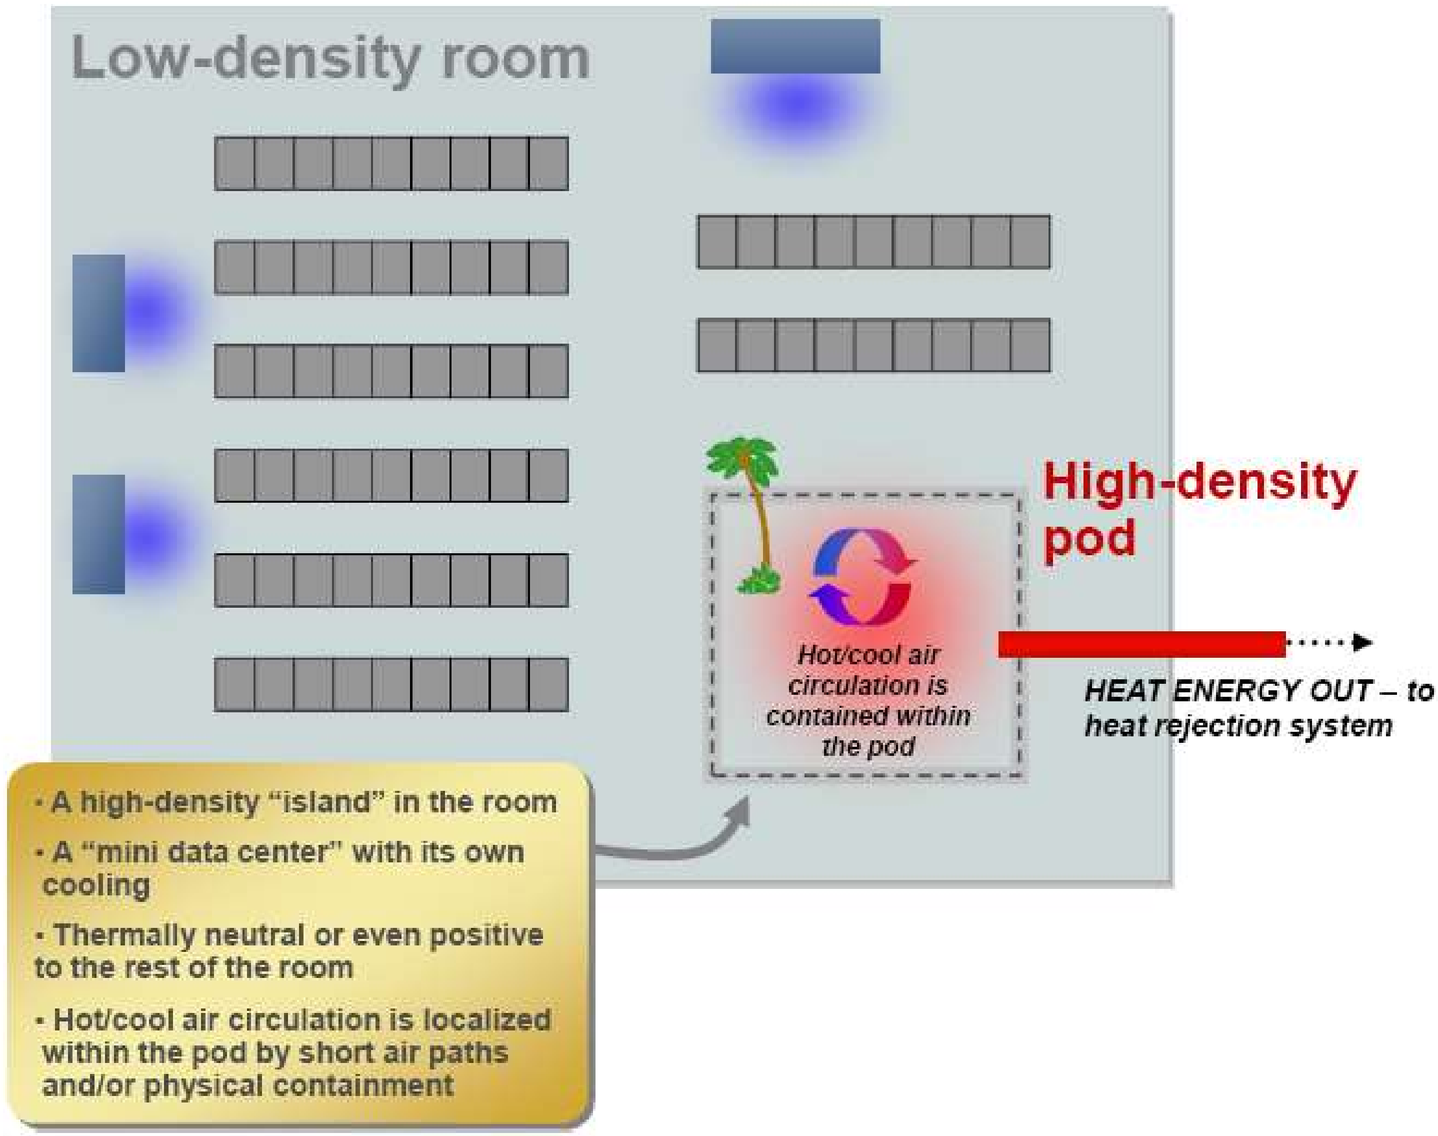
Low-density room
High-density pod
Hot/cool air circulation is contained within the pod
HEAT ENERGY OUT – to heat rejection system
• A high-density "island" in the room
• A "mini data center" with its own cooling
• Thermally neutral or even positive to the rest of the room
• Hot/cool air circulation is localized within the pod by short air paths and/or physical containment

The problem: unmanaged high density

- Traditional data center design uses a raised floor to distribute cooling to low-density IT equipment (Figure 2a). However, when high-density equipment is randomly installed throughout a low-density data center the cooling stability is upset and hot spots begin to appear (Figure 2b).

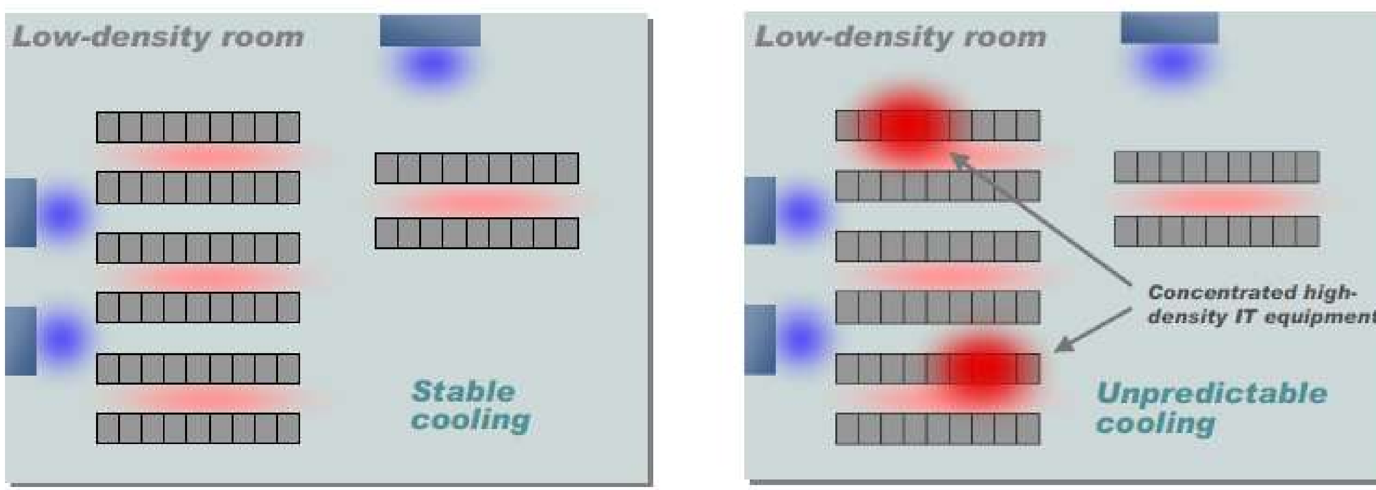

Figures 2a (left) and 2b

2a – Low-density data center

2b – High-density hot spots

The solution: high density zones

- This "drop in" solution eliminates the hot spots in Figure 2b by simply moving high-density equipment into the zone. The heat generated from the high-density IT equipment within this zone is rejected to the outdoors with no impact to the existing data center cooling system or the surrounding low-density IT racks.

Figure 3

Isolated (room-neutral) high-density zone

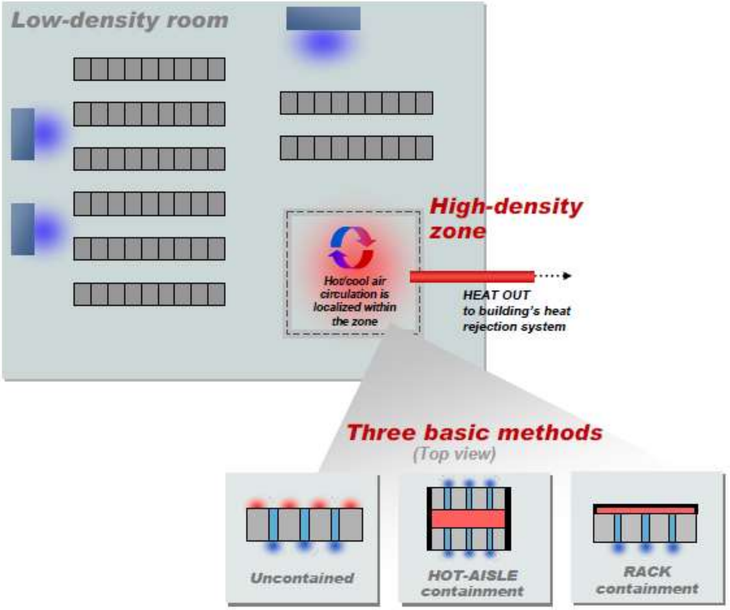

> High density enables high efficiency

In traditional data centers with room-based power and cooling, **unmanaged** high-density racks can cause destabilizing effects such as cooling inefficiency, loss of cooling redundancy, hot spots, thermal shutdown, and circuit overload.

However, with today's new power and cooling technologies, high-density racks offer an opportunity for dramatically increased efficiency and predictability, if deployed effectively and supported by "smart" row-based power and cooling.

The high-density zones described in this paper provide a way to deploy high density while at the same time achieving **increased overall data center efficiency** by targeted, scalable, localized power and cooling.

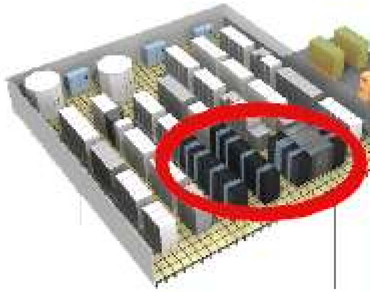

A high-density "island" in the room

A "mini data center" with its own cooling

Thermally "invisible" to the rest of the room (ideally)

Hot/cool air circulation is localized within the zone by short air paths and/or physical containment

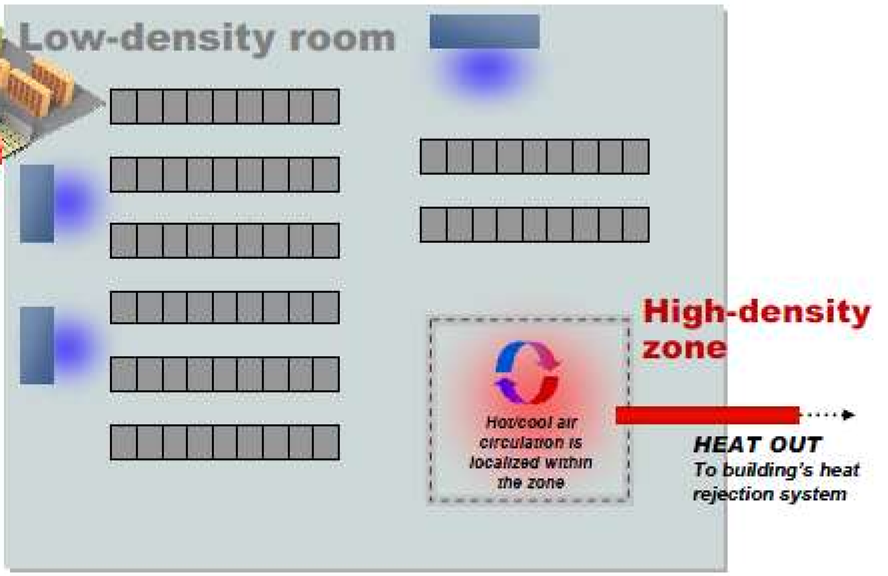

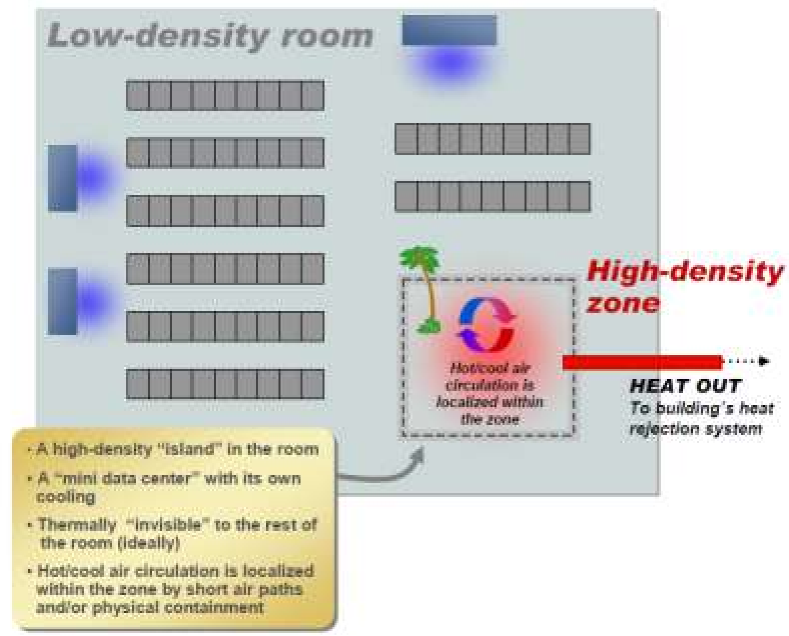

Figure 1

Basic concept of high-density zone

- The system in Figure 4 integrates a cluster of high-density IT racks with a high-density row based cooling system and high-density UPS and power distribution system in a remanufactured, pre-tested zone.

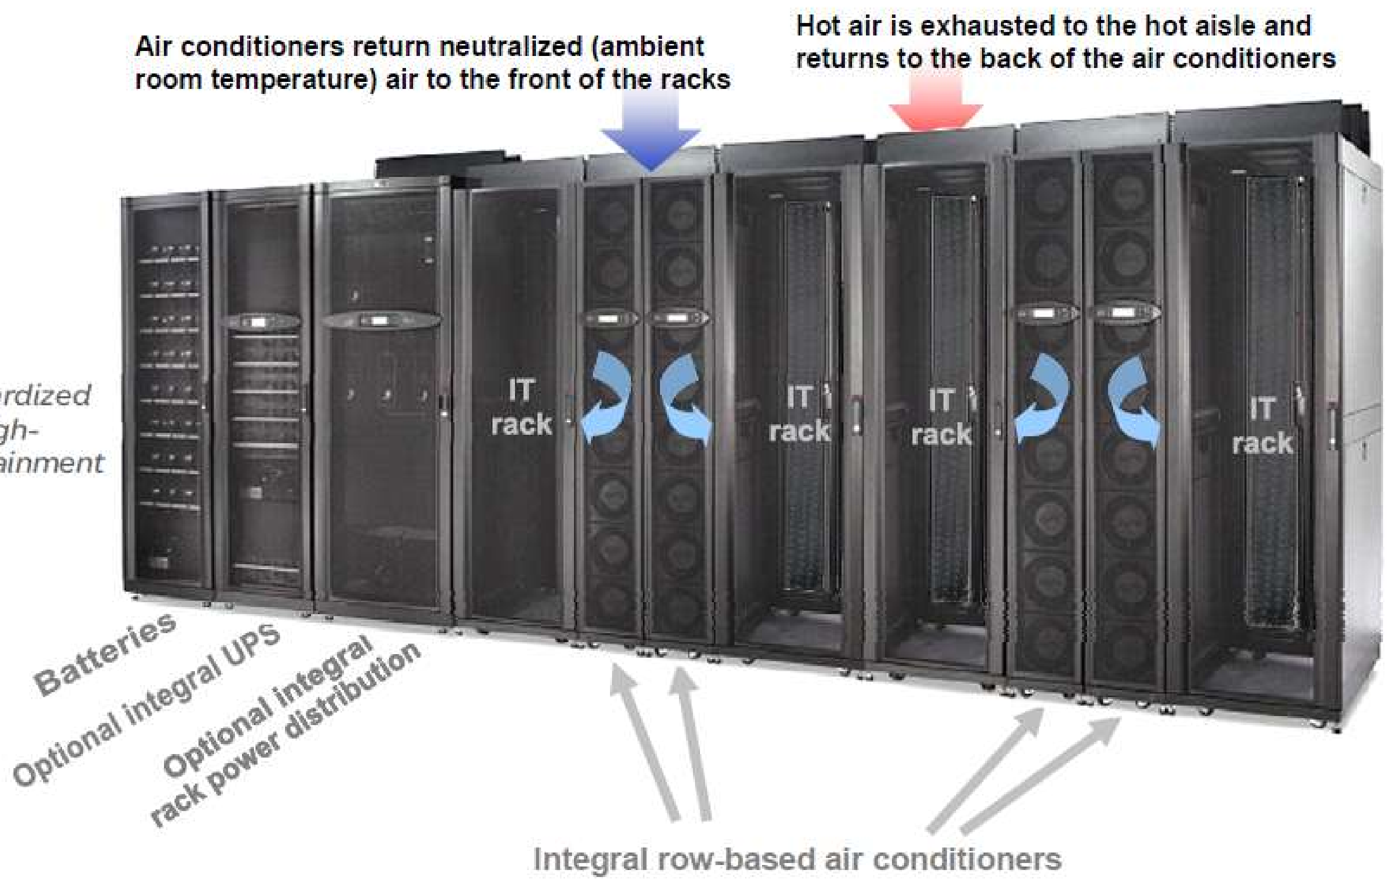

Figure 4

Front-view of a standardized modular multi-rack high-density zone (no containment in this example)

High-density zone containment methods

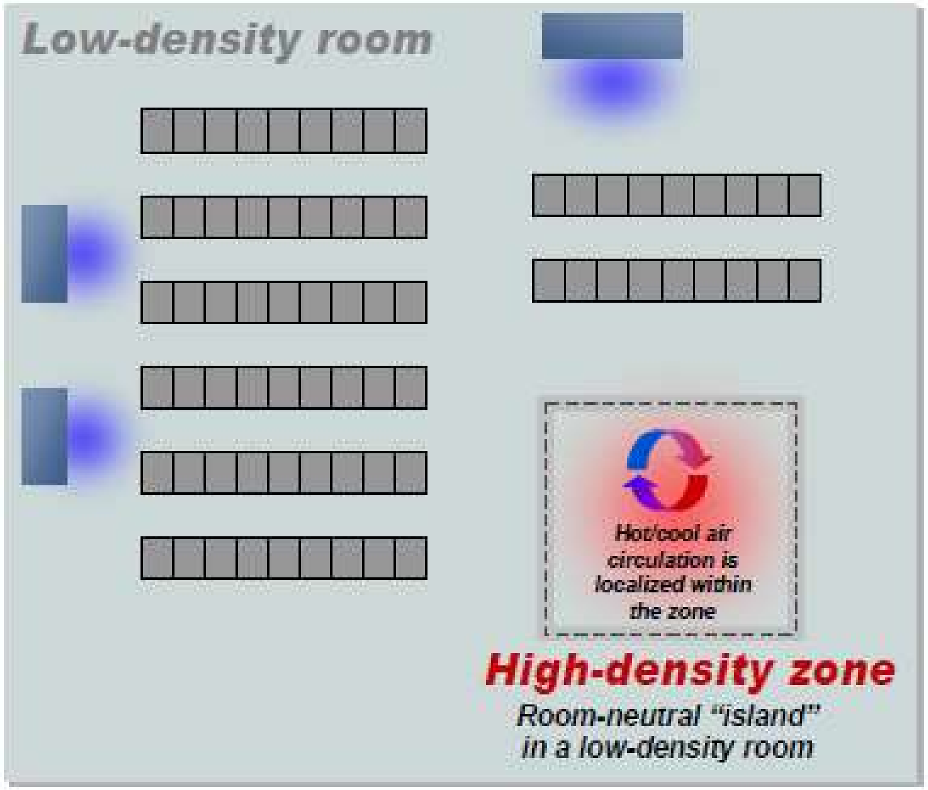

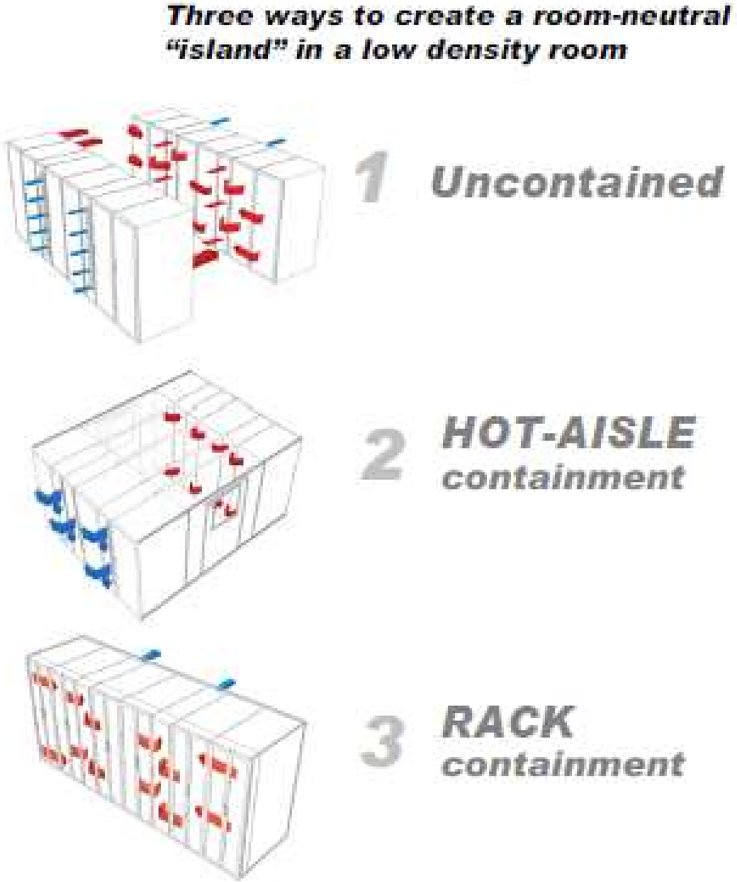

Uncontained zones

- Rely on the standard layout and widths of the common hot aisle and cold aisle arrangement to keep hot and cold air streams from mixing. For this reason, uncontained zones depend on multiple racks in a row and are not effective in cooling stand-alone IT racks. The hot and cold aisles formed by rows of racks (and in some cases walls) are what isolate the hot and cold air streams as illustrated in Figure 6.

Figure 6

*High-density zone with **no containment***

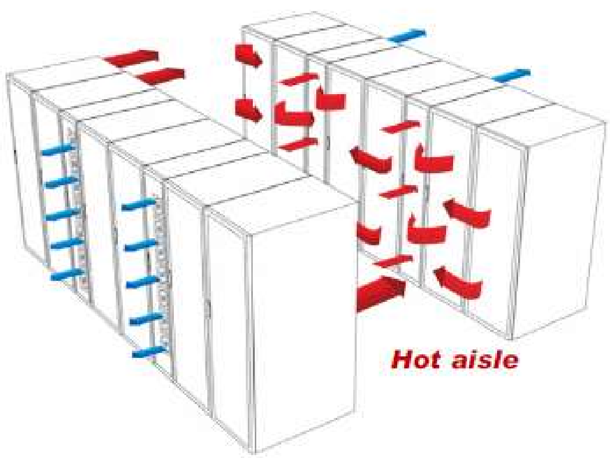

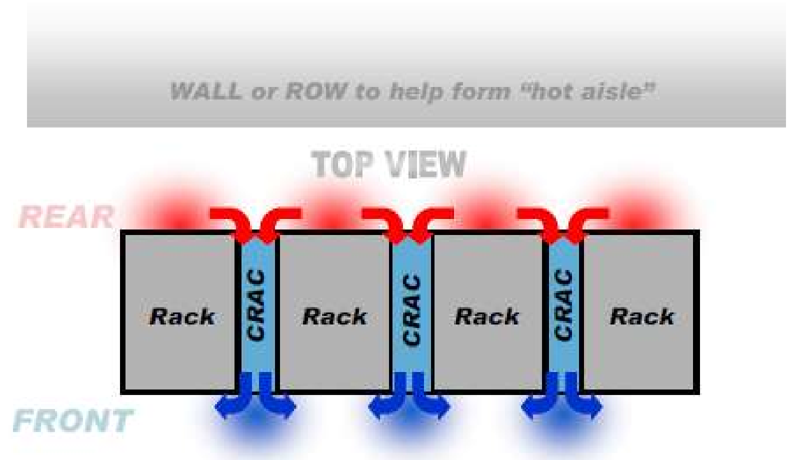

- When to use this method:
  - When IT racks designated for the zone are moved and relocated frequently
  - When IT racks are used form a variety of different vendors
- Trade-offs:
  - More row-based air conditioners required at lower densities in order to properly capture hot exhaust air from all IT racks.

APC White Paper #135
Hot-aisle vs. Cold-aisle Containment

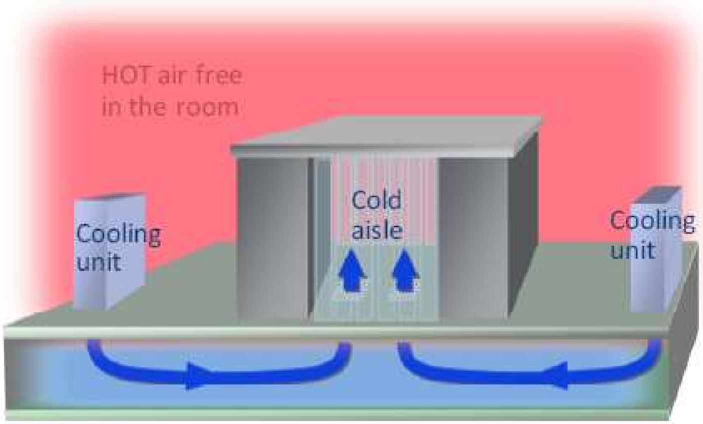

Figure 1

Cold-aisle containment system (CACS) deployed with a room-based cooling approach

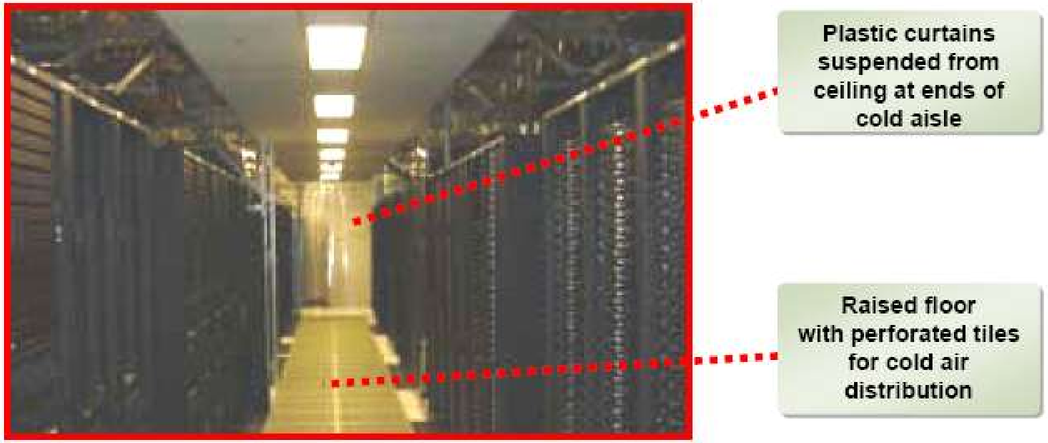

Figure 2

Example of a "homegrown" cold-aisle containment system

Figure 3

Hot-aisle containment system (HACS) deployed with row-based cooling

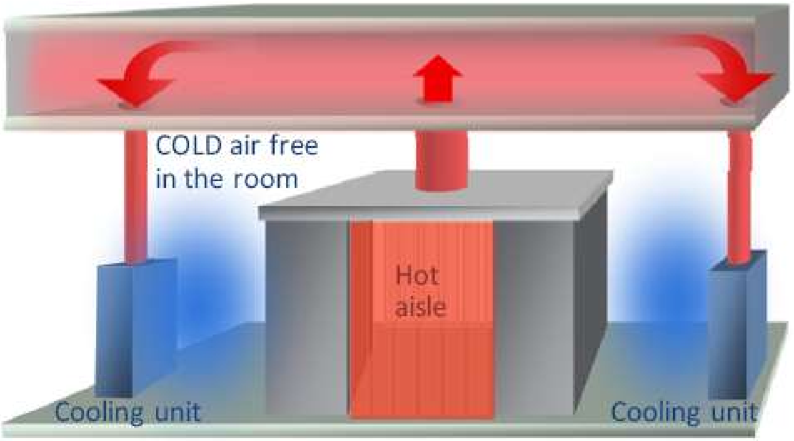

Figure 4

Example of a hot-aisle containment system (HACS) operating as an independent zone

Figure 5

Hot-aisle containment system (HACS) ducted to a remote air conditioner

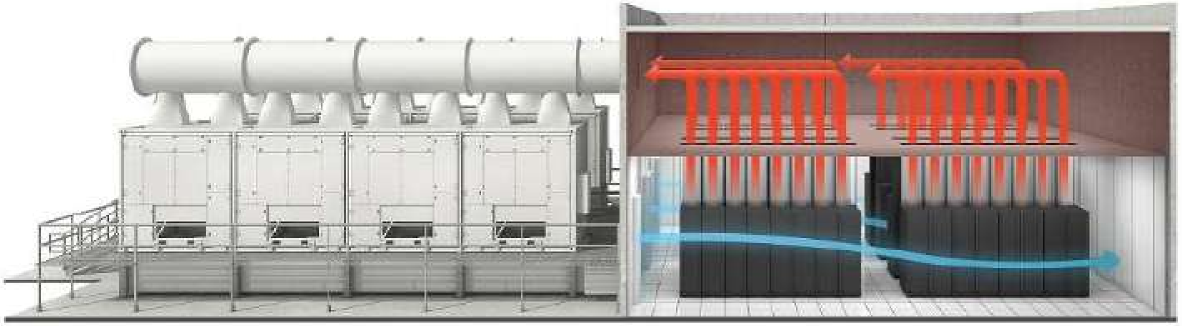

Benefits of Containment

APC white paper #135

- **Cooling systems can be set to a higher supply temperature (thereby saving energy and increasing cooling capacity) and still supply the load with safe operating temperatures.**
- **Elimination of hot spots.**
- **Economizer mode hours are increased.**
- **Humidification / dehumidification costs are reduced.**
- **Better overall physical infrastructure utilization, which enables right-sizing –which, in turn, results in equipment running at higher efficiencies.**

Table 1

Impact of controlling the uncontained area temperature for a CACS and HACS

Containment type	IT inlet air	Uncontained area DB	Uncontained area WBGT	Econ hours	m^3/s CFM[9]	PUE	Comments
Traditional uncontained	13-27°C 56-81°F	24°C 75°F	17°C 63°F	2,814	149%	1.84	Baseline with 49% cold and 20% hot-air leakage[10]
Scenario #1: IT inlet air temperature held constant at 27°C/80.6°F							
CACS Max ASHRAE IT inlet air temp and no limit on uncontained area temp	27°C 81°F	41°C 106°F	28°C 83°F	6,218	100%	1.65	WBGT only 2°C/3°F below OSHA max regulations. Includes 37% reduction in chiller power consumption. This is due to the increased IT supply temperature which allows for an increased CW supply temperature.
HACS Max ASHRAE IT inlet air temp and no limit on uncontained area temp	27°C 81°F	27°C 81°F	21°C 70°F	6,218	100%	1.65	WBGT 9°C/16°F below OSHA max regulations. Includes 37% reduction in chiller power consumption with increased CW supply temperature. *Note the hot-aisle temperature is 41°C/106°F.
Scenario #2: Temperature in uncontained area held constant at 27°C/80.6°F							
CACS 27°C /80.6°F max uncontained area temp	13°C 56°F	27°C 81°F	18°C 64°F	2,075	100%	1.86	Complies with OSHA, and complies with ASHRAE. Includes 5% increase in chiller power consumption. This is due to the decreased IT supply temperature which leads to a decreased CW supply temperature.
HACS 27°C /80.6°F max uncontained area temp	27°C 81°F	27°C 81°F	21°C 70°F	6,218	100%	1.65	Same results as HACS in Scenario #1.
Scenario #3: Temperature in uncontained area held constant at 24°C/75°F							
CACS 24°C /75°F max uncontained area temp	10°C 50°F	24°C 75°F	15°C 59°F	0	100%	1.98	Acceptable work environment but worse efficiency than baseline data center in first row. Includes 15% increase in chiller power consumption. This is due to the decreased IT supply temperature which leads to a decreased CW supply temperature.
HACS 24°C /75°F max uncontained area temp	24°C 75°F	24°C 75°F	18°C 65°F	5,319	100%	1.69	Higher efficiency, complies with OSHA, and complies with ASHRAE. Includes 28% reduction in chiller power consumption with increased CW supply temperature. *Note the hot-aisle temperature is 38°C/100°F.

[9] Total airflow (stated as % of IT airflow)

[10] Hot-air leakage occurs when hot exhaust air from servers mixes with the raised floor supply air, which increases server inlet temperature. Cold-air leakage occurs when cold air from gaps/voids in the raised floor mixes with return air, lowering return temperature and decreasing the cooling unit's efficiency.

Hot aisle containment

- Hot aisle containment zones are identical to uncontained zones except for the fact that the hot aisle in every pair of rows is contained.

Figure 7

High-density zone with ***hot aisle containment***

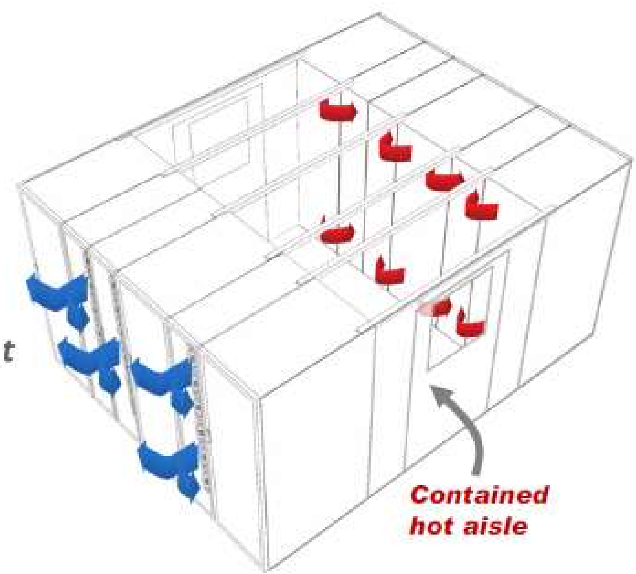

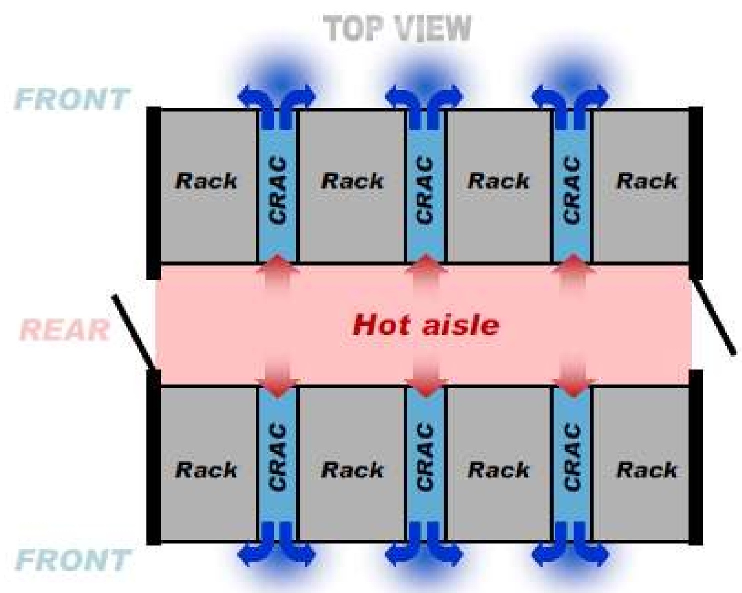

***Figure 6** – Hot aisle containment system (HACS) operating as an independent zone*

Hot aisle containment system (HACS) efficiency benefits

Efficiency – The efficiency of the HACS will be higher because the hot aisle is capable of maintaining higher temperatures. In a typical high density server environment the temperature difference between the server exhaust air and the room temperature is typically around 30° F / 17° C). If the room is maintained within ASHRAE TC9.9 standards at 72° F / 22° C, a 30° F / 17° C temperature difference, would yield a server exhaust air temperature of 102° F / 39°F.

In a typical HACS environment the cooling units tend to move slightly more air than the servers and draw in a small amount of room air into the hot aisle. The effect of this can cause a slight reduction (about 2° F) in the return temperature, yielding a return temperature to cooling units of 100° F / 38°C. The net effect of this elevated return temperature (i.e., 100° F / 38° C) to the cooling unit enables better heat exchange across the cooling coil, better utilization of the cooling equipment, and overall higher efficiency. **Figure 7** gives an example of the effect of elevated return temperatures on sensible cooling capacities (the ability of an air conditioning system to remove heat from the air).

The effect of increasing return temperatures on cooling unit capacity holds true for virtually all air conditioning equipment. Some equipment may have limits as to the maximum return temperature they can handle, but, in general, all cooling systems will yield higher capacities with warmer return air.

In the case of the HACS, hot aisles operating at 100° F / 38° C with high density servers is typical. Contrast this with a cold aisle contained room where the entire room space would have to be maintained at 100° F / 38° C in order to achieve the same level of efficiency. While CACS would enable higher return temperatures, the typical data center operator will not operate the entire data center room at 100 F / 38° C in order to achieve the same efficiency as a HACS.

***Figure 7** – Effect of increased return temperature on sensible cooling capacity**

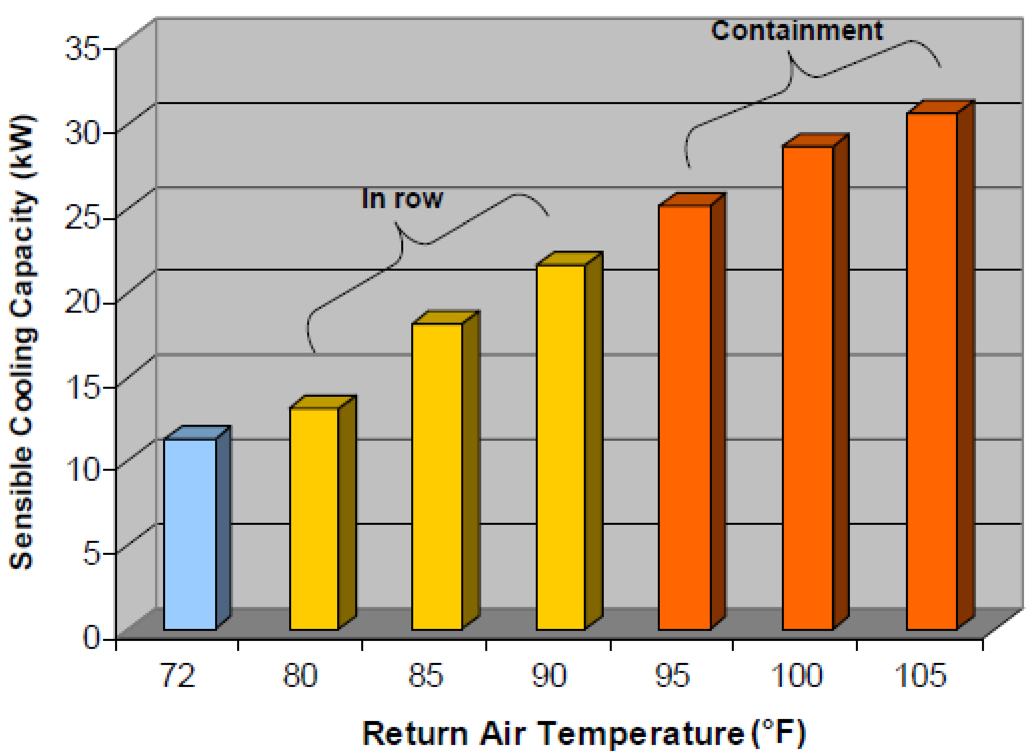

* APC In-row RC model air conditioner

Improved Flexibility – Unlike CACS, a HACS does not impact the temperature of the surrounding room. A HACS is, in effect, a room neutral solution. For example, if the temperature of the data center is set for 75° F (24° C), and a CACS system is implemented, the room temperatures outside of the cold aisle will rise because hot air will mix in with the air outside of the cold aisle on its way to the intake of the cooling system. The hot air inside of the HACS is contained from the rest of the room. The HACS does not deliver any hot air to the outside room; therefore the existing cooling system is not rendered less efficient.

A HACS can be "dropped in" to the data center without requiring any changes to the existing data center cooling architecture. When utilizing a row-based cooling approach (as opposed to a room-based approach), no need exists for the installation of specialized duct work and no adjustments need to be made to the existing HVAC systems to handle elevated return temperatures.

Higher Availability – The "Cold Air Volume Sample Calculation" in **Figure 5** demonstrates the differences in cold air volume when comparing CACS volume to room volume (uncontained cold air volume is 17 times greater than cold air in a contained cold aisle). This difference has a significant impact on the ability of the systems to support a cooling failure (i.e., runtime). A runtime that could be minutes in an uncontained room scenario might only be seconds if a CACS approach is deployed. With HACS only the hot air is contained, leaving the rest of the data center environment cool. Therefore, the servers will draw air from a larger pool of cool air outside the contained hot aisle, thereby extending available runtime.

Cold Aisle Containment

The Cold Aisle Containment System (CACS) is typically deployed in traditional perimeter-based cooling environments. Traditional cooling environments use the entire room as a hot air return plenum and use deliver cold air via the raised floor plenum to the cold aisles. The CACS encloses the cold aisle allowing the rest of the data center to become a large hot air return plenum. By containing the cold aisle, the hot / cold air streams within the data center are separated.

***Figure 3** – Example of a home grown cold aisle containment system*

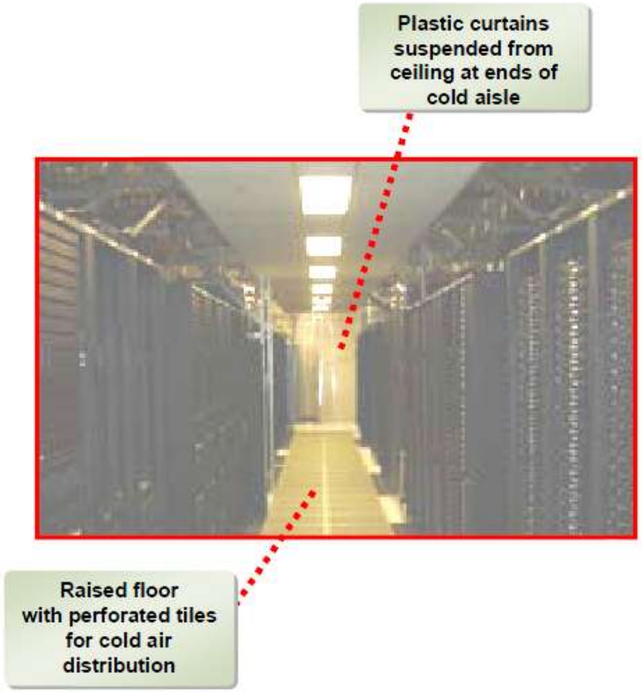

Figure 2 – Cold aisle containment (CACS) deployed with a perimeter cooling approach

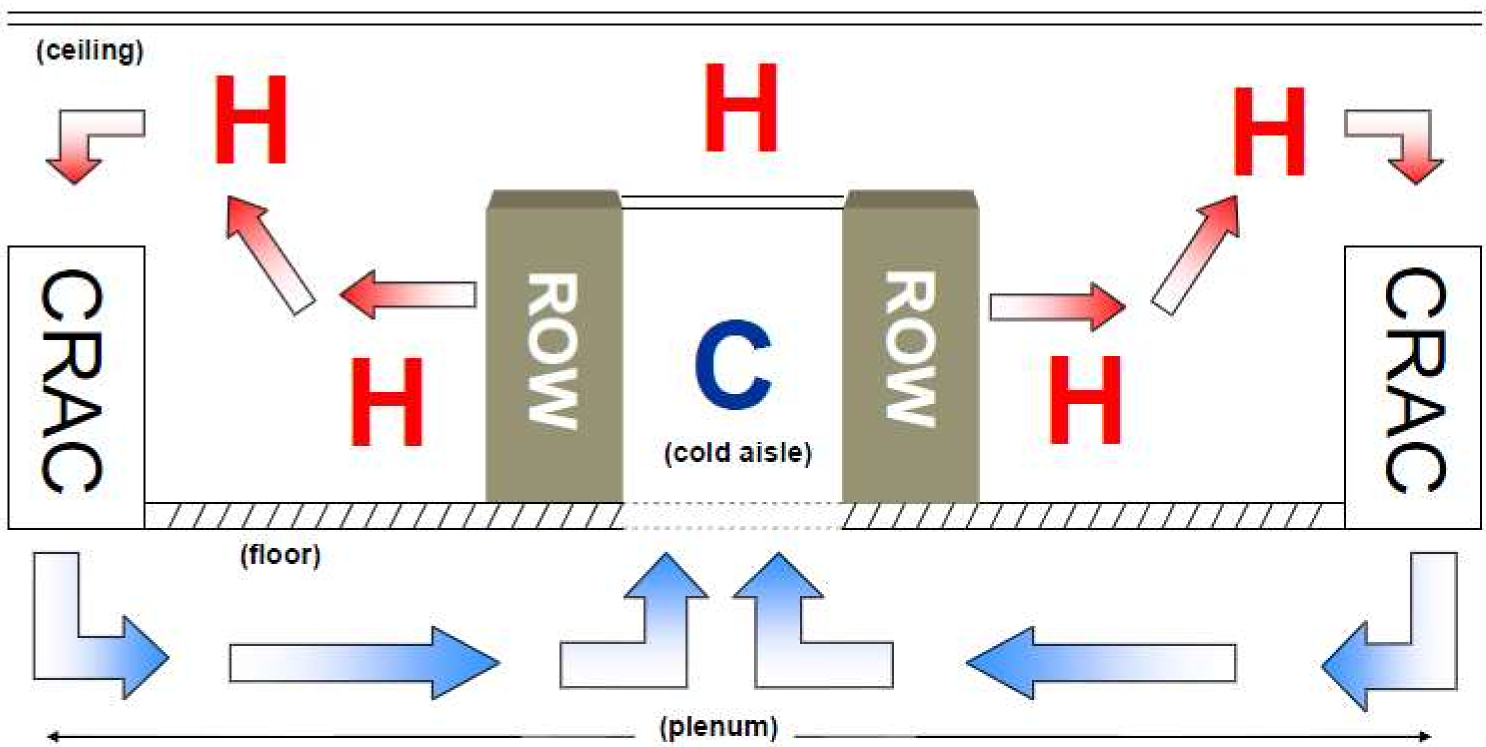

Cold Aisle Containment – efficiency limitation in Room based approach

- **Inefficiencies resulting from distances and pressures required for adequate air distribution -** The single largest contributor to inefficiency in a room-based approach is the requirement to move cold air from a perimeter CRAC unit to a distant load. A row-based cooling approach brings the source of the cooling in close proximity to the load. As a result, much less energy is required to deliver the cold air to its destination. This is not the case if the data center owner chooses to deploy CACS with row-based cooling. See APC White Paper # 130, "The Advantages of Row and Rack-Oriented Cooling Architectures for Data Centers" for a detailed comparison of row and room based cooling approaches.

- **Density limitations of using cold air distribution through raised floor -** The practical density limit when using a CACS approach is approximately 6 kW per rack. See APC White Paper # 46, "Cooling Strategies for Ultra High Density Racks and Blade Servers" for details regarding the reasons for this limitation. Higher densities can be achieved only if an investment is made in a customized design. To address some of the limitations of the raised floor plenum, some CACS solutions are offered with fan powered floor tiles. This improves airflow for higher density racks. The use of additional fan assisted devices further reduces the efficiency of the CACS. The extra fans contribute to the overall power consumption and add heat to the supplied cold air. The efficiency gains achieved by the CACS are therefore diminished by the added requirement for floor-based fans. This density limitation can be avoided if the data center owner chooses to deploy CACS with row-based cooling.

Cold Aisle Containment –

efficiency limitation in Room based approach

- **Predictability of the raised floor -** Cold aisle containment helps improve predictability through the elimination of hot and cold air mixing. However, it does not eliminate the variable of the raised floor. Cabling, piping, and other obstructions are added below the raised floor as the data center evolves. These obstructions limit the delivery of sufficient cool air to the IT equipment. **Figure 4** provides an example of how air dams in the plenum under the raised floor can hinder the

***Figure 4** – Raised floors become congested as data center requirements change*

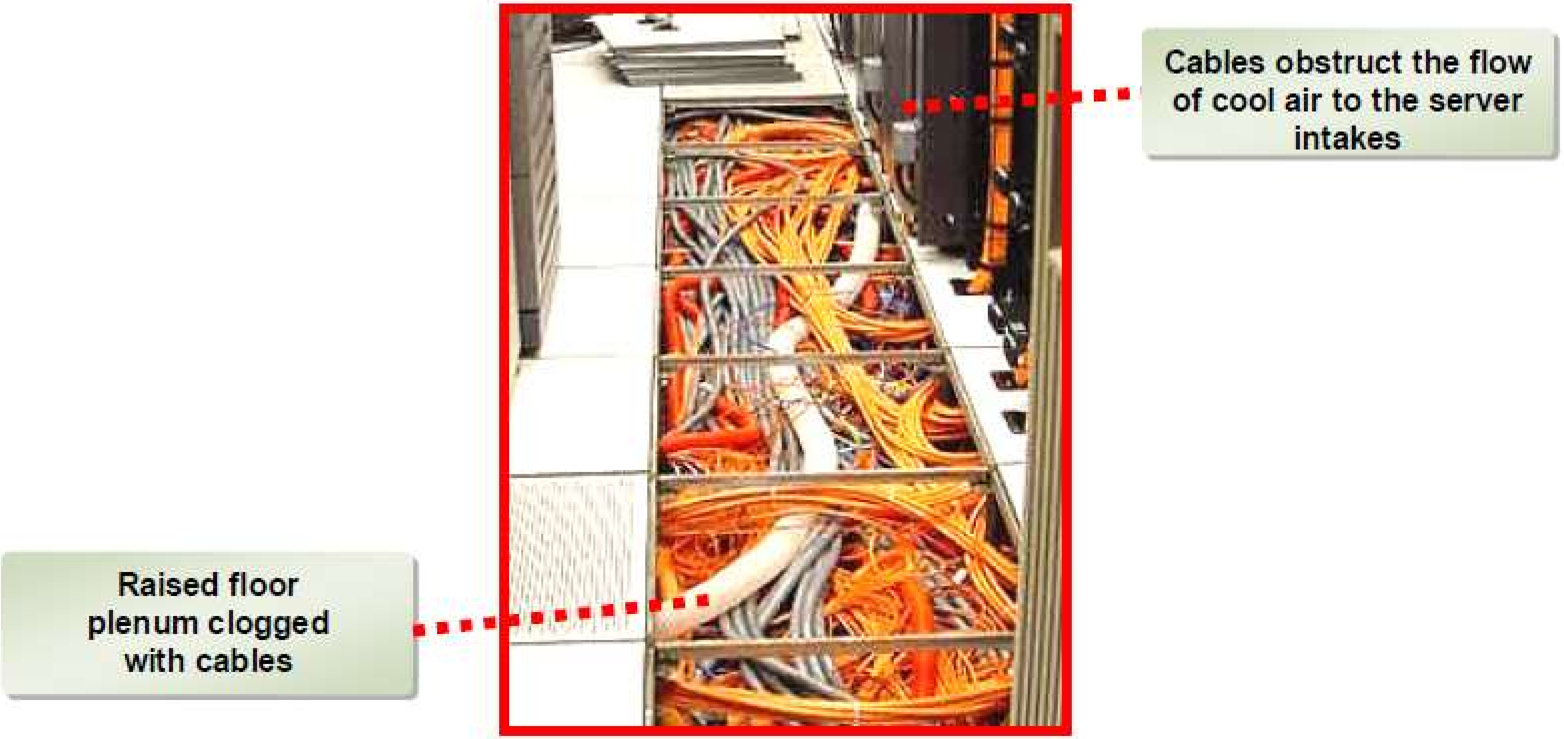

Cold Aisle Containment

efficiency limitation in row based approach

- **Availability of cold air during a loss of power / cooling -** Containing the cold aisle minimizes the overall pool of cold air available to the servers, should a loss of power and / or cooling occur. The reduced volume of cold air results in more rapid temperature increases in the event of a failure. **Figure 5** depicts a sample data center and compares the volume of air in a contained cold aisle to the volume of air in an uncontained cold aisle. The **uncontained** cold aisle shows a volume of cold air that is 17 times greater than that of the cold air volume found in the **contained** cold aisle. This reduced air volume shortens the amount of time (seconds instead of minutes) it would take for the servers to overheat if a failure were to occur.

***Figure 5** – Comparison of cold air volumes in contained and uncontained cold aisles*

Cold Air Volume Sample Calculation

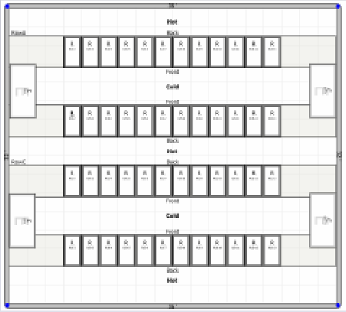

Room Dimensions:

- 36 ft(11m) x 31 ft (9.4M) x 10(3m) ft
- Cold Aisle Width: 4 ft (1.2m)
- Hot Aisle Width: 3 ft (0.9m)
- Rack Height: 42U – 6.5 ft (1.99M)
- Rack Width: 1.97 ft (0.6M)

Contained Cold Aisle Volume =

4ft x (1.97 ft x (12 racks per row)) x 6.5 ft = 614.6 ft^3 (17.2 M^3)

Room Volume (without hot aisles) =

(36 ft x 31ft x 10 ft) – (3 ft x (1.97 x12) x 10) = 10,450.8 ft^3 (3,185.4 M^3)

Uncontained cold air volume is **17 times** greater than cold air in contained cold aisle scenario [614.6 ft^3 (17.2 M^3) x 17 = 10,450.8 ft^3 (3,185.4 M^3)]

Cold Aisle Containment – efficiency limitation in row based approach

- **All of the cold aisles in the entire data center must be contained in order to realize benefits -** Containing only some of the cold aisles in the data center will yield little benefit because any other cold air that is allowed to mix with hot air will diminish any expected savings. Mixing will cause the cooling system to operate in a less efficient manner (a smaller difference between return "hot" air and cooling coil temperatures). To minimize mixing and to maximize cooling system efficiency, all cold aisles must be contained. Only then will hot return air temperatures reach maximum potential thereby allowing the cooling equipment to operate at much higher efficiency levels.

- **Overall perception and operation of a hot data center –** ASHRAE Standard TC9.9 recommends that server inlet temperatures range from 68-77° F (20-25° C). When cold aisles are contained, the air in the rest of the room becomes hotter (well above 80° F / 27° C and in some cases as high as 100° F / 38° C), and anyone entering the data center is exposed to unusually high temperatures. People are generally alarmed when entering such hot conditions, and tours become impractical. People's expectations need to be adjusted so that they understand that the higher temperatures are "normal" and not a sign of impending system breakdown. This cultural change can be challenging for workers not accustomed to entering a data center operating at higher temperatures.

  When operating a data center at elevated temperatures, special provisions need to be made for non-racked IT equipment. This is equipment that cannot be integrated into a CACS. Since, with a CACS system, the room is a reservoir for hot air, miscellaneous devices (such as tape libraries and standalone servers) will need to have unique ducting in order to enable them to pull cold air from the contained cold aisles. In addition, electric outlets, lighting, fire suppression, and other systems within the room will need to be evaluated for suitability of operations at elevated temperatures.

Table 1 – *Cold aisle containment vs. hot aisle containment summary*

Characteristic	Cold Aisle Containment	Hot Aisle Containment	Comment
Efficiency improvements	Yes	Yes	HACS is more efficient than CACS because HACS typically will operate at higher return temperatures due to isolation of the hot air from the rest of the room.
Ability to increase cold air supply set point without impacting entire data center	No	Yes	With HACS cooling set points can be set higher while still maintaining a comfortable work environment. Increasing CACS cooling set points results in uncomfortably high data center temperatures.
Leverages maximum number of potential free cooling days	No	Yes	By increasing cooling set point containment systems allow for increased free cooling. However, Increasing the set point of CACS results in increased room temperatures which is undesirable from a free cooling days perspective.
Room neutral solution	No	Yes	A HACS deployment is a "drop in" solution. CACS impacts the surrounding data center infrastructure.
Ease of deployment with room cooling	Yes	No	CACS is preferred when using room level cooling with a free return system which draws its return air from the room. A HACS without in-row cooling would require special return duct work or ceiling plenum.
Ability to scale for high density	No	Yes	CACS is often implemented with raised floor and inefficient fan assisted floor tiles are needed in order to achieve higher density.
Room neutral design	No	Yes	HACS is room neutral – it doesn't impact the outside room temperature in any way. CACS makes the air outside of the contained rows hotter.
Adverse temperature impact on non-racked equipment	Yes	No	With CACS, because the cold aisles are contained, the rest of the data center is allowed to become hot. Equipment outside of contained areas would have to be evaluated for operation at elevated temperatures.

> Why shorter air paths increase data center efficiency

With traditional room-based perimeter cooling, cool supply air must travel further to each rack, requiring extra fan power. In addition, if there is no raised floor there can be substantial mixing of the cool supply air with warm exhaust air in the room, which will require lowering the supply temperature far below what is needed at the IT racks. Hot spots can force lowering the supply temperature even further to control overheating. Too-low supply temperature can risk condensation on the cooling coil, which results in wasteful dehumidification-rehumidification and reduction of system cooling capacity. Long return paths similarly can cause air mixing, which lowers the return temperature to the cooling unit. Lower return temperature slows the rate of heat transfer to the coil, so heat is removed less efficiently.

The much shorter air paths in row-based cooling dramatically lessen mixing of supply and return air (and with containment and blanking panels, virtually eliminate mixing). On the supply side, this allows operation at a higher coil temperature, which takes less chiller energy to maintain and is much less likely to cause wasteful condensation. On the return side, it produces a higher return temperature which increases the heat removal rate. Compared to long-path cooling systems, these combined short-path effects (1) increase operational cooling efficiency and (2) increase the cooling capacity of the heat rejection system.

- When to use this method:
  - In cases where floor space must be conserved. This method is popular because it consumes the same space as two rows of low-density racks.
  - In data centers with hot aisle / cold-aisle layouts
- Trade-offs:
  - Hot aisle containment panels increase capital cost
  - Hot aisle containment may exceed work environment policies due to high temperature
  - Incompatible with some types of cabling, power strips, labels, and other materials that are not rated for high temperatures
  - Not possible with a single row of racks
  - Authority having jurisdiction (AHJ) may require fire suppression in hot aisle

Rack containment

- Rack containment (also called rack air containment) is similar to hot aisle containment except that the hot exhaust air is contained using the back frame of the equipment racks and a series of panels to form a rear air channel.

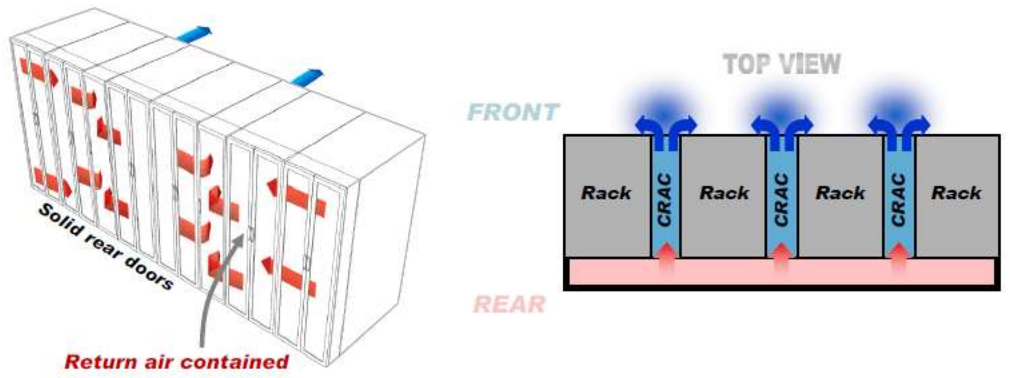

Figure 8

High-density zone with ***rack containment***

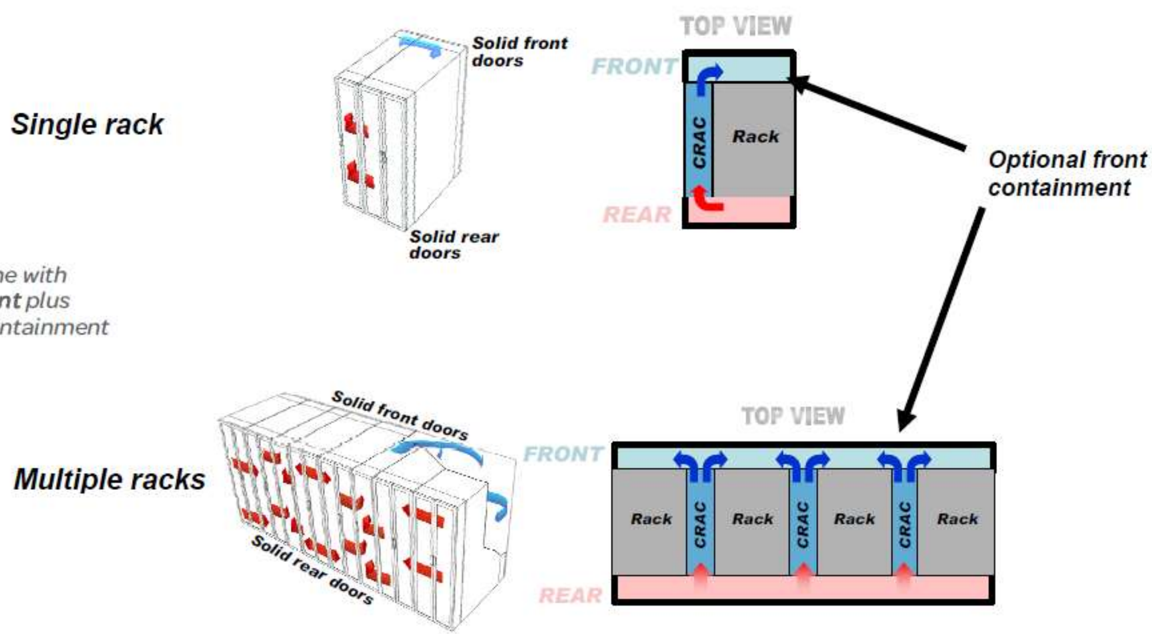

Figure 9

High-density zone with ***rack containment*** *plus optional front containment*

- When to use this method:
  - In cases where hot aisle containment is the preferred method, but a single odd row is left uncontained
  - When frequent access to and easy management of communication cables is required
  - For complete isolation in cases such as stand-alone open data center environments or mixed layouts – only when optional front containment is used
  - In wiring closets that lack any form of cooling, exposing high-density equipment to high temperatures – only when optional front containment is used
  - When sound attenuation is required – only when optional front containment is used
- Trade-offs:
  - Front and rear containment panels increase capital cost
  - In a single rack configuration, cost increases substantially when cooling redundancy is required

Additional high density zone benefits

- Standardization of design elements
- Compatibility with any data center, new or existing
- Configurability with dedicated UPS and power distribution
- Configurability with any level of redundancy
- Configurability with any number of IT racks

Table 2

Deploying high-density equipment: traditional vs. zoned approach

Attribute	Traditional approach	Zoned approach	Comments
Positioning the data center as a source of corporate competitive advantage	Difficult	Easier	Simple economics – cost of doing business is lower per unit of computational output
Just-in-time IT deployments	Very difficult	Easy	Deployments are highly dependent on modular and predictable power and cooling which affects management and ability to quickly deploy
Predictability of performance	High	Low	Strongly linked to data center infrastructure efficiency
Likelihood of hot spots	High	Very low	Management applications insure optimal placement of equipment in zones to prevent hot spots
Cooling efficiency	Poor	Excellent	Room based cooling units are oversized to overcome under-floor obstacles, distance, air mixing, demand fighting, etc.
Ability to plan	Poor	Excellent	Standardization / predictability facilitate "what-if" scenarios before moves, adds, and changes

Figure 11

Packaged standalone high-density zone

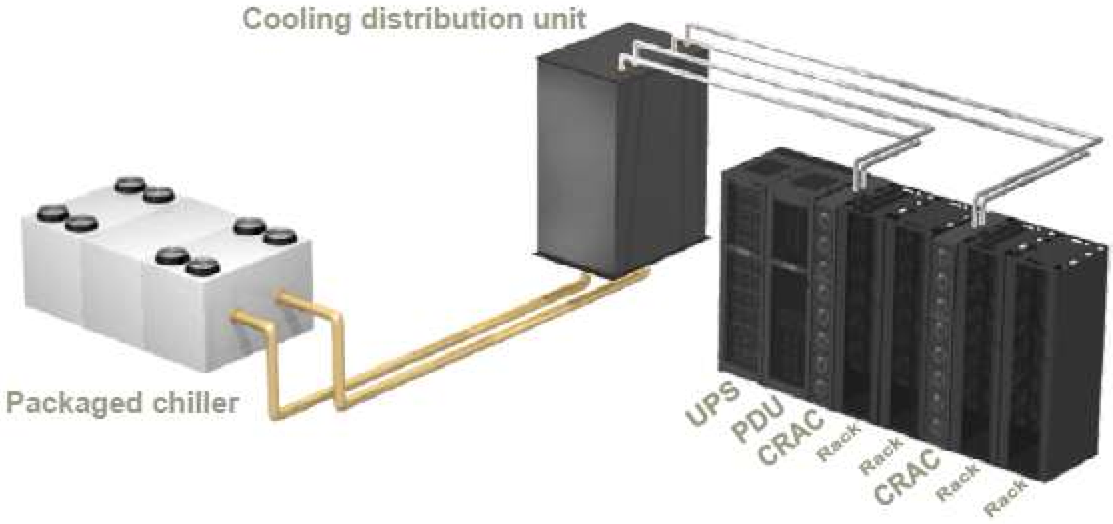

Containment Advantages

- Energy savings in servers.
- Energy savings in the facility chilling process
- Energy savings in the facility air handling units

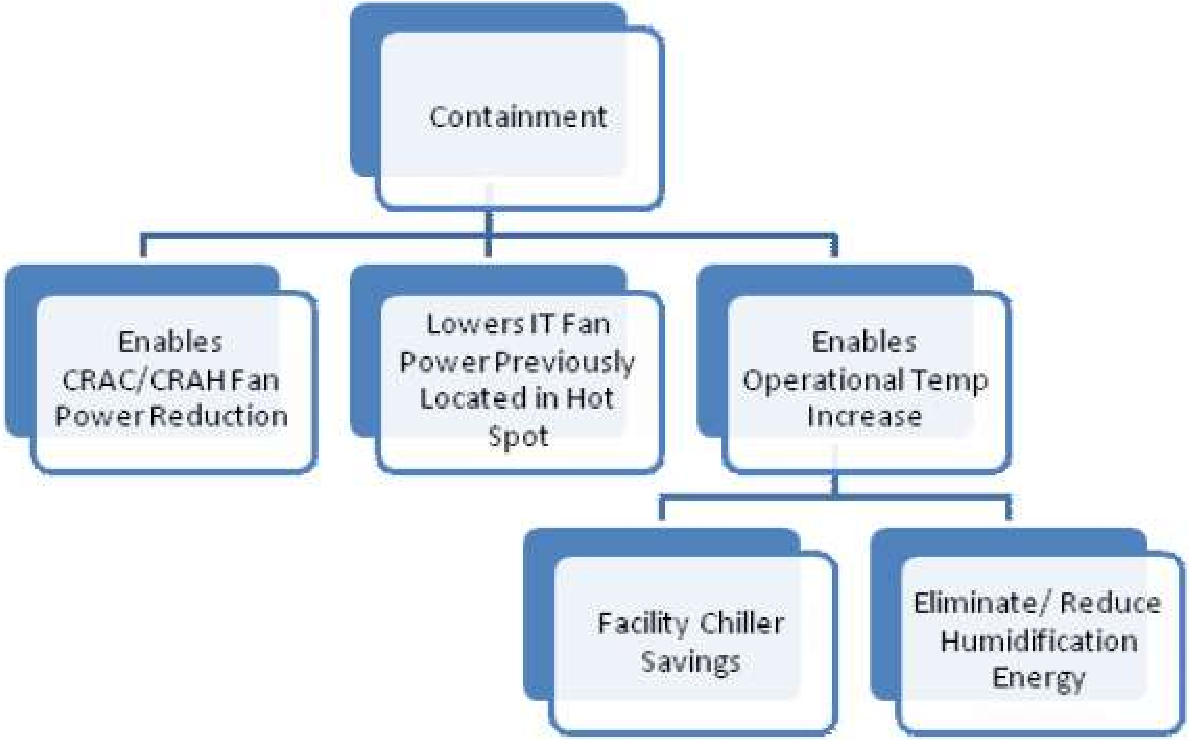

Table 1

Comparison of zone containment methods

Selection criteria	No containment	Hot aisle containment	Rack air containment	Comments
Minimize footprint	Good	Good	Moderate to poor	•NO containment and HOT AISLE containment provide minimum row spacing •RACK air containment adds 8 inches to the depth of the rack but may be acceptable in consolidation applications •Front AND rear containment adds 16 inches to the depth of the rack – should be weighed against available floor space
Ease of change management	Good	Moderate to poor	Moderate to poor	•Taking racks in and out of an existing row is more difficult when containment systems constrain the rack with hardware, especially with front containment
Minimize energy consumption	Moderate	Good	Good	•NO containment layout is closely linked to the existing data center layout which could increase the number of row-based units
Ease of redundancy	Moderate	Good	Moderate to poor	•HOT AISLE containment row-based CRAC positions are independent of redundancy •More row-based CRACs needed to maintain redundancy in rack containment
Minimize # of row-based CRACs (particularly at low density)	Poor to moderate	Good	Moderate to good	•RACK air containment and RACK air containment with front containment may be limited since not all rack air can be shared among all row-based coolers as with HOT AISLE containment •NO containment depends heavily on rack power density where high densities require less row-based coolers •RACK air containment and rack air containment with front containment highly-influenced by redundancy (more coolers needed)
Sound attention	Poor	Moderate to poor	Good	•Poor to moderate with RACK air containment only •Good when using RACK air containment with front containment •Will reduce the decibel level of the cooling equipment but will not completely eliminate the noise
Installation in thermally unstable or non-data center space	Poor	Poor	Good	•Poor to moderate with RACK air containment only •Good when using RACK air containment with front containment •Examples include wiring closets, offices, and commercial spaces
Cost	Dependent upon variables such as rack power density and number racks			•Although the hot aisle containment has additional panels that increase cost, it will require fewer row-based coolers than no containment, particularly at lower rack power densities

Row Layout Improvements

- Airflow path lengths are minimized to reduce fan power
- Airflow resistance is minimized to reduce fan power
- IT equipment exhaust air is returned directly at high temperature to the air conditioner to maximize heat transfer
- Air conditioners are located so that airflow capacities are balanced to the nearby load airflow requirements

Summary and analysis

- The modular rack-oriented architecture is the most flexible, fast to implement, and achieves extreme density, but at the cost of additional expense.
- Room-oriented architecture is inflexible, time consuming to implement, and performs poorly at higher density but has cost and simplicity advantages at lower density.
- The modular row-oriented architecture provides many of the flexibility, speed, and density advantages of the rack-oriented approach, but with a cost similar to the room-oriented architecture.

Usable Capacity

- Most users naturally assume that if they have 500 kW of cooling units installed, they can install and cool 500 kW of IT loads. This is simply not the case. While a group of air conditioning units taken together may have in total the claimed capacity, this does not mean that they are able to deliver this cooling to the load. The fraction of the actual capacity that can be obtained in the real world cooling IT loads is called the “usable capacity”.

Typical Usable Capacity

Figure 5 – Usable air conditioner capacity as a function of average rack power density for the three cooling architectures

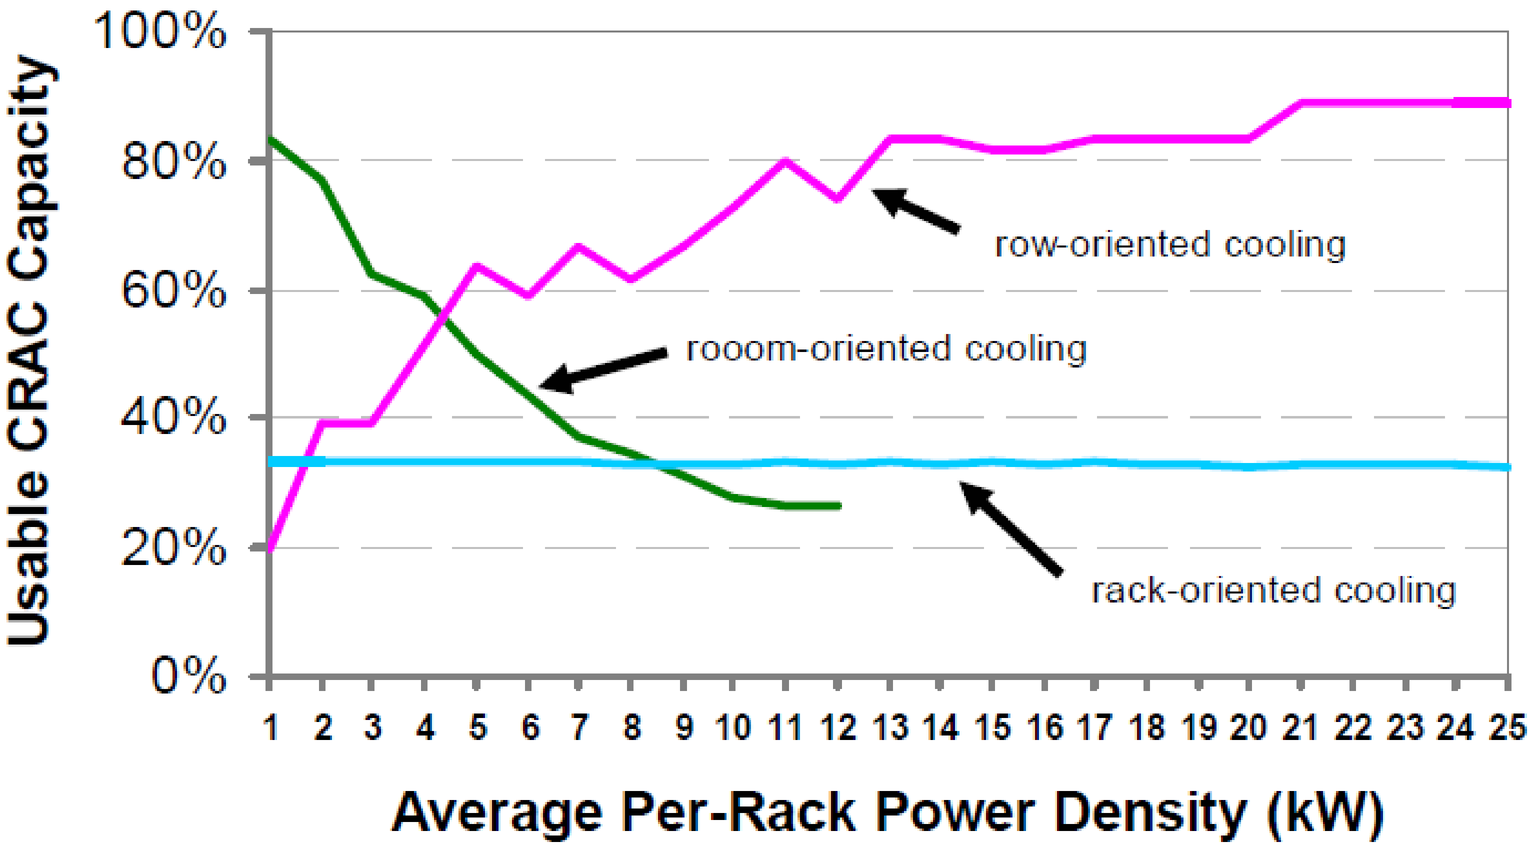

- The usable capacity in a room-oriented architecture appears on the surface to be 100%, because it appears that all the capacity is pooled and sharable at the room level. In fact, at very low power densities such as 1-2kW per rack, this is a reasonable assumption as shown in the Figure 5. However, this assumption breaks down quite dramatically as the power density increases. This loss of capacity is due to the inability of the system to deliver the required cool air to the load. The result is that the system must be oversized compared with the load, resulting in a reduction in the effective usable capacity. **The lack of predictability of the room oriented architecture creates a practical cutoff of around 6 kW per rack** as shown in Figure 5.

- **Row-oriented offers the highest usable capacity** across the broadest power density range. Due to the close coupling of the CRAC units to the load, all of the capacity can be delivered to the load up to power densities on the order of 25 kW, or approximately 4X the practical density capacity of room-oriented architecture. In addition, CRAC units can share cooling with nearby racks, which reduces the stranded capacity problem discussed earlier which is associated with rack-oriented architecture. However, **the usable capacity of row oriented architecture falls at very low power densities, because air conditioning units must be assigned to every row no matter how low the density becomes**. The unusual jagged nature of the usable capacity curve for the row-oriented architecture is due to quantization effects, due to finite row lengths combined with the need to assign CRAC units to specific rows and the lack of fractional sizes for the CRAC units.

- The usable capacity in a **rack-oriented** architecture is **typically significantly less than 100%.** In this architecture, **each rack has a dedicated air conditioner and therefore dedicated capacity**. Whenever the actual load in a rack is less than the rated capacity of that rack, the remainder of the capacity of that rack is not utilized, and furthermore cannot be utilized by any other rack. For example, **if a rack has 10 kW of cooling but only a 6 kW IT load, the rack has 4 kW of stranded capacity that cannot be used by any other rack**. This stranded capacity cannot be borrowed by neighboring racks for redundancy maintenance, or any other purpose. Since real-world racks vary significantly in power density, **usable capacity may be 50% or even lower of the rated capacity**.

Typical Electricity Usage

***Figure 6** – Annual CRAC electrical costs per megawatt of IT load as a function of average rack power density for the three cooling architectures*

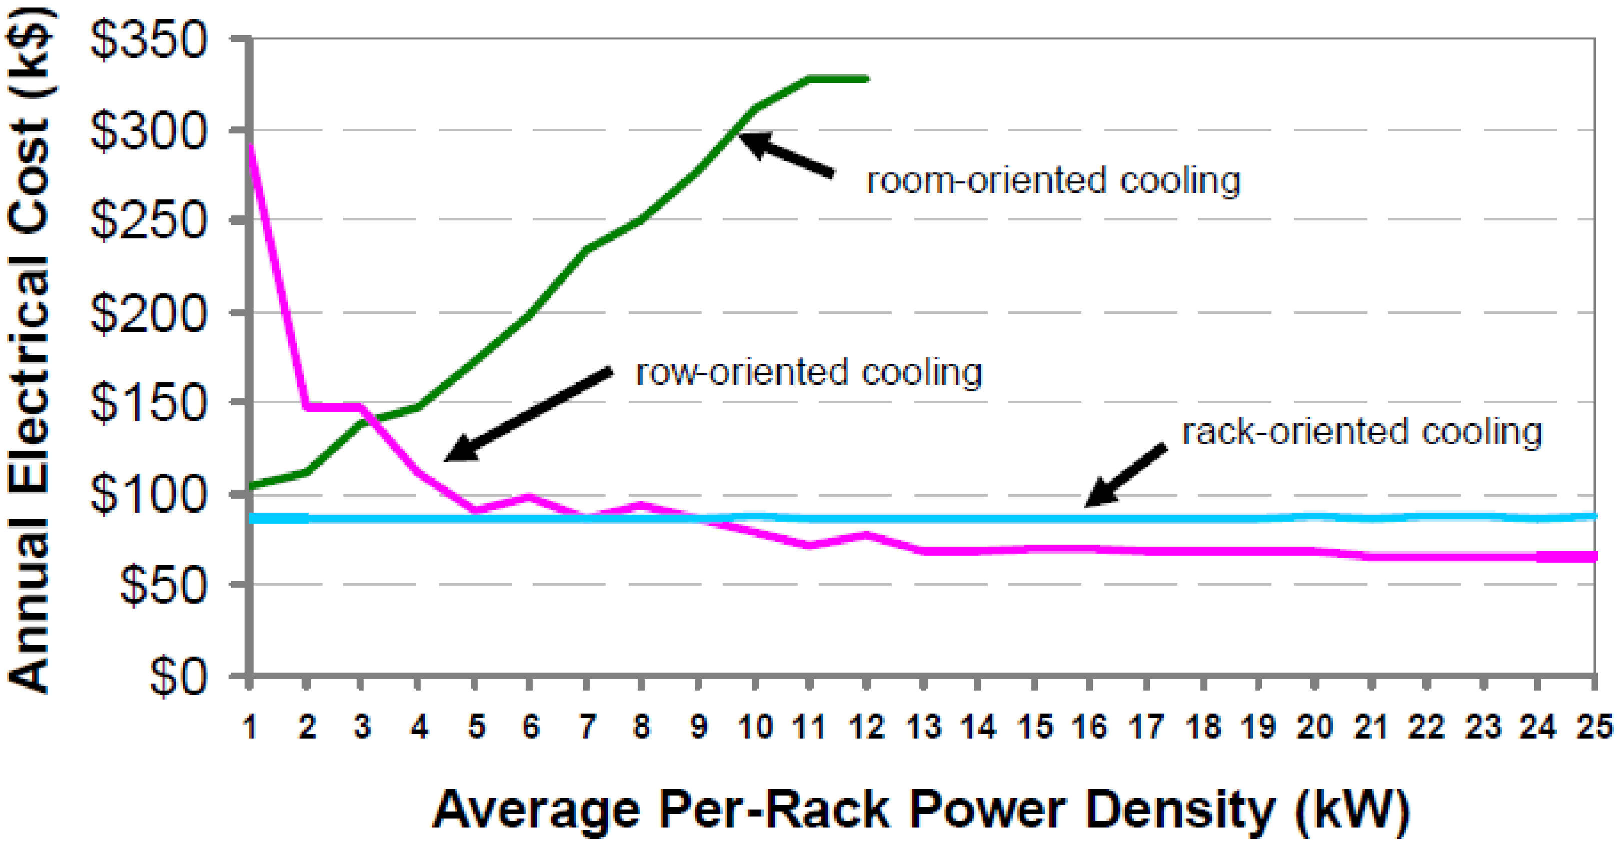

Five strategies for deployment of high density enclosures and blade servers

APC White paper #46

- 1. *Load spreading.*
- 2. *Rules-based borrowed cooling.*
- 3. *Supplemental cooling.*
- 4. *Dedicated high-density areas.*
- 5. *Whole-room cooling.*

- Cooling requirement Cooling an ultra-high density enclosure is a much more difficult problem than providing power. The blade server system described above would require approximately 2,500 cfm (cubic feet per minute) (1,180 L/s cool air at the intake (based on a common value of 20°F [11°C] exhaust air temperature rise) and would exhaust the same amount of heated air from the rear of the enclosure.

4 key elements in design blade server

- Supply **2,500 cfm** (1,180 L/s) of cool air to the enclosure
- Remove 2,500 cfm (1,180 L/s) of hot exhaust air from the enclosure
- Keep the hot exhaust air out of the intake
- Provide all of these functions in a redundant and uninterrupted manner

Figure 3 – Available rack enclosure cooling capacity of a floor tile as a function of per-tile airflow

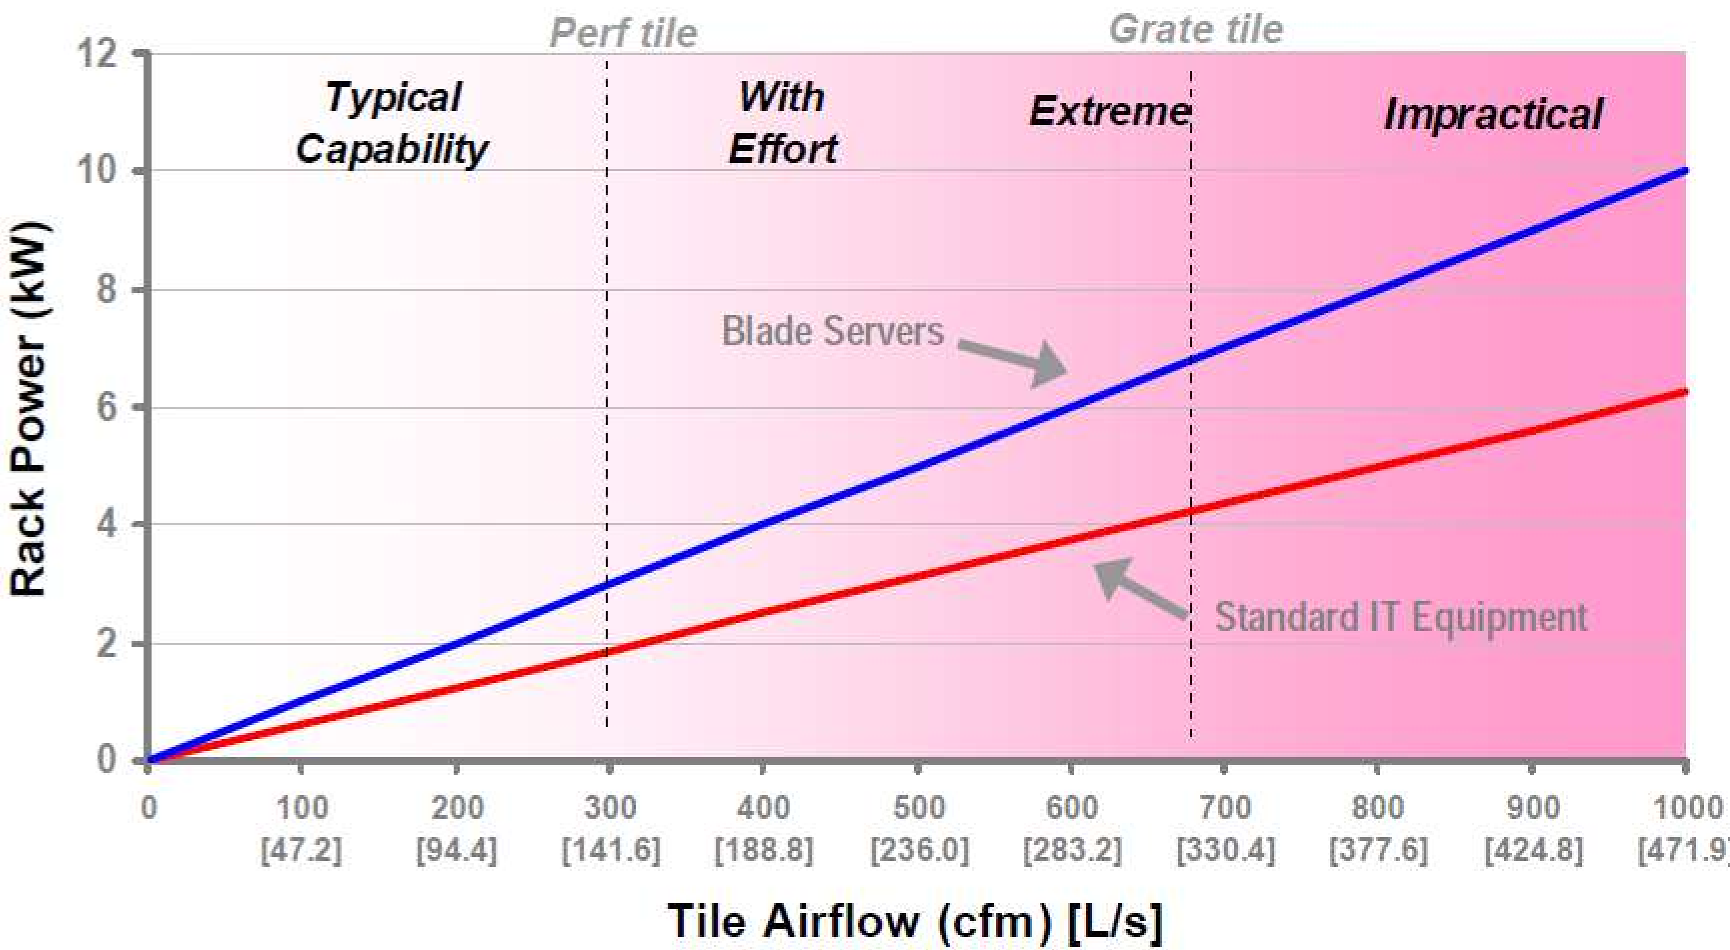

Method 1: Load spreading

- *Provide the room with the capability to power and cool to an average value below the peak potential enclosure value, and spread out the load of any proposed enclosures whose load exceeds the design average value by splitting the equipment among multiple rack enclosures.*
  - This is the most popular solution for incorporating high-density equipment into today's data centers. Fortunately, 1U servers and blade servers do not need to be installed closely spaced in the same enclosure, and can be spread out across multiple racks. By splitting equipment among racks no rack need ever exceed the design power density and consequently cooling performance is predictable.

Method 2: Rules-based borrowed cooling

- *Provide the room with the capability to power and cool to an average value below the peak enclosure value, and use rules to allow high density racks to borrow adjacent underutilized cooling capacity.*
  - This approach takes advantage of the fact that some racks draw less power than the average design value. The cooling delivery and return capacity that was available to the underutilized enclosures is available to other enclosures in the same vicinity. A simple rule like “do not locate high density racks near each other” has some beneficial effect, but more sophisticated rules can be put in place that can reliably and predictably permit enclosures to be cooled to over twice the average design value.

Method 3: Supplemental cooling

- *Provide the room with the capability to power and cool to an average value below the peak enclosure value, and use supplemental cooling equipment to cool racks with a density greater than the design average value.*
  - This solution normally requires that the installation be planned in advance to be able to utilize supplemental cooling equipment when and where needed. When a room has been planned in this way, a variety of techniques can be used to supplement rack cooling.

- A variety of techniques can be used to supplement rack cooling. These include:
  - Installation of specialty (grate-type) floor tiles or fans to boost the cool air supply from the CRAC to an enclosure
  - Installation of specialty return ducts or fans to scavenge the hot exhaust air from an enclosure for return to the CRAC
  - Installation of special racks or rack mounted cooling devices with the capability to provide the required cooling directly from the rack

Method 4: Dedicated high-density area

- *Provide the room with the capability to power and cool to an average value below the peak enclosure value provide a special limited area within density enclosures to that area.*
  - This approach requires prior knowledge of the fraction of high density enclosures, and the ability to segregate those enclosures into a special area, and under these constraints can achieve optimal space utilization.

Figure 7 – Example of a modular power and cooling system for a dedicated high density area within a data center. Modules of 2-12 IT racks, rated at 20 kW per rack.

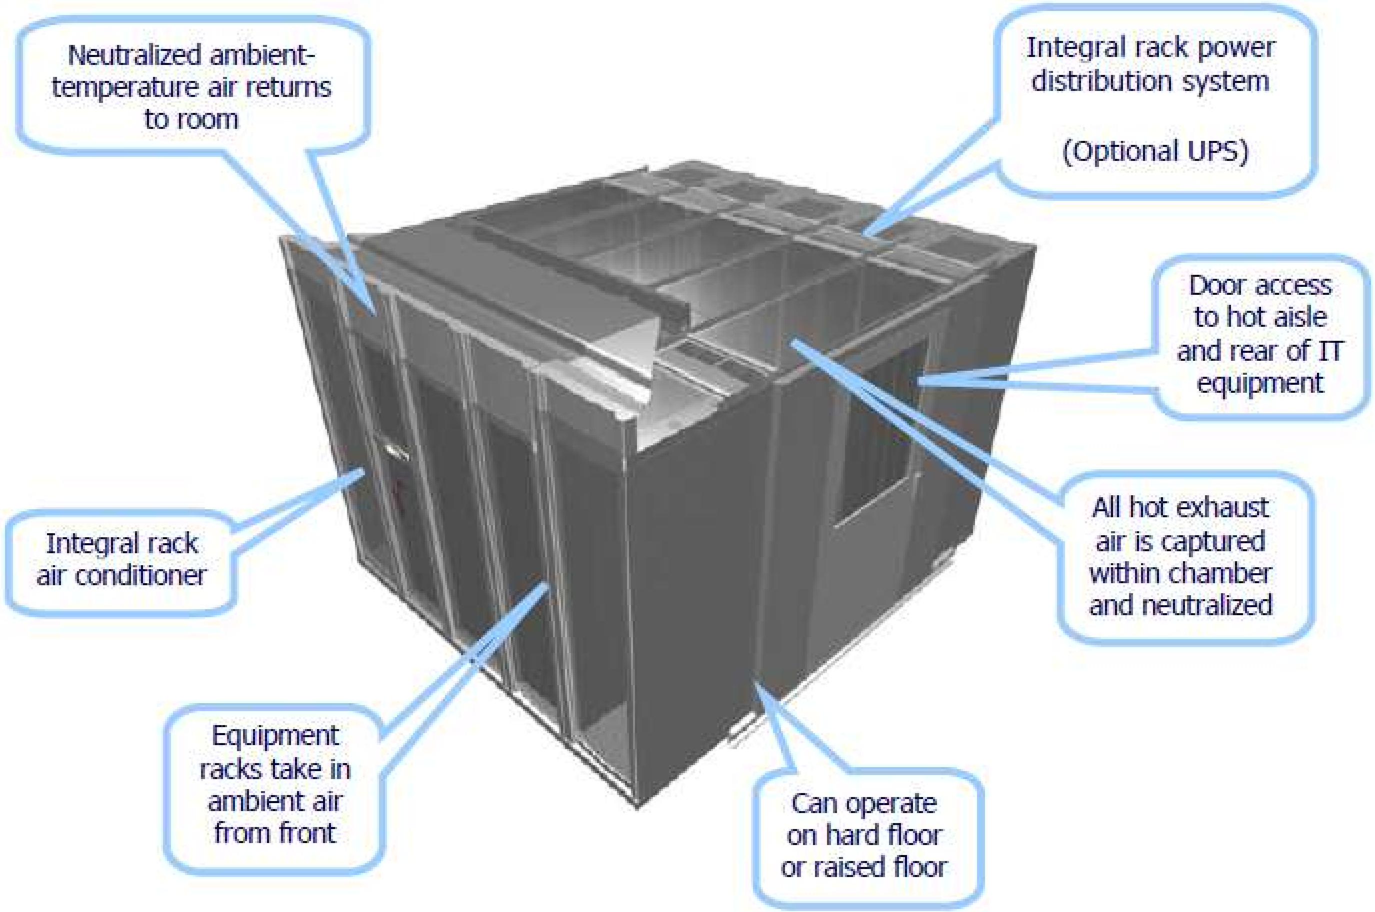

Method 5: Whole-room cooling

- *Provide the room with the capability to power and cool any and every rack to the peak expected enclosure density.*
  - This is conceptually the simplest solution but is never implemented because data centers always have substantial variation in per-rack power and designing for the worst case is consequently wasteful and cost prohibitive. Furthermore, to design for an overall rack power density of over 6 kW per rack requires extremely complex engineering and analysis.

Table 2

Application of the five approaches to cooling high density enclosures

Approach	Advantages	Disadvantages	Application
1 Spread the load **Split equipment among enclosures to keep peak load down**	Works anywhere, no planning needed Essentially free in many cases	High density equipment must be spread out even more than approach 2 Uses more floor space Can cause data cabling issues	Existing data centers, when high density equipment is a small fraction of the total load
2 Borrowed cooling **Provide average cooling capability with rules to allow borrowing of underutilized capacity**	No new equipment needed Essentially free in many cases	Limited to about 2X the design power density Uses more floor space Requires enforcement of complex rules	Existing data centers, when high density equipment is a small fraction of the total load
3 Supplemental cooling **Provide average cooling capability with provision for supplemental cooling equipment**	High density where needed and when needed Deferred capital costs High efficiency Good floor space utilization	Limited to about 10 kW per enclosure Racks and room must be designed in advance to support this approach	New construction or renovations Mixed environment Location of high density equipment is not known in advance
4 High density area **Create a special high density row or zone within the data center**	Maximum density Optimal floor space utilization High density equipment does not need to be spread out High efficiency	Need to plan a high density area in advance, or reserve space for it Must segregate high density equipment	Density 10-25 kW per rack When there is a requirement to co-locate high density devices New construction or renovations
5 Whole Room **Provide high density cooling capability to every rack**	Handles all future scenarios	Extreme capital and operating costs of up to 4X alternative methods May result in extreme underutilization of expensive infrastructure	Rare and extreme cases of large farms of high density equipment with very limited physical space

Figure 8 – Data center area per kW of capacity as a function of rack power density

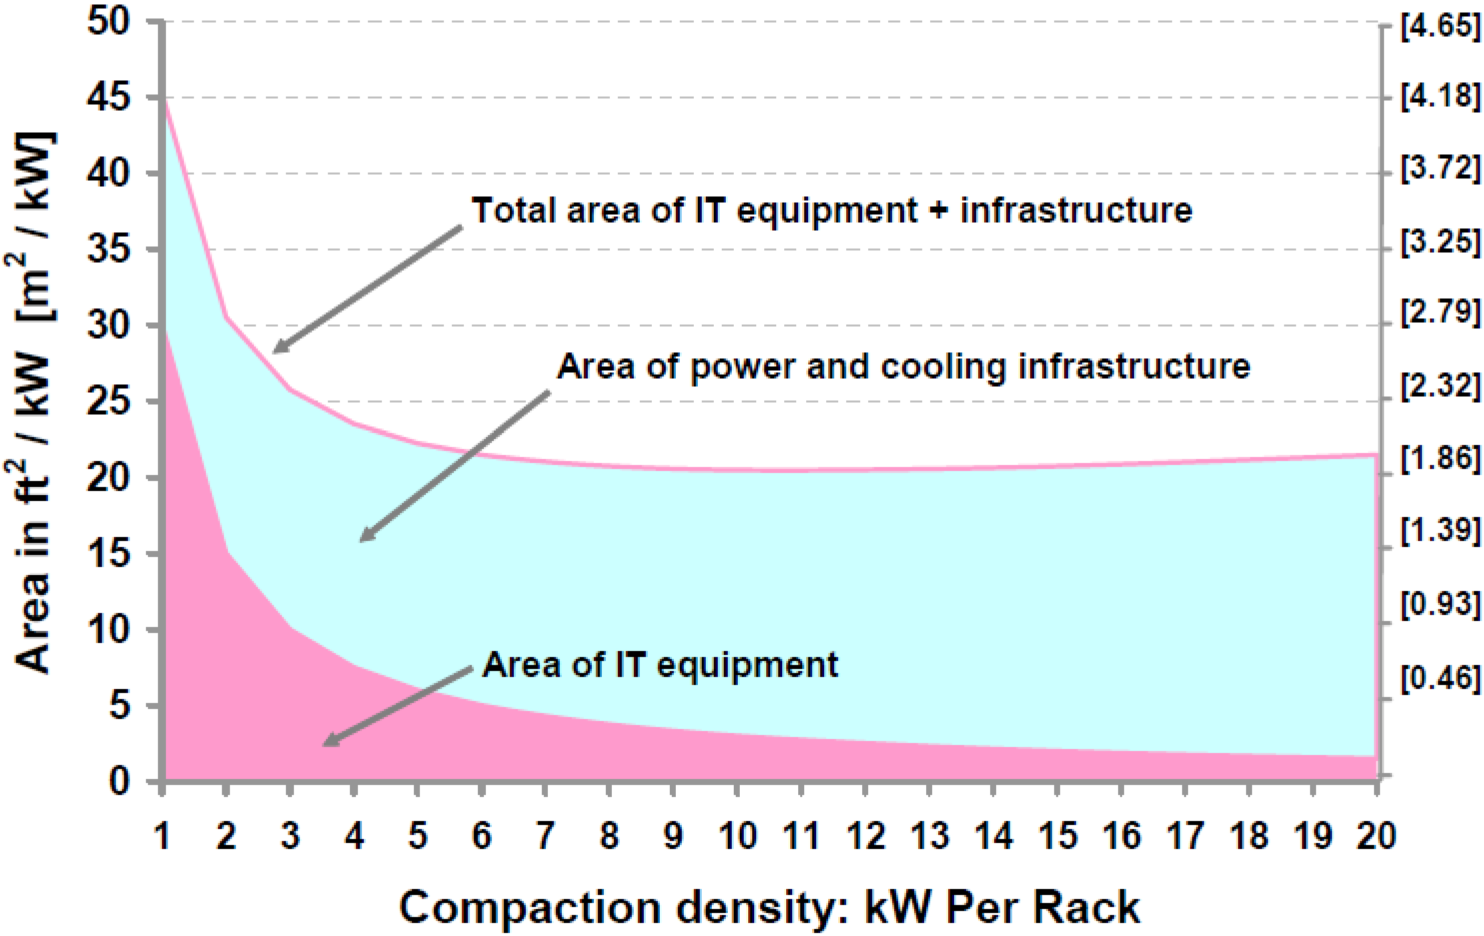

Practical strategy to optimize cooling when deploying high density computing equipment

IT equipment improvement	Analysis
1) Ignore physical size of IT equipment and focus on functionality per Watt consumed	This is an effective way to minimize area and TCO.
2) Design the system to permit later installation of supplemental cooling devices	This allows supplemental cooling equipment to be installed later where and when needed on a live data center, in the face of uncertain future requirements.
3) Choose a baseline power density for new designs between 40 and 100 W / ft^2 [0.4 – 1.1 kW / m^2], with 80 W / ft^2 [0.9 kW / m^2] (2800 W / enclosure average) being a practical value for most new designs	The baseline power density should be selected to avoid significant waste due to oversizing, and by keeping it below 100 W / ft^2 (1.1 kW / m^2) the performance and redundancy capability becomes predictable.
4) When the fraction of high density loads is high and predictable, establish and equip special high density areas of 100-400 W / ft^2 [1.1-4.3 kW / m^2] (3-12 kW per enclosure) within the data center	When it is known in advance that an area of high density is needed, and spreading the load out is not feasible. This can add significant cost, time, and complexity to the data center design. These areas will use specialized cooling equipment and not the typical raised floor design.
5) Establish policies and rules, which determine the allowable power for any enclosure based on its location and the adjacent loads	When understanding of the design capabilities is combined with power monitoring, the implementation of rules for new equipment installs can reduce hotspots, help ensure cooling redundancy, increase the system cooling efficiency and reduce electrical consumption. More sophisticated rules and monitoring can enable higher power density.
6) Use supplemental cooling devices when indicated	Installing supplemental cooling devices where needed and when needed can boost the cooling capacity of a region of a data center by up to 3X the design value to accommodate high density equipment.
7) Split equipment up that cannot be installed to meet the rules	The lowest cost and lowest risk option, but can consume considerable space when there are more than a small fraction of high density loads. Many users who do not have significant area constraints choose this as their primary strategy.

Physical design of data center systems

(IBM J. RES. & DEV. VOL. 49, 2005, pp. 709-723)

- Layout of computer rack equipment
  - Data centers are typically arranged into hot and cold aisles.
- Air distribution configurations
  - Hot rack exhaust air recirculation—A major cooling problem.

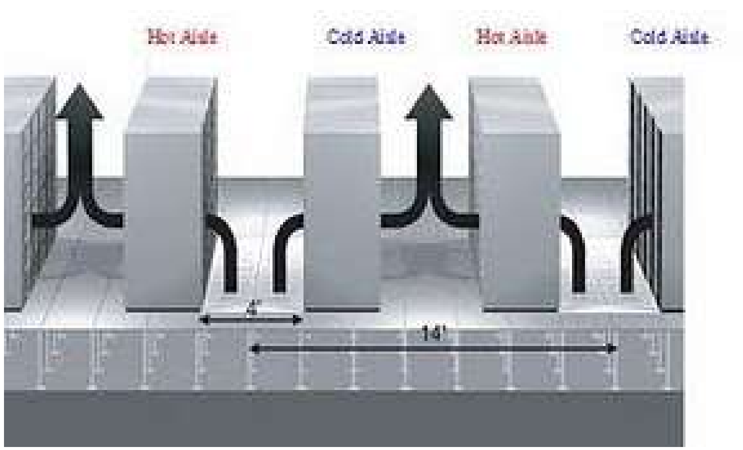

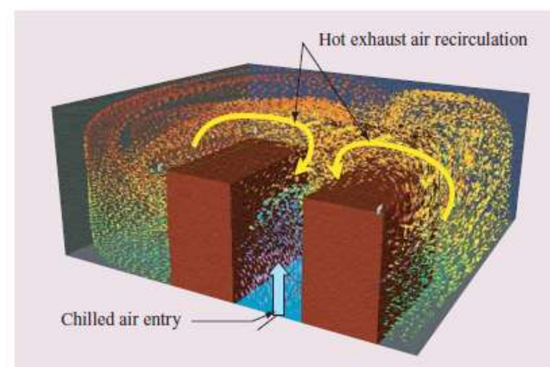

Figure 5

Computer-based simulation of hot-air recirculation in a raised-floor data center.

Figure 7 – Rack arrangement with no separation of hot or cold aisles

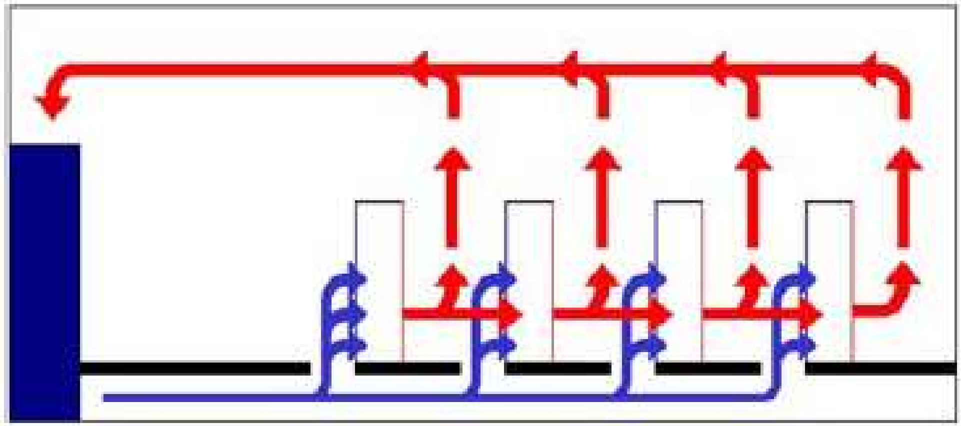

Figure 8 – Hot aisle / cold aisle rack arrangement

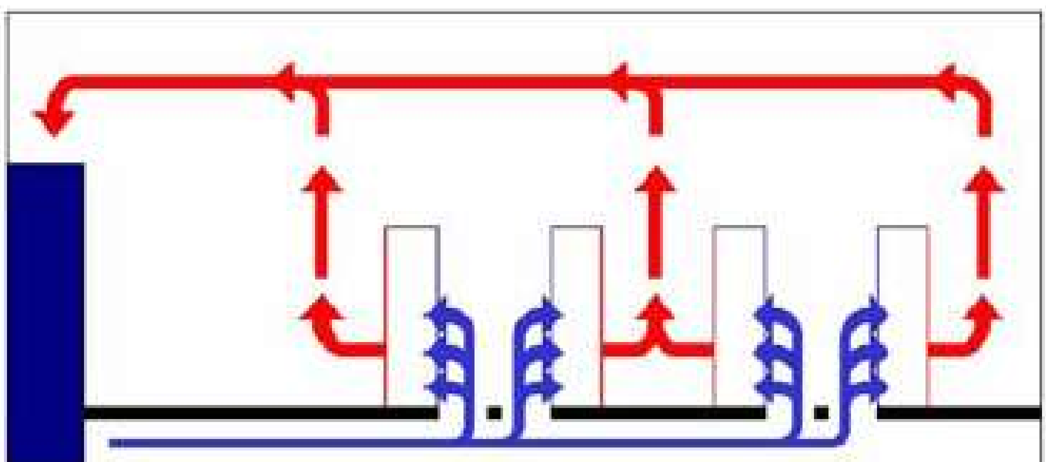

APC white paper #49

Summary of rack airflow design flaws with consequences

Design flaw	Availability consequences	TCO consequences	Solution
• No blanking panels • Equipment on shelves • Use of 23 inch (584mm) racks without rail-brushes	• Hot spots, particularly at the tops of racks • Loss of cooling redundancy	• Electricity costs • Reduced capacity of CRAC • Humidifier maintenance • Water consumption	• Use blanking panels • Do not use shelves • Use racks that have no open space outside the rails • Add brushes outside rails on wide racks
• Under-rack wire openings without brushes	• Hot spots, particularly at the tops of racks • Loss of static pressure in raised floor • Loss of cooling redundancy	• Reduced efficiency of CRAC	• Use brushes or gasketing on under-rack wire openings
• Glass doors • Doors with low ventilation	• Overheating • Amplification of problems relating to blanking panels	• Decreases space and rack utilization	• Use fully vented doors front and rear
• Use of fan trays and roof fans	• Very little benefit • Same investment could have been used for useful purpose	• Wasted capital • Wasted electricity	• Do not use fan trays or roof fans
• Shallow racks	• Cable obstructions cause overheating	• Decreases space and rack utilization	• Use racks with enough depth to allow free air around cables

APC white paper #49

Table 2

Summary of rack layout design flaws with consequences

Design flaw	Availability consequences	TCO consequences	Solution
• Racks all facing in the same direction • Hot-aisle-cold-aisle not implemented	• Hot spots • Loss of cooling redundancy • Loss of cooling capacity • Humidifier failures	• Excess power consumption • Water consumption • Humidifier maintenance	• Use hot-aisle-cold-aisle layout
• Not in rows	• Same problems	• Same	• Arrange racks in rows
• In rows but not tightly butted	• Same problems	• Same	• Bay racks together • Do not space racks out
• Racks all facing in the same direction • Hot-aisle-cold-aisle not implemented	• Hot spots • Loss of cooling redundancy • Loss of cooling capacity • Humidifier failures	• Excess power consumption • Water consumption • Humidifier maintenance	• Use hot-aisle-cold-aisle layout

Table 3

Summary of load distribution design flaws with consequences

Design flaw	Availability consequences	TCO consequences	Solution
Concentrated Loads	Hot spots Loss of cooling redundancy	Excess power consumption	Spread loads out evenly as possible

APC white paper #49

Table 4

Summary of cooling setting design flaws with consequences

Design flaw	Availability consequences	TCO consequences	Solution
Humidity set too high	Hot spots Loss of cooling redundancy	Excess power consumption Water consumption Humidifier maintenance	Set humidity at 35-50%
Multiple CRAC units fighting to control humidity of the same space	Loss of cooling redundancy Loss of cooling capacity	Excess power consumption Water consumption Humidifier maintenance	Set all units to the same setting Set 5% dead-band on humidity set points Use centralized humidifiers Turn off unnecessary humidifiers

APC white paper #49

Table 5

Summary of air delivery and return design flaws with consequences

Design flaw	Availability consequences	TCO consequences	Solution
• Hot air return location not over hot aisle • Dropped-ceiling lamp with integral air return located over cold aisle	• Hot spots, particularly at the tops of racks • Loss of cooling redundancy	• Electricity costs • Reduced capacity of CRAC • Humidifier maintenance • Water consumption	• Locate hot air returns over hot aisle • Do not use lamps with integral air returns over cold aisles, or block the return
• Overhead delivery vents over hot aisles • Vented floor tile in hot aisle	• Hot spots • Loss of cooling redundancy	• Electricity costs • Reduced capacity of CRAC • Humidifier maintenance • Water consumption	• For overhead delivery always locate delivery vents over cold aisles • For raised floor delivery always locate delivery vents in cold aisles
• Vented floor tile near no load • Overhead delivery vent open above no load • Peripheral holes in raised floor for conduits, wires, pipes	• Small	• Electricity costs • Reduced capacity of CRAC	• Close vents or openings located where there is no loads
• Low height of return vent in high ceiling area	• Loss of CRAC capacity • Loss of cooling redundancy	• Electricity costs • Reduced capacity of CRAC • Humidifier maintenance • Water consumption	• Use dropped ceiling for return plenum, or extend duct to collect return air at high point

Air velocity from floor grilles

- It is important to understand that the cooling capacity of the cabinet is directly related to the airflow volume delivery stated in CFM (cubic feet per minute). IT equipment is designed to raise the temperature of the supply air by 20-30°F (11-17°C). Using the equation for heat removal, the amount of airflow required at a given temperature rise can be quickly computed.

CFM or m^3/s = the volume of airflow required to remove the heat generated by IT equipment

Q = the amount of heat to be removed expressed in kilowatts (kW)

Δ°F or Δ°C = the exhaust air temperature of the equipment minus the intake temperature

$$CFM = \frac{3{,}412 \times Q}{1.085 \times \Delta°F} \qquad m^3/s = \frac{Q}{1.21 \times \Delta°C}$$

APC white paper #46

- Variations in tile airflow are undesirable but some averaging and sharing of airflow do take place within the data center; for this reason, **individual tile airflow variations of 30% should be considered acceptable**. However, the huge variations shown for grate-tiles in **Figure 4** are not acceptable since a fraction of rack locations would not receive sufficient cooling capacity.

Figure 4 – Airflow variation as a function of raised floor depth, for two different tile types

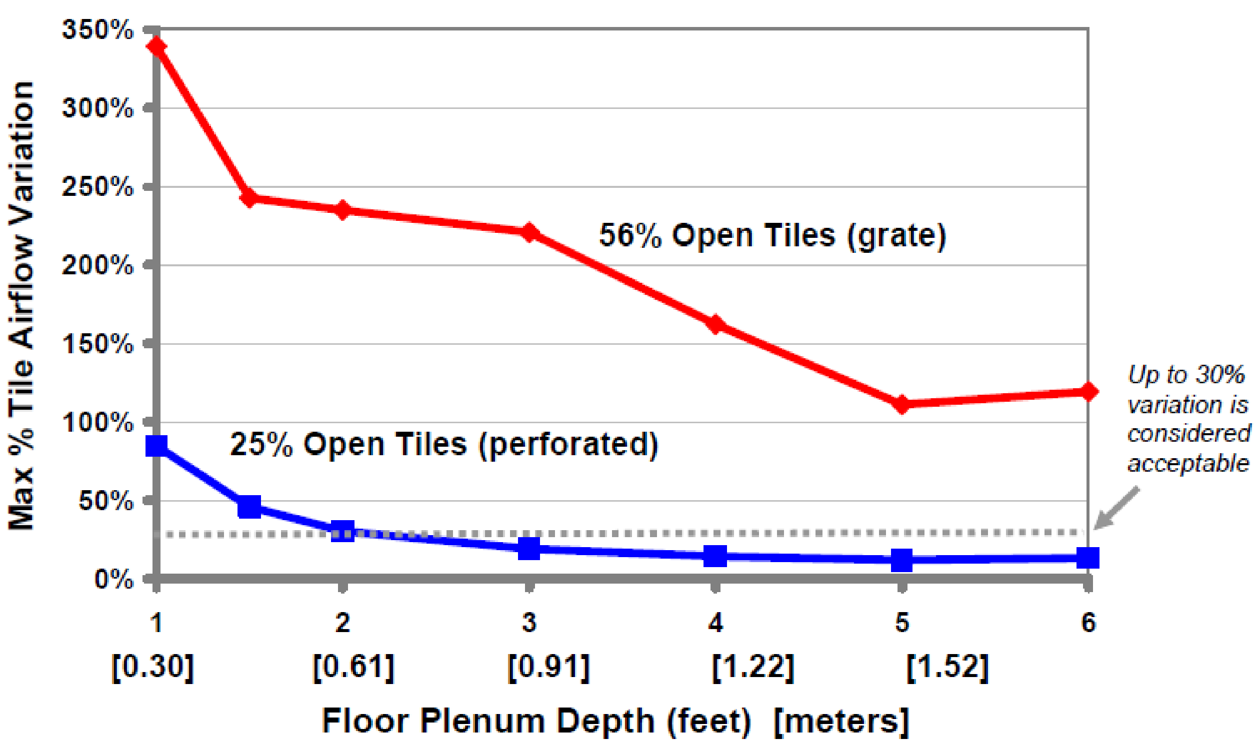

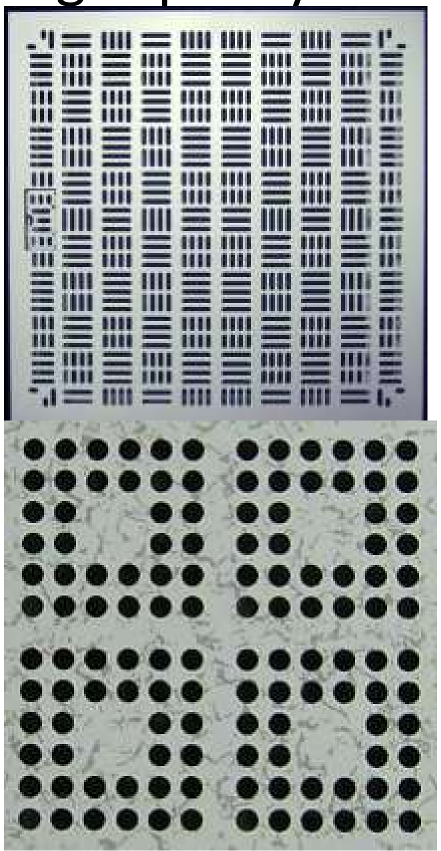

Figure 3 – Available rack enclosure cooling capacity of a floor tile as a function of per-tile airflow

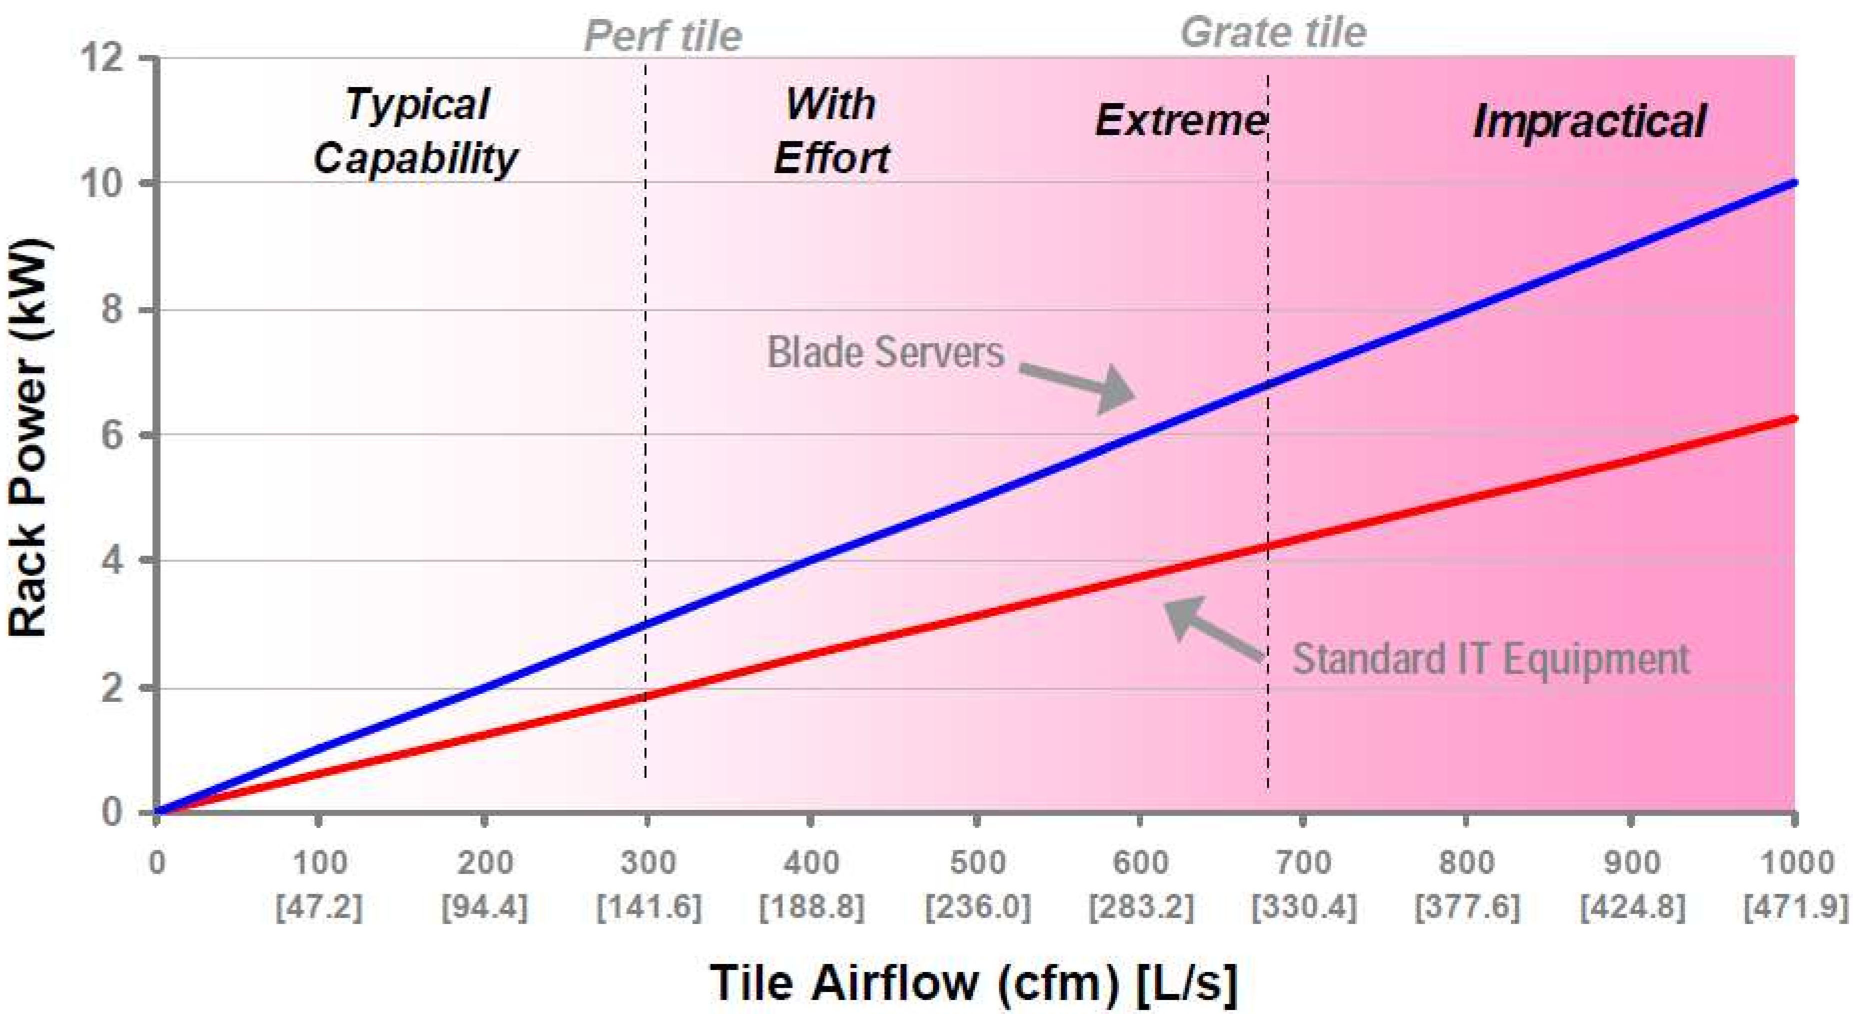

Air cooling for datacenter cooling - common questions

- Which one is the best ventilation scheme for air cooling the datacom equipment?
- What underfloor plenum height should one provide to allow proper distribution?
- What ceiling height is the most optimum for both raised floor (underfloor supply) and non-raised floor (overhead supply)?
- Where should one place the cabling trays and piping under a raised floor so as to minimize airflow distribution problems?
- How should one design for future growth, which accommodates increases in datacom equipment power and, therefore, increased cooling needs?
- Where should the computer room air-conditioning (CRAC) units be placed for the most efficient use?

Raised Floor Design

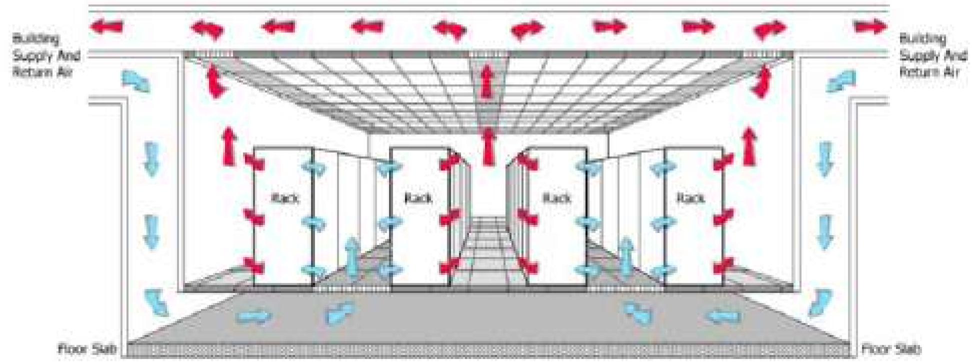

Raised floor implementation using building air from a central plant.

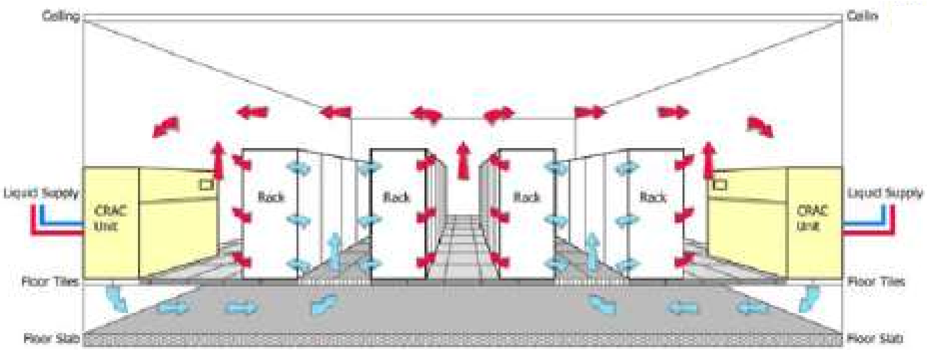

Raised floor implementation most commonly found in data centers today using CRAC units.

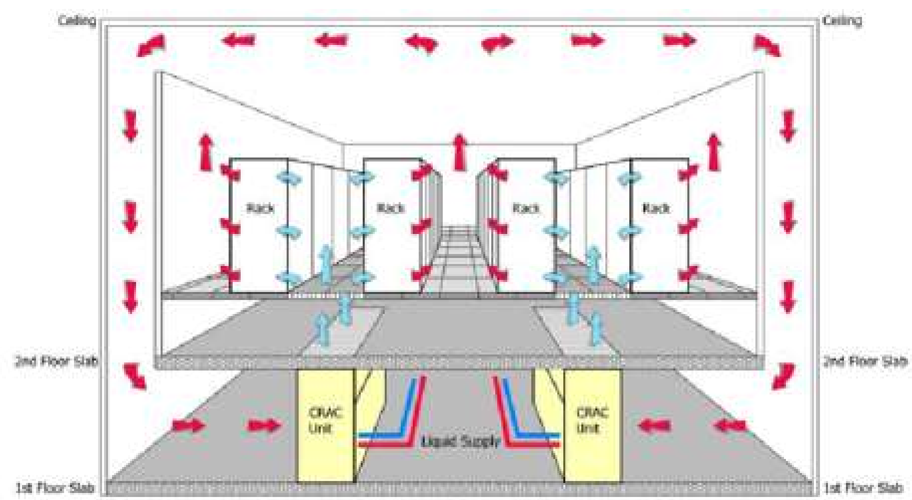

Raised floor implementation using two-story configuration with CRAC units on the lower floor.

Elements of the raised floor

- The raised floor was developed and implemented as a system intended to provide the following functions:
  - A cold air distribution system for cooling air Tracks, conduits, or supports for data cabling
  - Conduits for power cabling
  - A copper ground grid for grounding of equipment
  - A location to run chilled water or other utility piping

The problems of using a raised floor

(APC white paper #19)

- Earthquake
- Floor loading
  - The buckling (lateral) strength of the floor depends on the presence of the tiles. Routinely pulled in a data center when frequently required cabling changes or maintenance is performed. This situation can give rise to unexpected catastrophic collapse of the raised floor.
- Access
  - Modern data center is around two years gives rise to the situation where data and power cabling is subject to nearly continuous change. The difficulty of access to this cabling when it is under the raised floor results in delay and costs associated with changing needs.

- Headroom
  - The loss of headroom for some data centers is not acceptable.
  - In Japan, it is common for the floor of the next overhead level of the building to be cut out to create the headroom required.
- Conduit
  - When cabling is run under a raised floor it becomes subject to special fire regulations. consider a fire in an air plenum to be a special risk. Therefore, cabling under the raised floor is required to be enclosed in fire-rated conduit, which may be metal or a special fire rated polymer. The result is costly and complex to install this conduit, and a particularly difficult problem when conduit changes are required in an operating data center.

- Power distribution
  - The number of branch circuits per square foot in the modern data center is much greater than. The density of the resulting conduits associated with this dramatic increase in branch circuits represent a serious obstacle to underfloor air flow. This can require a raised floor height of 4 feet (1.2 meters) to ensure the needed airflow. Increasing the height of the raised floor compromises the structural integrity and compounds cost, floor loading, and earthquake issues.
- Security
  - The raised floor is a space where people or devices may be concealed.
  - In the case of data centers which are partitioned with cages, such as co-location facilities, the raised floor represents a potential way to enter and access caged areas. This is a reason why many co-location facilities do not use raised floor systems.

- Cost
  - The raised floor represents a significant cost.
- Cleaning
  - The raised floor is an area which is not convenient to clean. Dust, grit, and various items normally accumulate under the raised floor and are typically abandoned there since the difficulty and accident risk associated with cleaning this area are considered to be serious obstacles.
- Safety
  - A tile left open poses a severe and unexpected risk to operators and visitors moving in the data center.

Modifications of raised floor design

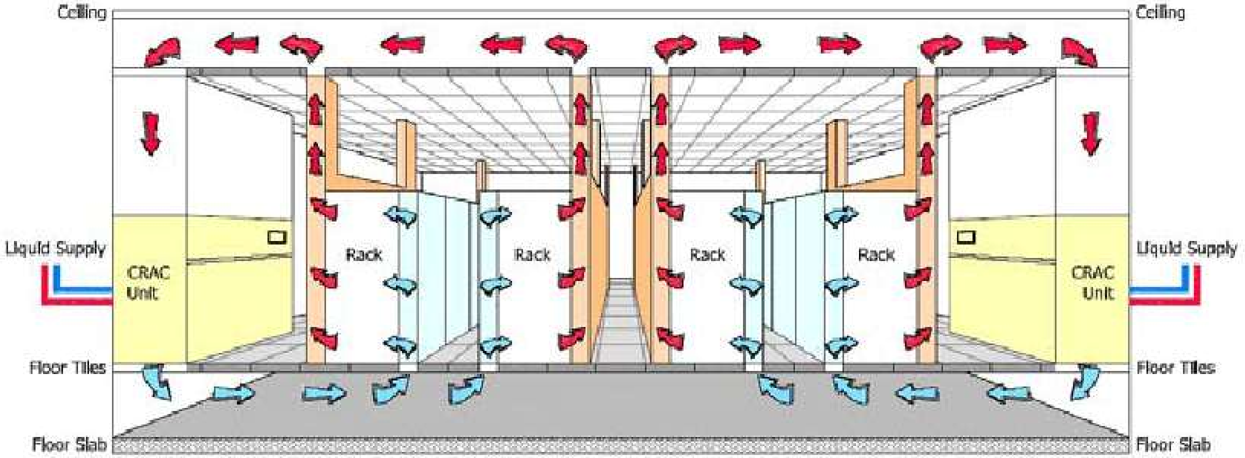

Raised floor implementation using inlet and outlet plenums/ducts integral to the rack.

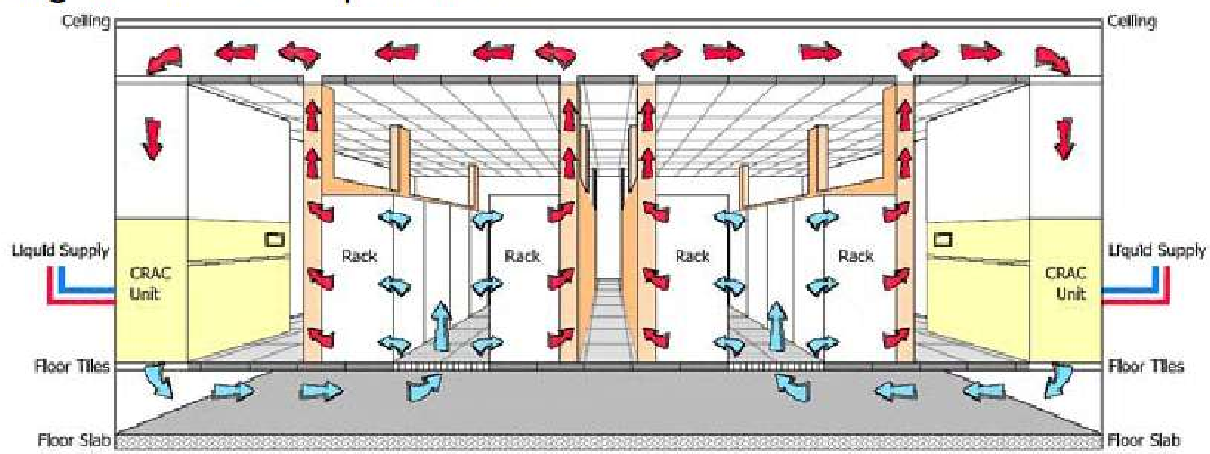

Raised floor implementation using outlet plenums/ducts integral to the rack.

Local Distribution Design

- Introduce cold air to the aisle as close as possible

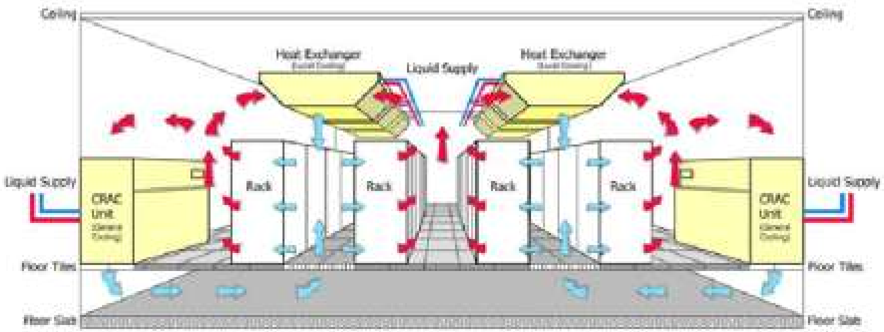

Local cooling distribution using overhead cooling units mounted to the ceiling.

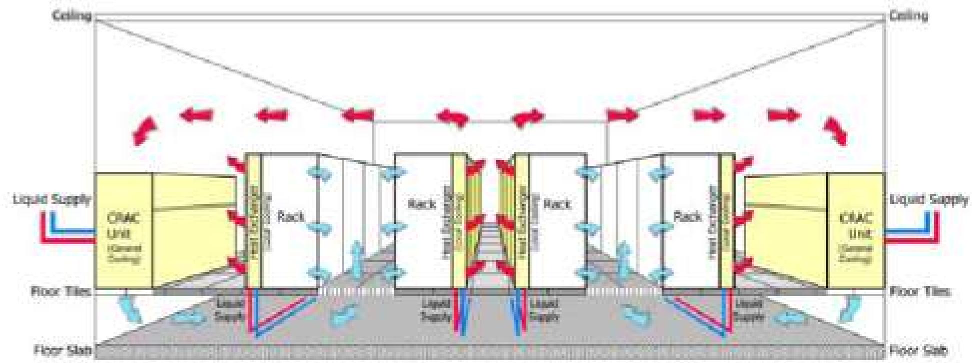

Local cooling via integral rack cooling units on the exhaust side of the rack.

- Reduce poor air distribution and mixing

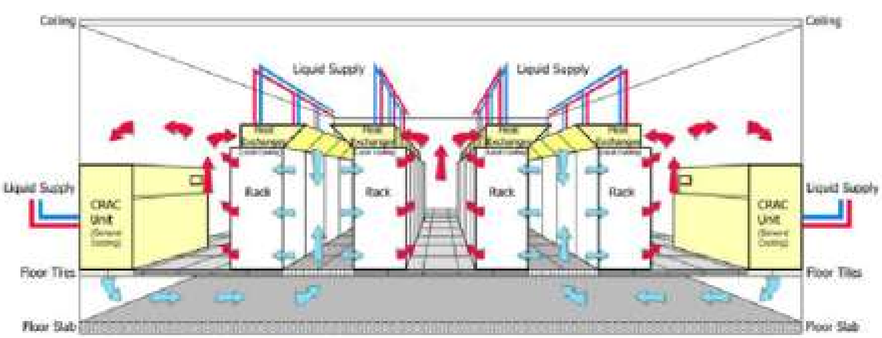

Local cooling distribution using overhead cooling units mounted to the rack.

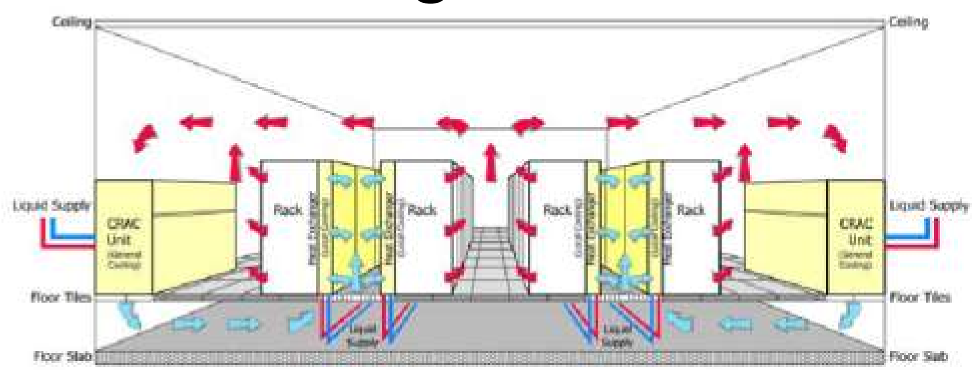

Local cooling via integral rack cooling units on the inlet side of the rack.

- Cooling @ exhaust is preferred to eliminate condensate carryover

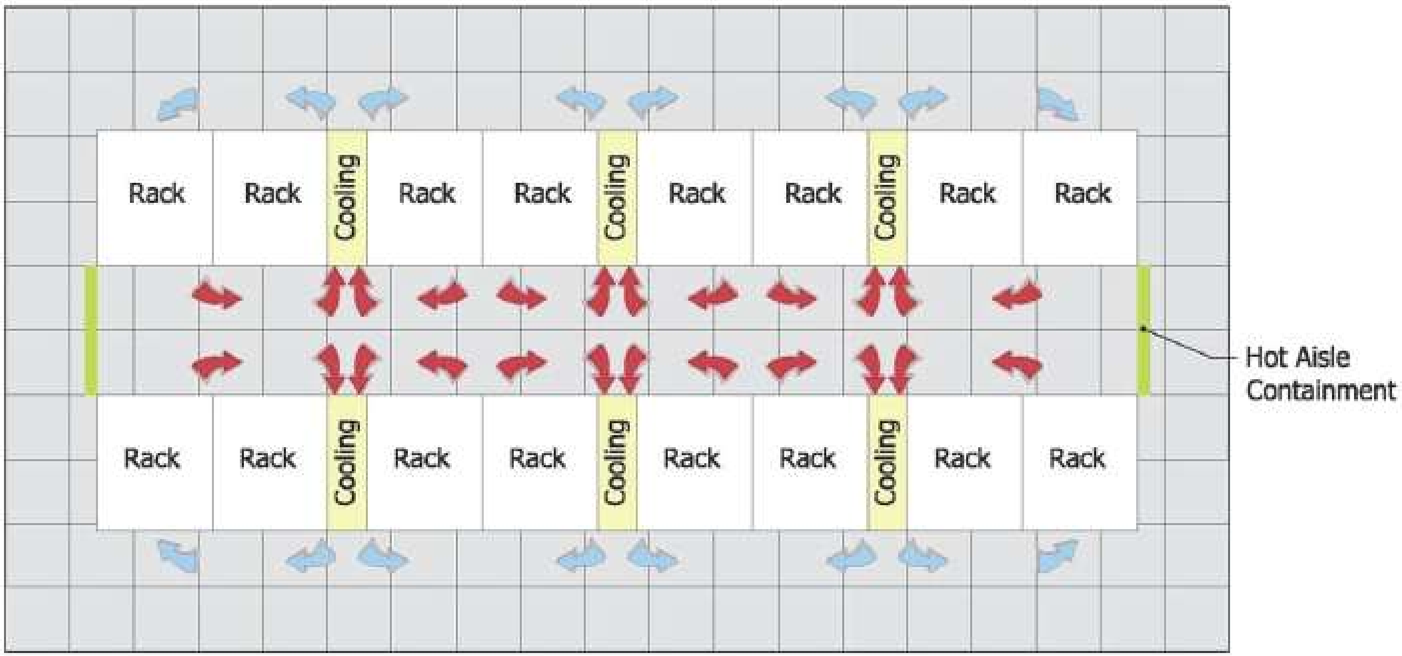

Figure 5.16 Local cooling units interspersed within a row of racks.

Overhead Design

- Overhead design provide a higher static pressure which normally give rise to a better air flow distribution.

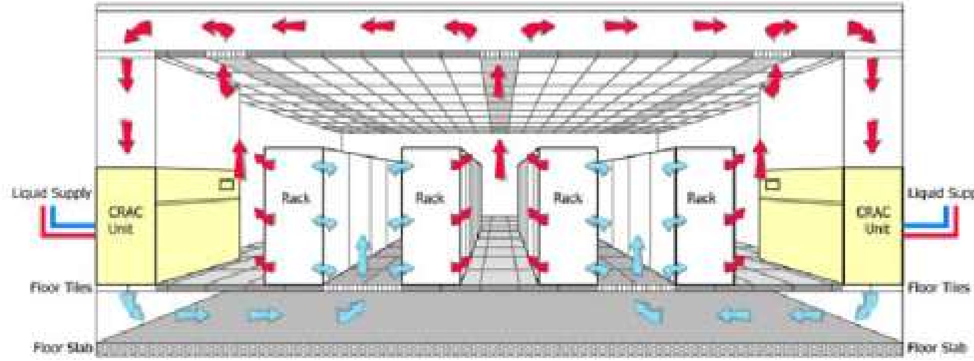

Raised floor implementation using a dropped ceiling as a hot air return plenum.

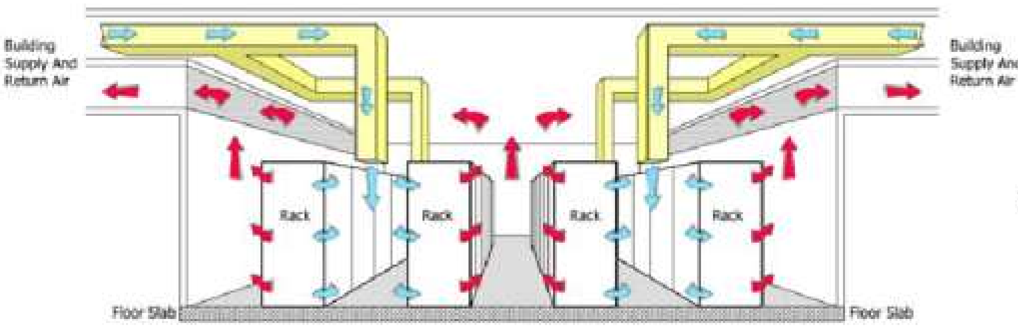

Overhead cooling distribution commonly found in central environments.

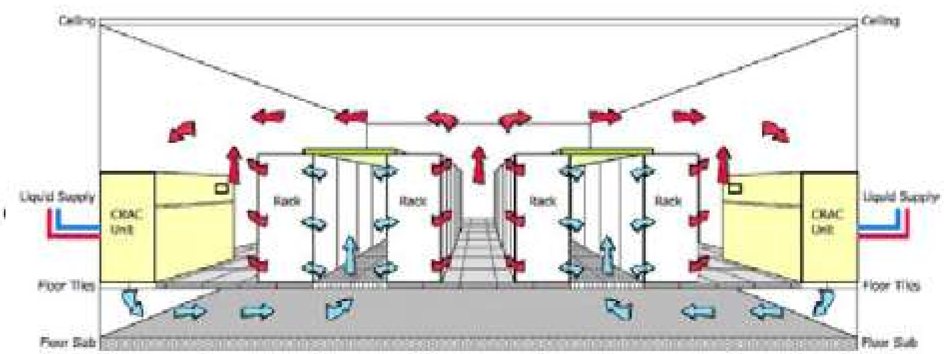

Raised floor implementation using baffles to limit hot-aisle/cold-aisle "mixing."

Some Typical Cabinets Designs

ASHRAE Transactions, 2005, Vol. 111, pt. 2, pp. 715-724

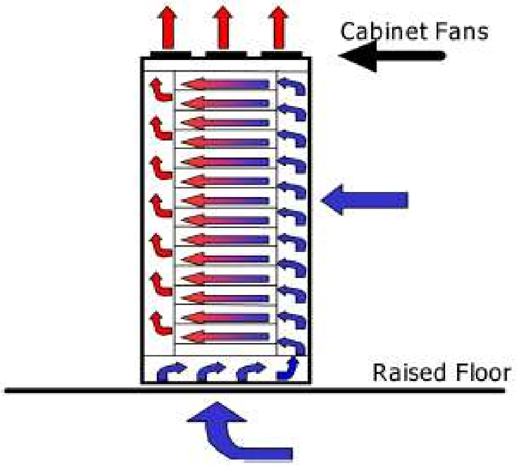

Figure 2 *Fans on top.*

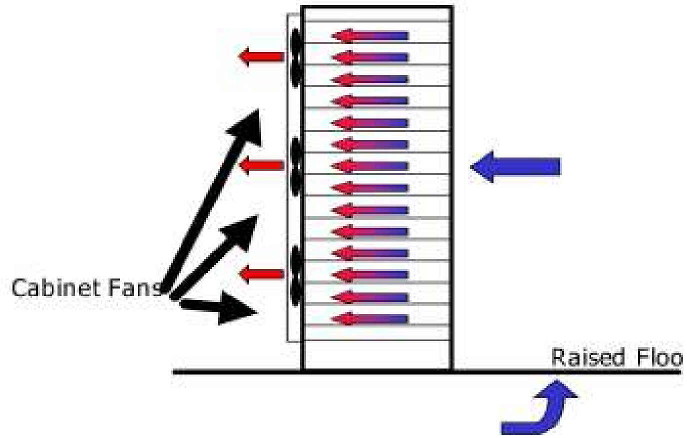

Figure 4 *Fans on rear (horizontal discharge).*

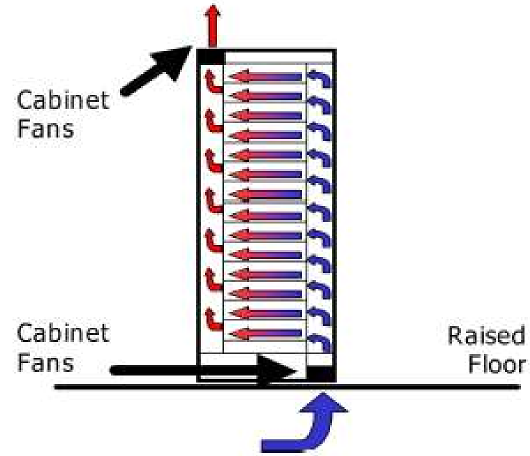

Figure 6 *Fans on bottom and top.*

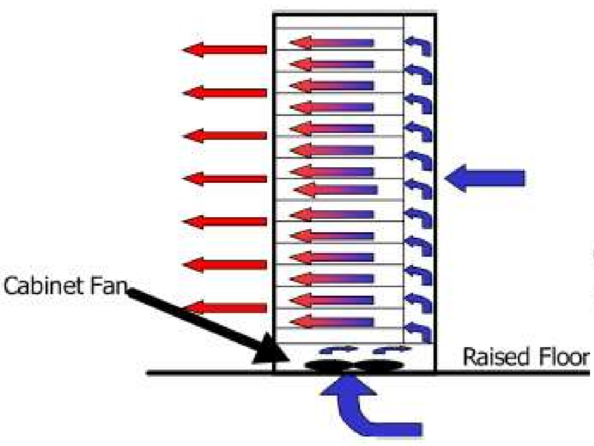

Figure 3 *Fans on bottom.*

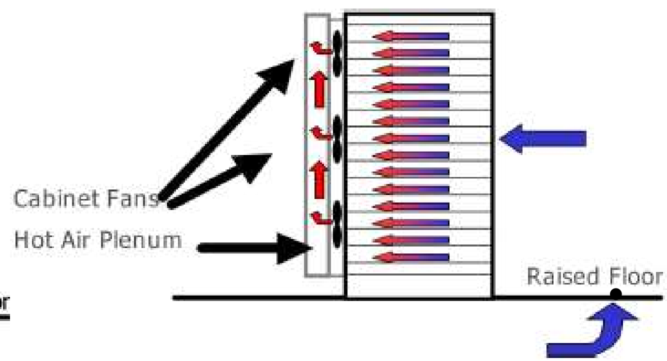

Figure 5 *Fans on rear (vertical discharge).*

- Cabinet fan in Figures 3 and 4 lend themselves to the “hot aisle/cold aisle” solution and were designed with that concept in mind.

  Cabinet arrangements in Figures 2, 5, and 6 lend themselves to the use of the ceiling plenum space as a return air plenum.

Data center air stream

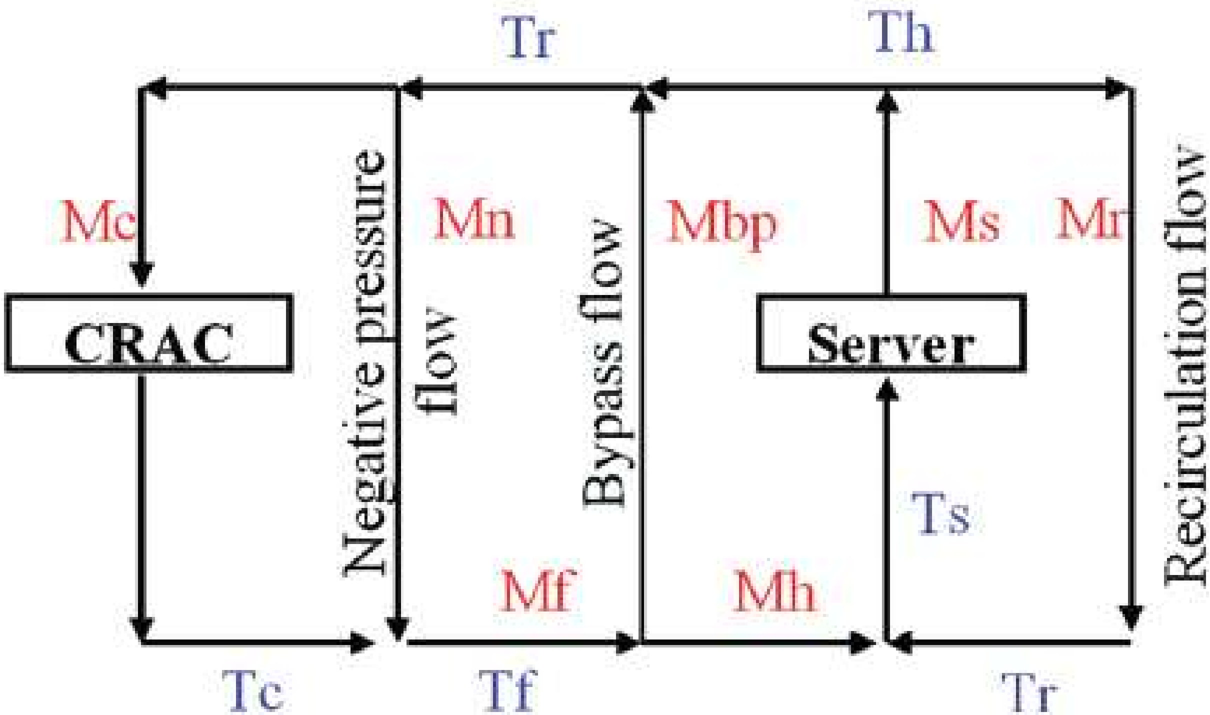

Data center air streams. (Note: CRAC: computer room air conditioning unit; Server: IT equipment server; bp: bypass; c:CRAC; f:floor; h:hall; M:mass flow rate; n: negative pressure; r: recirculation; s: server; T: temperature.)

- ***CRAC flow (Mc):*** total air-flow rate produced by all operating CRAC units in the data center. This air flow rate is normally much more than what is needed by the servers.
- ***Negative pressure flow (Mn):*** air that is induced into the floor void due to Venturi. With negative static pressure under the floor, air will be induced into the floor void from the space above
- ***Bypass air flow (Mbp):*** this is air that leaves the floor grills and returns directly to the CRAC unit without cooling servers.
- ***Recirculation airflow (Mr):*** air that is discharged from servers, which returns and mixes with air entering the servers to cool them.

Typical temperatures for a legacy data center :

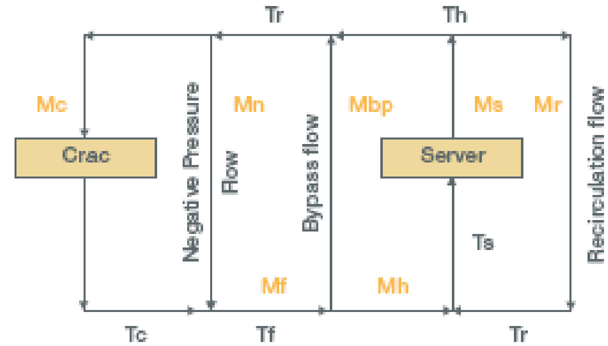

- Tr = 21°C [69.8°F], return air temperature to CRAC (at CRAC), normally the CRAC return air set point
- Tc = 14°C [57.2°F], discharge air temperature from CRAC (at CRAC)
- Tf = 14.1°C [57.4°F], floor void temperature (after room air is drawn in), very close to Tc
- Ts = 21°C [69.8°F], server inlet air temperature (mixture of grill and recirculation air)
- Th = 28°C [82.4°F], server outlet air temperature (before cold air mixes with it)

Bypass Air:

Methods to reduce

- Seal air gaps in the raised floor, particularly cable cutouts within cabinets at the discharge side of servers
- Locate floor air grills so that they supply the inlet of servers, i.e. remove floor air grills from hot aisles and other areas where cooling isn't required
- Ensure air velocities from floor grills are not too high, which can cause air to overshoot the top of the cabinet
- CRAC units turned off are a source of air bypass (consider air isolation with dampers)

Recirculation Air:
Some methods for improvement

- Fit blanking plates in cabinets where servers are not installed
- Close gaps between cabinets (where warm air can make its way to the server inlet)
- Ensure as best as possible the adequate supply of cold air to the server inlets
- Ensure air velocities and flows from floor grills are sufficient to reach the top of cabinet servers
- Review air return path to CRACs and consider return air plenum or ductwork if necessary
- Consider means to physically isolate the supply (cold) and return (hot) air streams
- Remove obstructions under the floor that restrict cold air to supply the server inlets

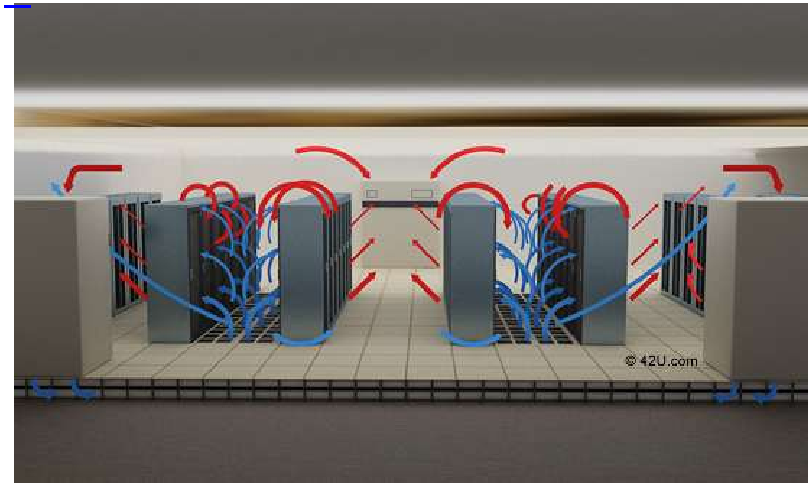

Diagram of a Data Center with mixing of hot and cold air - a very common cause of power loss.

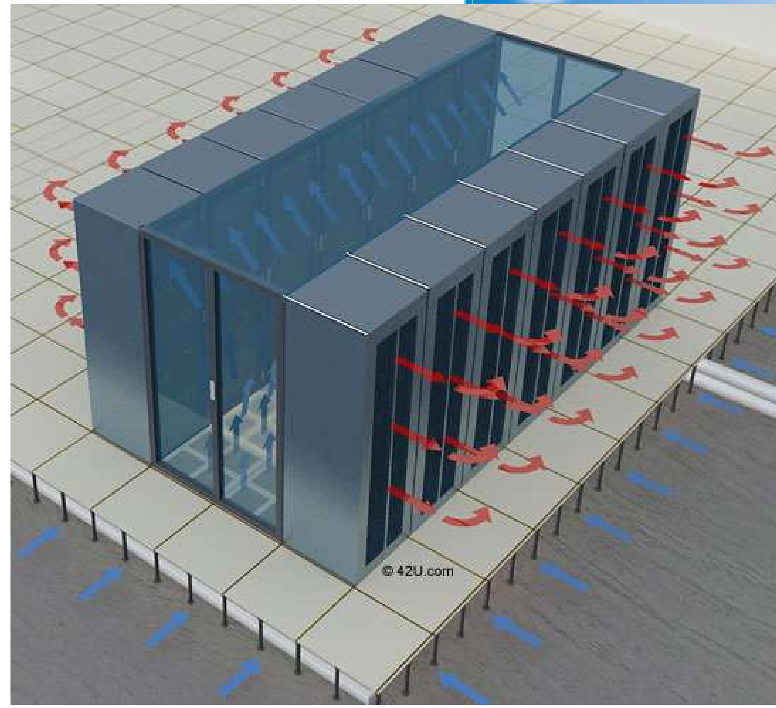

Cold Aisle Containment

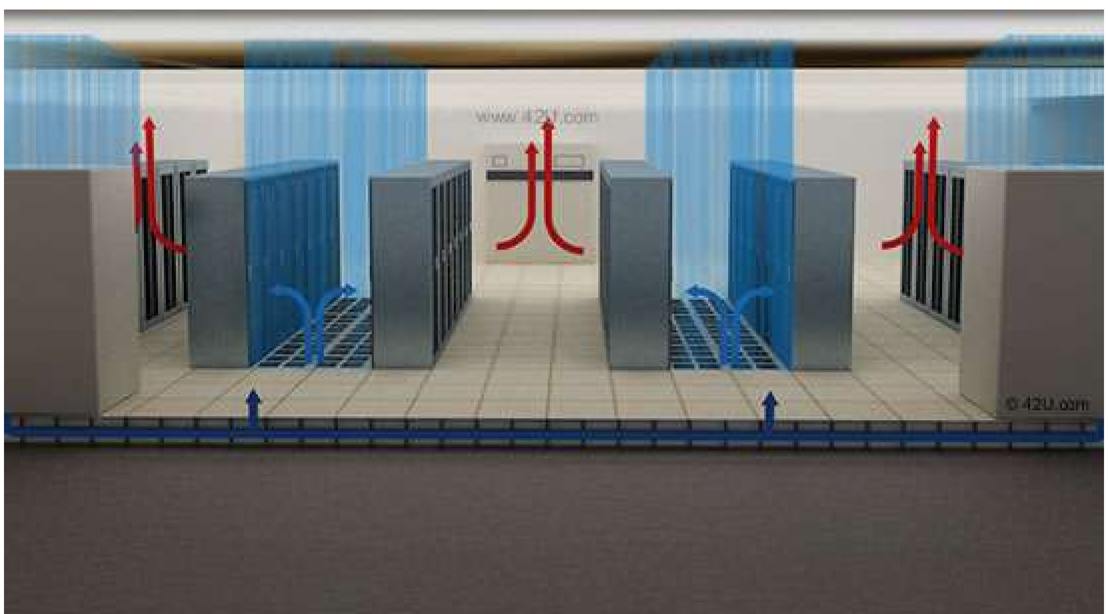

Diagram of a Data Center Curtains containing the airflow in the Cold Aisles

Conventional vs. Modern Hot aisle Containment Design (APC white paper #123)

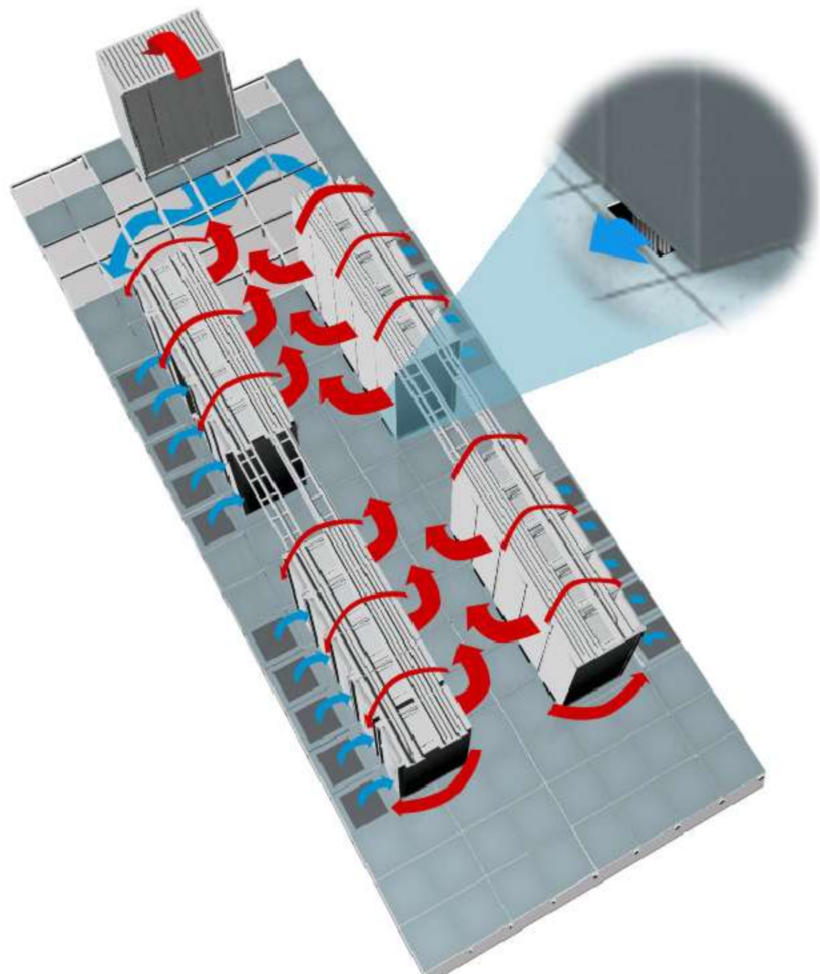

Figure 1

Air flow pattern with traditional cooling approach

Table 2

Results of CFD analysis

	Average hot aisle dry bulb temperature (°F / °C)	Worst case hot aisle wet bulb temperature (°F / °C)	Wet bulb globe temperature (OSHA / ISO max = 86°F / 30°C)	Average IT equipment inlet temperature (°F / °C)
Raised floor room	81.9 / 27.7	68.7 / 20.4	72.7 / 22.6	73.8 / 23.2
Enclosed hot aisle room	89.8 / 32.1	66.5 / 19.2	73.5 / 23.1	70.7 / 21.5

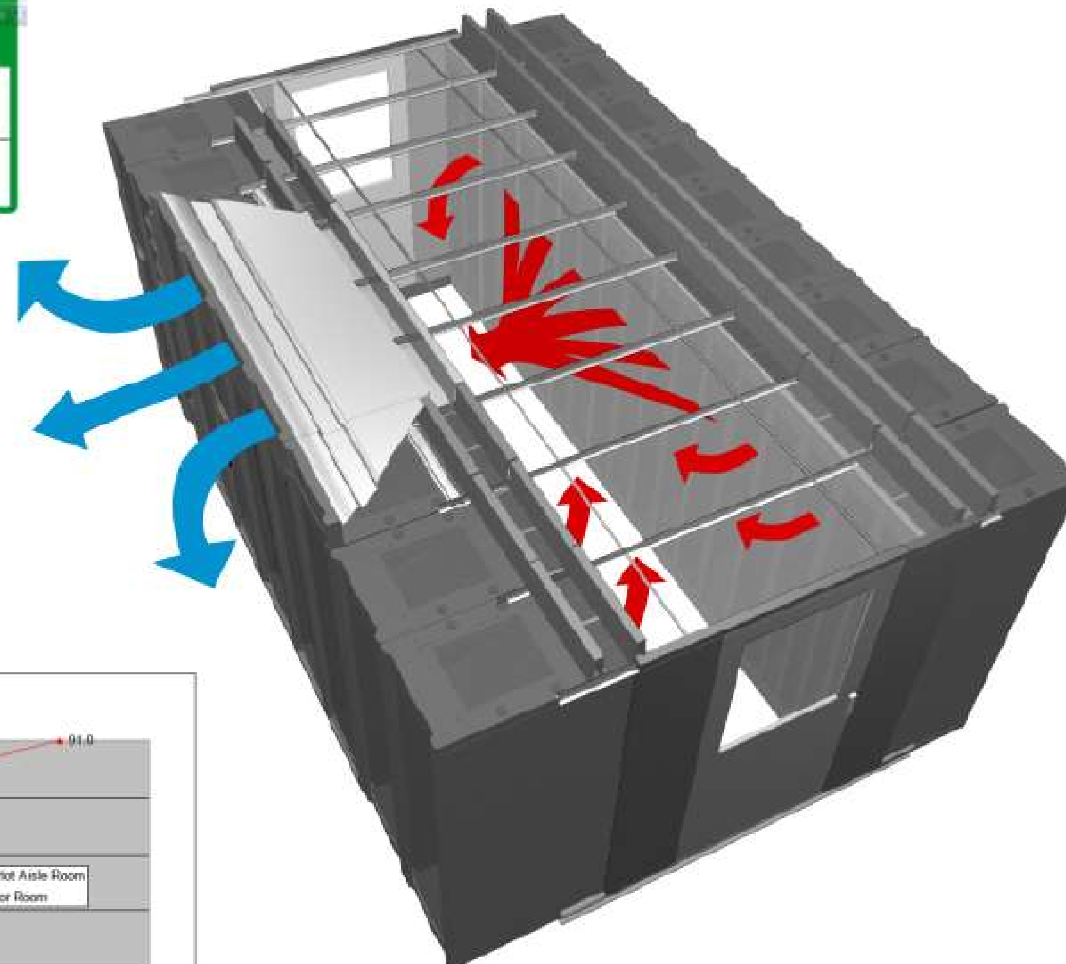

Figure 2

Modern cooling approach

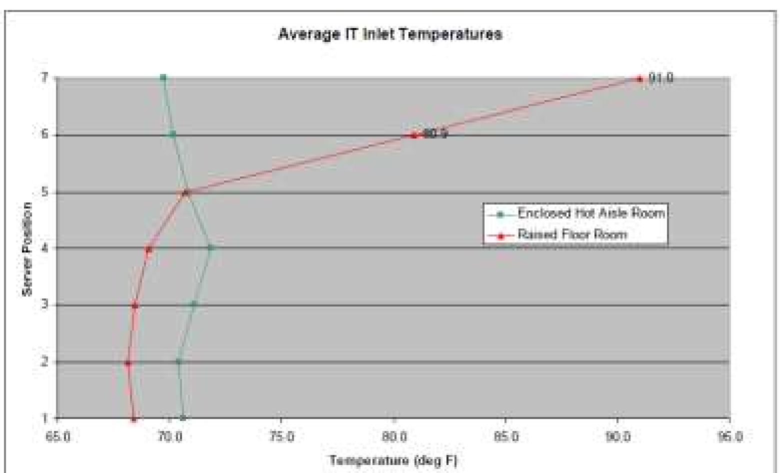

Figure 6

IT equipment inlet temperature as a function of mounting position in rack

Ventilation Design

Some guidelines

- The best ventilation scheme is a raised floor for chilled air supply (Figure 1b) with exhaust hot air removal from vents in the ceiling or top parts of the walls or venting via CRAC units (Figure 1a) that reside on the raised floor.
- The worst ventilation scheme is with overhead chilled-air supply (Figure 1c) with exhaust vents at the floor or at the bottom part.
- The typical underfloor supply design (Figure 1a) can result in hot spots at the very top part of the rack inlet. Which is due to hot air recirculation patterns. This does not occur in overhead supply designs. where the chilled air is supplied from the top, leading to good mixing at the top of the rack

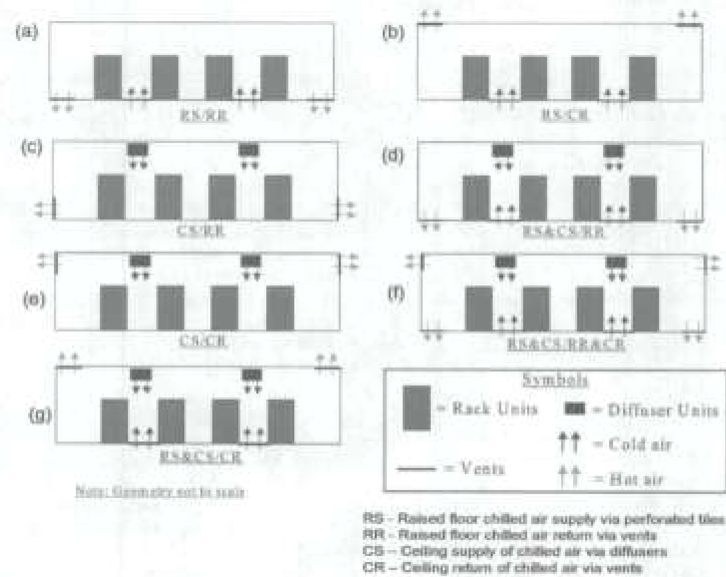

Figure 1 Different ventilation schemes studied by Shrivastava et al. (2005a, 2005b).

ASHRAE Transactions, 2007, pt. 1, pp. 206-218.

Ventilation Design

Some guidelines (Conti..)

- If hot exhaust air from equipment and cabinets can be directed upward into a ceiling return plenum, this may be superior to simply having a high ceiling.
- If flexibility exists in the orientation of rows of equipment, a layout that allows hot air unobstructed access to the return of CRAC units (or other cooling system returns) should be superior to a layout with rows perpendicular to the CRAC units
- A cold aisle/hot aisle arrangement should be followed in laying out racks within a data center: the fronts of the racks drawing in chilled air either from overhead or from the raised floor should face the chilled air exhausting into the cold aisle.

Raised Floor Plenum Height

- For low raised floors (6-12 in. [0.15-0.3 m]), do not place datacom equipment close to CRAC units since low airflow or reverse flow can occur from the perforated tiles.
- Airflow from a large number of perforated tiles can be made uniform if perforated tiles of varying percent openings.
- Partitions can be placed underneath the raised floor to direct air into the desired areas.
- It is suggested that raised floors should allow a free flow height of at least 24 in. (0.61 m); that is, if piping and cabling take up 6 in. (0.152 m) then the raised floor height should be 30 in, (0,76 m) A large underfloor depth of 24 in. (0.61 m) was also recommended.

Room Ceiling Height
- some guidelines

- The ceiling height will depend on the type of ventilation system that is chosen.
- When the supply chilled air from the perforated tiles exceeds the rack flow rates (110%), then increasing the ceiling height reduces the datacom equipment intake temperatures for three cases:
  - underfloor chilled air supply and room CRAC hot air return with a hood at the top of the CRAC.
  - underfloor air supply and ceiling hot air return that vents to the CRAC.
  - Overhead air supply and room CRAC hot air return at the bottom of the CRAC.
- For flows from the perforated tiles equal to or less than the datacom equipment flow rates, increasing the ceiling height can result in increased inlet temperatures for underfloor air distribution without a ceiling. A hot recirculation cell intensifies over the datacom equipment with increased height.

Room Ceiling Height
- some guidelines (Conti..)

- Increasing the ceiling height from 9 to 12 ft (2.74 to 3.66 m) reduces the rack inlet temperature by as much as 6 to 12°C (42.8°F to 53.6°F) in hot spot regions, with small impact (inconsistent) in lower flux regions for 6 ft (1.83 m) tall racks arranged in a hot aisle/cold aisle fashion on a raised floor.
- Increasing the ceiling height beyond 12 ft (3.66 m) for 6 ft (1.83 m) racks does not seem to have any impact on rack inlet temperatures for racks arranged in a hot aisle/cold aisle fashion on a raised floor with underfloor supply and room CRAC return.

Effect of Rack Placement
- some guidelines

- Placement of high-density datacom equipment at floor locations that have high static pressure allows the highest possible airflow in the cold aisle adjacent to the equipment. Typically, the highest static pressures are farther away from the CRAC units or where the flows from two or more CRAC units are in collision with each other.
- The recirculation of the hot air can be reduced by load spreading. By placing lower-powered equipment near high-powered equipment, the effect of the hot air recirculation can be reduced.
- Flow near the end of an aisle should be considered in detail, as recirculation can occur both around the sides and over the top of the cabinet.

Effect of Rack Placement
- some guidelines (Conti..)

- Some data centers have employed plastic stripes at the end of an aisle to prevent air recirculation but allow ease of access.
- If server inlet temperatures as measured by Thermal Guidelines for Data Processing Environments (ASHRAE 2008, 2nd. Ed.) are much less than those required by the manufacturer, then reducing airflow by turning off or reducing the air conditioning can result in significant energy savings.
- The inlet air temperature to high-powered racks can be reduced significantly by removing an adjacent rack. Racks in the vicinity of the removed racks also experience a reduced inlet air temperature.
- For some layouts, the best position for high-powered racks is near the end of the aisles. However, a more recent study has found that the outside racks at the ends of the aisles can experience more hot air recirculation.

Detrimental Underfloor Blockages

- some guidelines

- If possible, chilled water pipes and cabling should be kept away from the exhaust of the CRAC units.
- Underfloor blockages have the biggest influence on flow rate uniformity through the perforated tiles.
- Blockages that are parallel to the hot and cold aisles have much lower impact than those that run perpendicular to the aisle. This conclusion was found when the CRAC units are located parallel to the computer rack equipment aisles.
- Blockages occurring under the cold aisle have the effect of reducing tile flow.

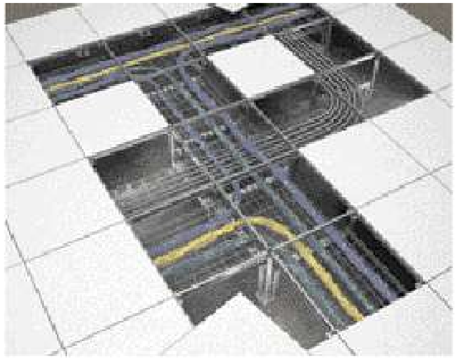

Reduce blockage effect

Seal cable cutouts

Integral – for new installations

Surface mount for existing installations

Cable management

Wrong

Correct

APC white paper #119

Figure 1

Establishment of order through modular deployment zones

Wire management in back of the rack

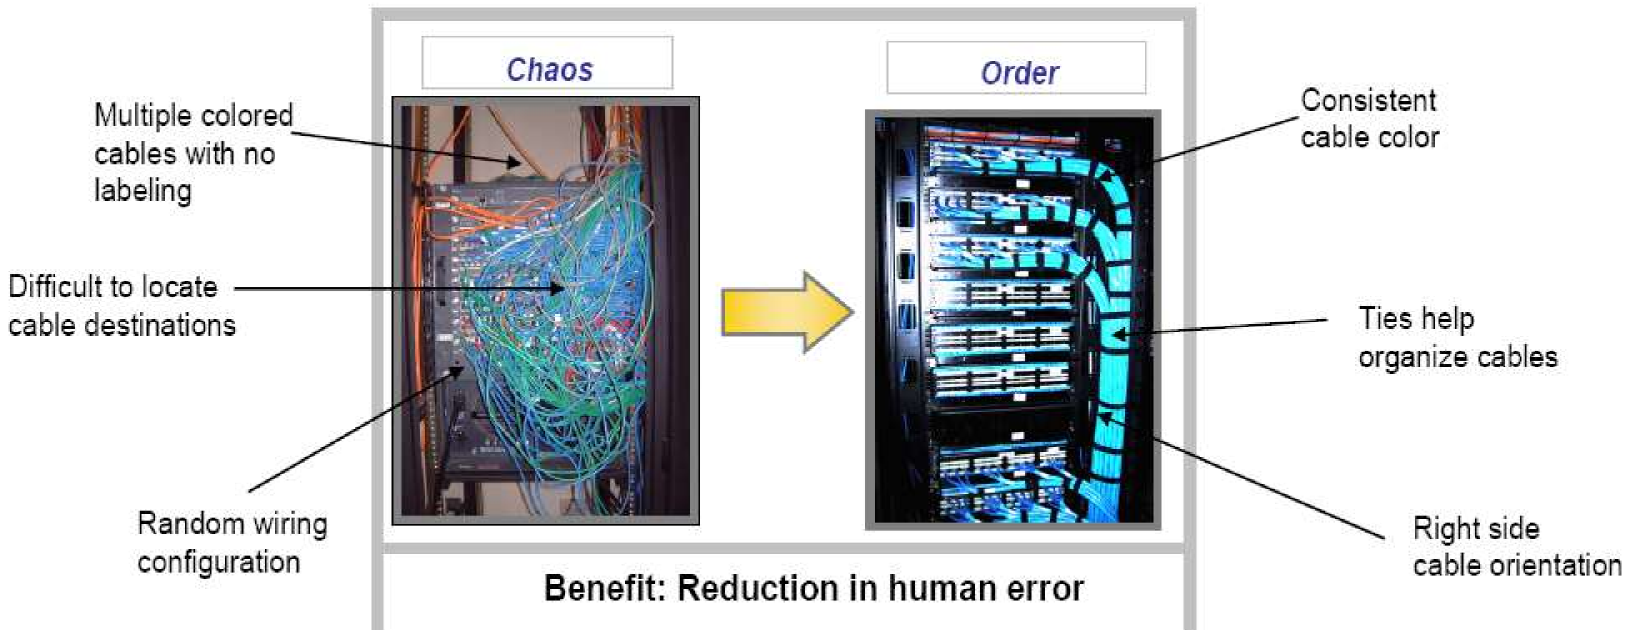

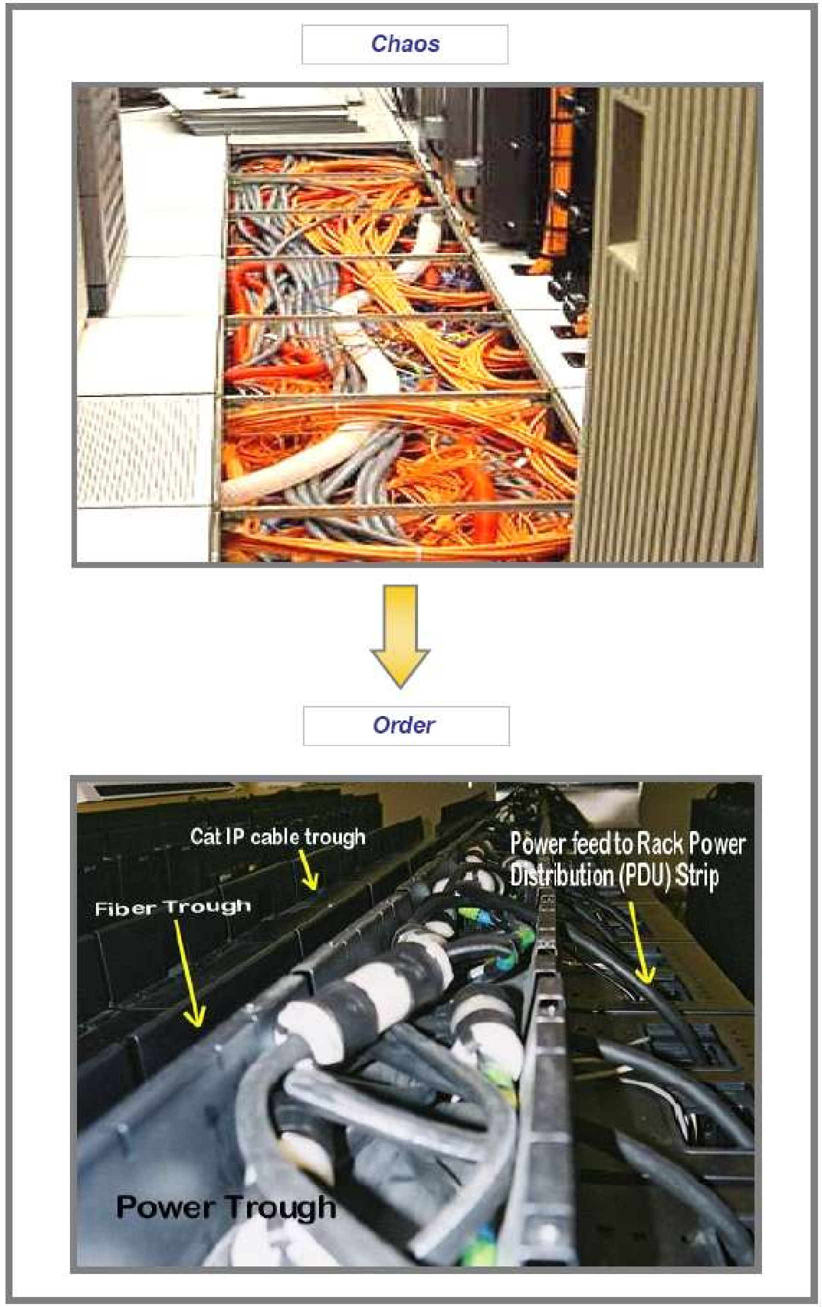

Figure 4

Overhead cabling provides a way out of the plenum mess

Installation of blanking panels

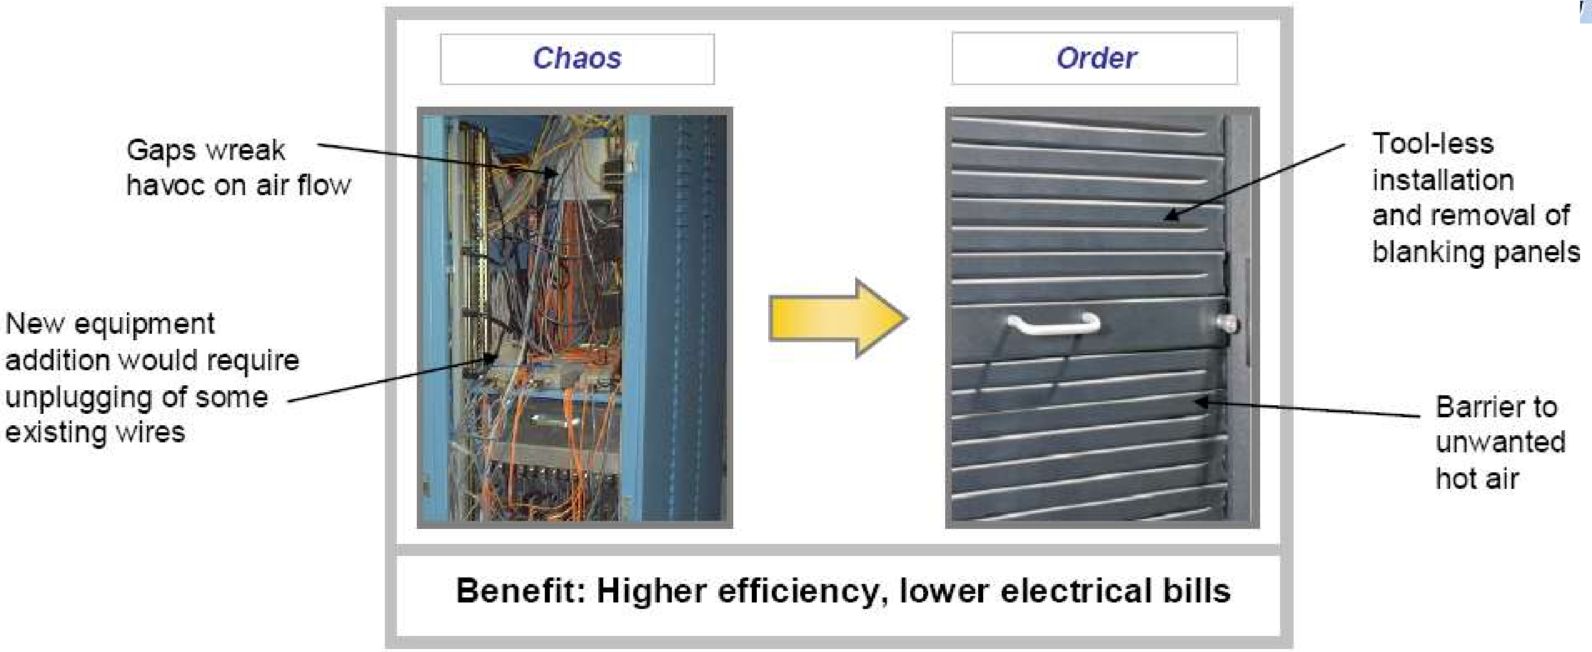

Migration to contained close-coupled™ cooling

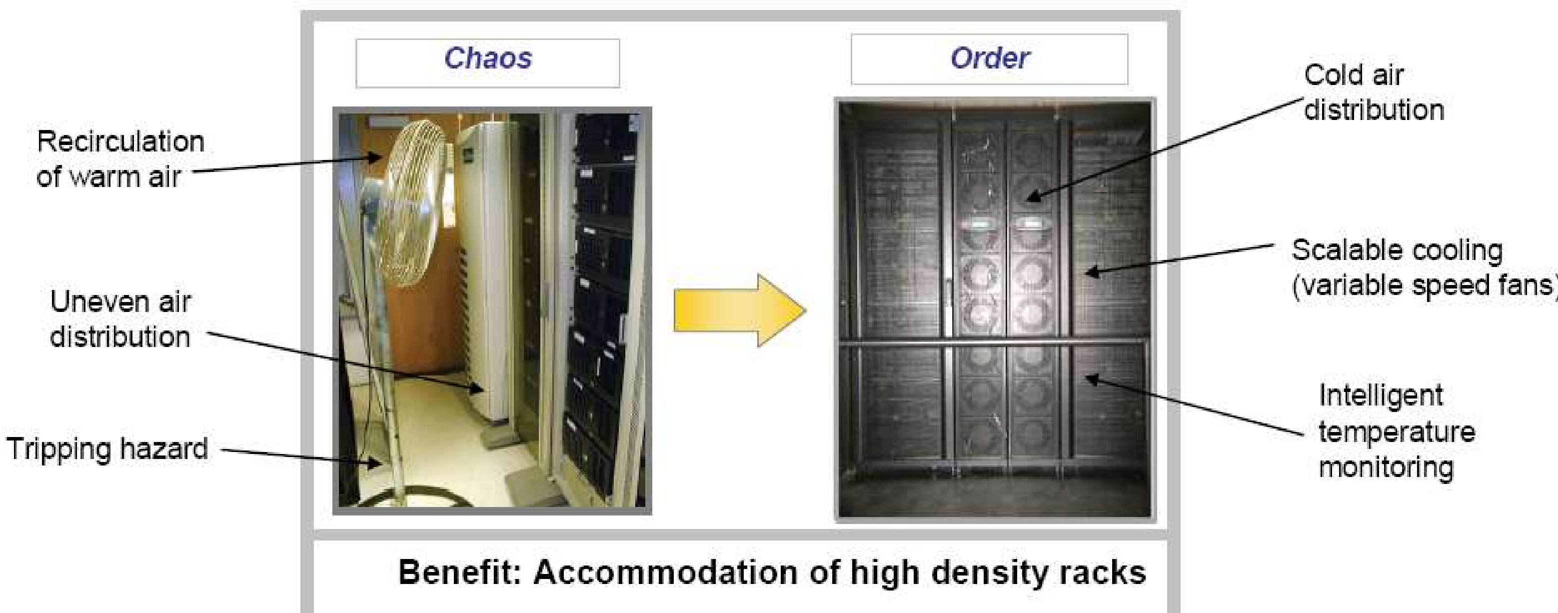

CRAC Unit Placement and Configuration – some guidelines

- If flexibility exists in the placement of the CRAC units. place them facing hot aisles rather than cold aisles, as the underfloor velocity pressure should be minimized in cold aisles.
- If CRAC units are aligned in parallel rows on a raised floor, then each row of CRACs should exhaust air in a direction that increases the static pressure across the floor rather than CRACs exhausting such that their plumes collide, causing decreased static pressure in these regions and overall loss of chilled air to the raised floor.

CRAC Unit Placement and Configuration
– some guidelines (Conti..)

- Airflow rate distribution in the perforated tiles is uniform when all the CRAC units are discharging in the same direction and this distribution is poor (non-uniform) when the CRACs discharge air in collision with each other.
- Turning vanes and baffles appeared to reduce the CRAC airflow rate by about 15%. It is, thus, preferable that turning vanes (scoops) not be used in CRAC units. However, when turning vanes are used in CRAC units facing each other, their orientation should be such that the airflow from the CRAC is in the same direction.

CRAC Unit Placement and Configuration – some guidelines (Conti..)

- Racks that have a clear path of hot air back to the intakes of the CRAC units generally show low rack air temperatures.
- Integrating sophisticated thermal instrumentation and control of a data center environment with the operational parameters of CRAC units e.g., volumetric airflow rate or chilled-air set point temperature, can result in significant energy savings of around 50%. Variable-frequency drives can be used to change fan speeds and, thus, CRAC airflow rates, and the chilled air set point temperatures can be changed via controlling the condenser conditions of the CRAC unit.

Effect of Aisle Spacing

- some guidelines

- The most commonly found aisle spacings are 4 ft (1.22 m) aisles that are consistent with two floor tiles for the hot aisle and cold aisle.
- With the increased server rack powers, Beaty and Davidson (2005) suggest that the aisle widths be increased to allow for more air in the vicinity of the IT equipment and at a possible reduced velocity.

Effect of Aisle Spacing
- some guidelines

- Limiting the velocity of the air supply by the perforated tiles has two benefits
  – high-velocity air tends to blow by the inlet grilles of the servers.
  – high-velocity air near the center of the aisle tends to blow by the intake of the servers and out the top of the aisle, not serving any cooling benefit.
- Wider cold aisles will increase chilled air to the servers and lower the velocity exiting the tiles, thereby eliminating a potential "blow by" of the high-velocity chilled air.

Partitions: At the End of Aisles, Above Racks, Within Racks, and Between Racks

- Placed ceiling partitions may help prevent recirculation of the hot exhaust air into the inlets of the servers, but consideration in the placement of these partitions must take into account local fire codes.
- To prevent recirculation within a rack, it is important to install blanking panels for those areas of the rack that do not have servers installed and also in those areas where there is a clear path to the front of the servers of hot air from the rear. Gap partitioning between the different units in a rack is helpful in reducing rack inlet temperature.
- Gaps between low-powered racks may be acceptable, since not much airflow is needed for low-powered racks and most if not all the air can come from the perforated tiles.

Partitions: At the End of Aisles, Above Racks, Within Racks, and Between Racks

- Gaps between high-powered racks that are arranged in a hot aisle/cold aisle fashion need to be eliminated. They can cause significant infiltration of hot exhaust air from the racks directly into the neighboring cold aisles. This can lead to higher rack inlet temperatures by as much as 6°C.
- For data centers with lower-powered racks (1.5 kW/rack), the facility's air-conditioning efficiency can be improved by about 5% by restricting the rack air intake opening to the bottom one-third potion of the front of the rack. Thus, the location of the intake and exit openings can influence energy efficiency.
- For data centers with lower-power racks (1-3 kW/rack), using fully populated racks results in less hot exhaust air recirculation than using partially filled.

Partitions: At the End of Aisles, Above Racks, Within Racks, and Between Racks

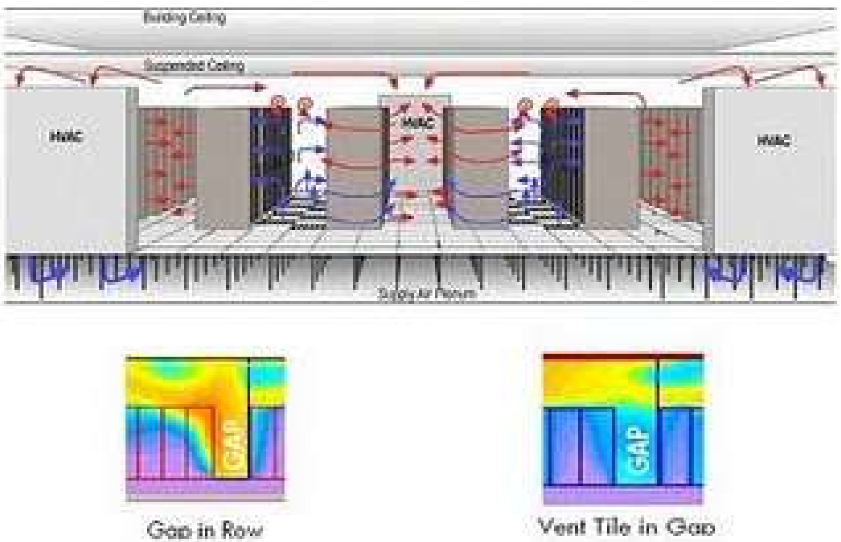

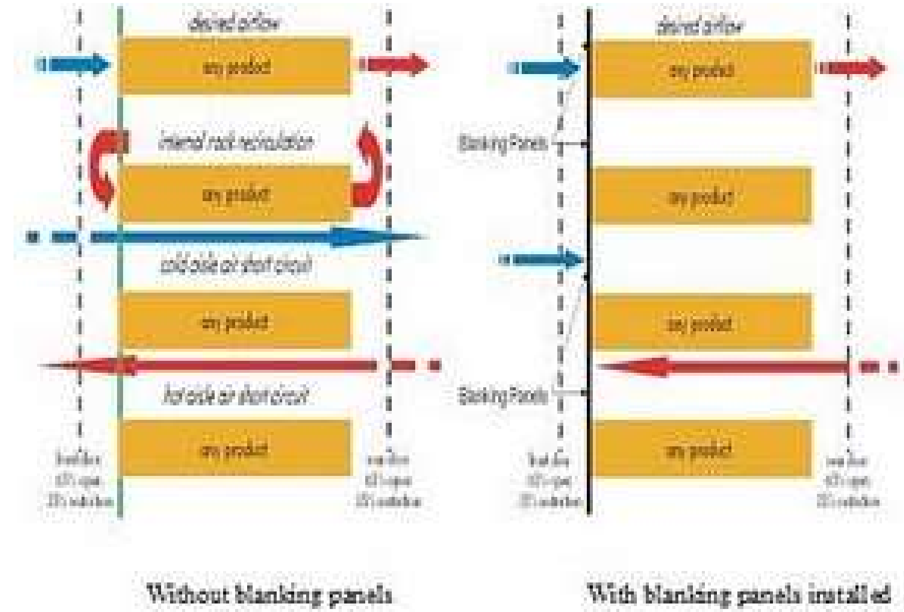

Schematic of the effects of gaps and blanking panels

Localized Cooling

- Placing the cooling near the source of heat shortens the distance the air must be moved. This increases the capacity, flexibility, efficiency, and scalability of the cooling.

NON-RAISED-FLOOR DESIGNS

EFFECT OF VARIOUS PARAMETERS

- Use ~30° diffusers to reduce the temperatures at the top of the racks, and do not use diffusers (blowing air straight down) if the temperatures at the bottom of the racks need to be cool.
- Using the air supply diffuser close to the top of the racks helps the bottom of the racks, and a larger clearance helps the tops of the racks.
- High rack heat loads can mean high rack flow rates and better mixing, and this can sometimes drop the inlet air temperature to the rack.
- The air at the top of the rack is usually hotter than at the bottom, even though this effect is not as pronounced as in underfloor air supply.

DATA CENTER ENERGY MANAGEMENT AND EFFICIENCY

- Use an air-side economizer when outside air is cooler than data center hot exhaust air to the CRAC unit.
- Using centralized air handlers means using larger sized, more efficient fans. Centralized units can also respond better to the varying thermal needs of different parts of the building and can also simplify the implementation of air-side economizers.
- Optimize the refrigeration plant using higher building chiller water set points, variable-flow evaporators and staging controls, lower condenser water temperature set points, high-efficiency variable-frequency drives in chiller pumps, and thermal storage units to handle peak loads.

DATA CENTER ENERGY MANAGEMENT AND EFFICIENCY, Conti..

- Localize liquid cooling of racks to augment air-handling capabilities of existing cooling infrastructure.
- Use free cooling via a water-side economizer to cool building chilled water in mild outdoor conditions and bypass the refrigeration system.
- Use humidification control to prevent over-control of humidity and use a wide set point range of 30%-70% RH. Humidity control is very energy intensive and is not really needed in current day servers. It is also difficult to sustain due to susceptibility to sensor drift. Waste heat in the return airstream can be used for adiabatic humidification.
- A common control signal can be used to ensure all CRAC units are set to the same humidity set point.

- Use high-efficiency power supplies and do not oversize the units. This could result in savings.
- Power supplies need to meet the requirements of the Server System Infrastructure (SSI) Initiative, which is an industry-wide initiative to promote standardization of interfaces between server components such as the board., the chassis, and the power supplies.
- Use high-reliability generation units as the primary power source with the grid as backup, and also use waste heat recovery systems such as adsorption chillers. This can allow the elimination of backup power sources such as uninterruptible power supplies (UPSs).
- Use high-efficiency UPS systems. For battery-based power backup, use as high a load factor as possible with at least 40% or higher of the rated capacity. This may require the use of smaller battery-operated UPS systems in parallel.
- Use power conditioning to operate the system in the most line-efficient mode for line reactive systems.

Effect of ancillary equipment

APC white paper #139

- ancillary equipment: any stand-alone equipment not arranged in standard alternating hot aisle / cold aisle rows.
- In typical data centers, this equipment includes devices such as switches, storage equipment, networking gear, and power distribution equipment. This equipment is generally lower in power density and accounts for approximately 5-10% of the IT load power for most data centers; however, depending on the business application, (e.g., a company whose primary application is data storage), this can be significantly higher.

Figure 2

Sample data center layout with row cooling and ancillary IT equipment

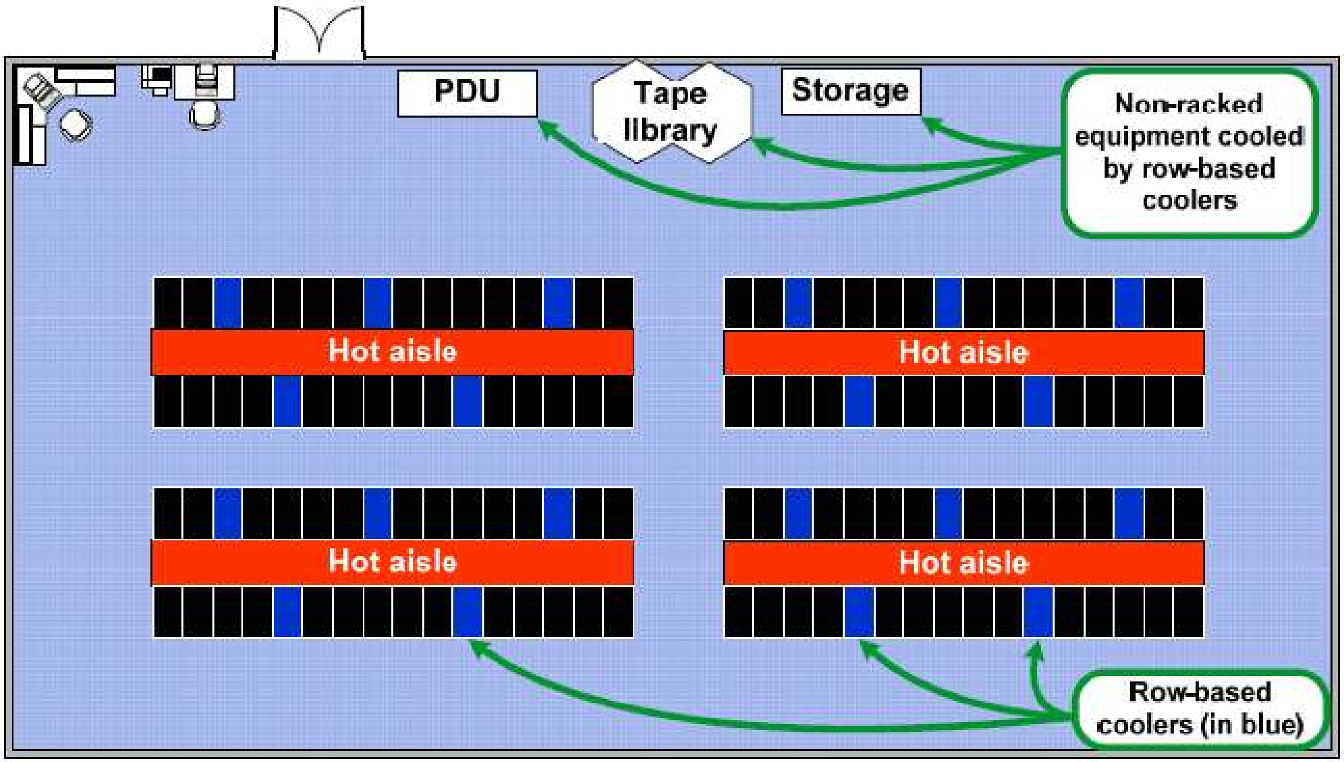

- The temperature distribution remains fairly uniform after the ancillary equipment is added. While there are the expected hot areas at the exhaust locations of the ancillary equipment, the intake air remains relatively cool as denoted by the light blue areas. The ancillary equipment is being effectively cooled by the row cooling units.

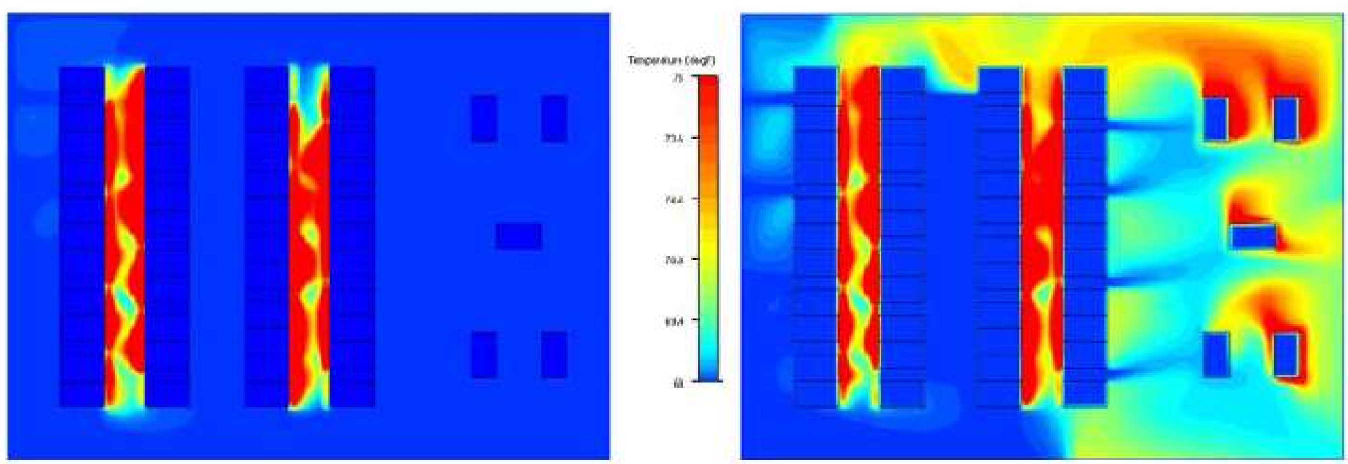

Without Ancillary Load | With Ancillary Load

Figure 4

CFD images of data center with and without ancillary load. With ancillary load, temperatures are different but reasonably uniform over respective row-based and ancillary areas.

- The placement of the ancillary loads have an impact on the temperature of the ancillary area.
- When the room is oriented like option "d", the ancillary equipment is more isolated (in an alcove) from the main row-based area. In this case, the aspect ratio, or ratio of the width (W) to length (L) has a significant impact on the temperature in this area.
- For the other three layouts, the aspect ratio has negligible impact.

Figure 5

Placements of ancillary equipment in the data center

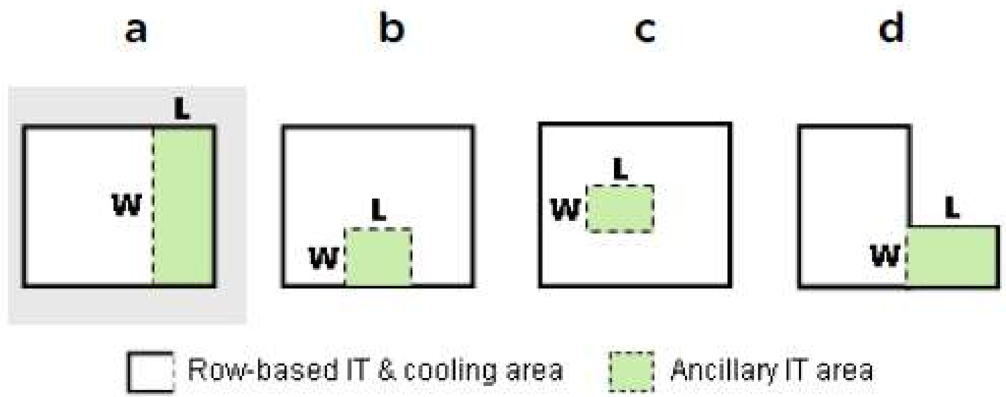

Effect of blanking Panels

- In the case of front-to-back airflow, the rack, the equipment, and the blanking panels provide a natural barrier, which greatly increases the length of the air recirculation path and consequently reduces the equipment intake of hot exhaust air as illustrated in Figure 3B. However, these key functions do not hold for equipment with side-to-side airflow.

Figure 3

Air circulation through a missing blanking panel (front-to-back airflow)

3a. (left)

Without blanking panels

3b. (right)

With blanking panels

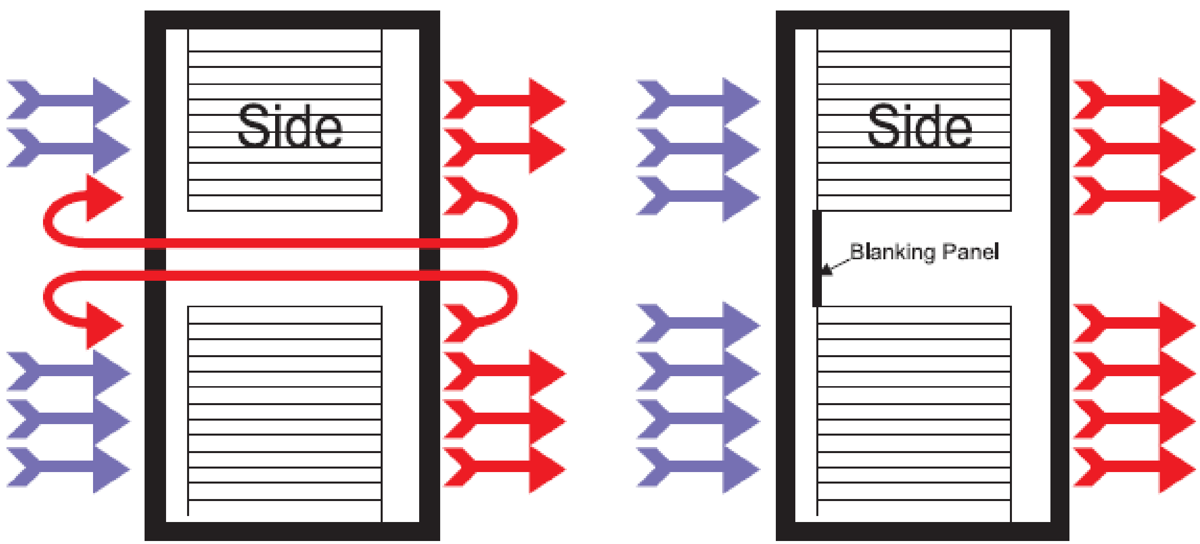

Cooling Options for Rack Equipment with Side-to-Side Airflow (APC white paper #50)

- Frequently, equipment with side-to-side airflow is mounted in open-frame racks to facilitate cooling. However, this has a very serious consequence when racks are placed in a row with equipment in adjacent racks. Under this condition equipment may have its air intake in direct alignment with the exhaust of adjacent equipment.

Figure 4

Router with side-to-side airflow

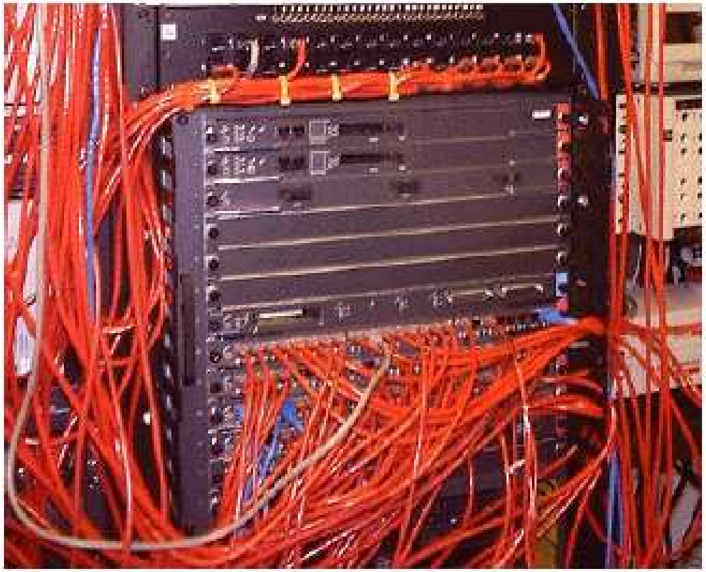

Figure 5

Equipment with side-to-side airflow in open frame racks

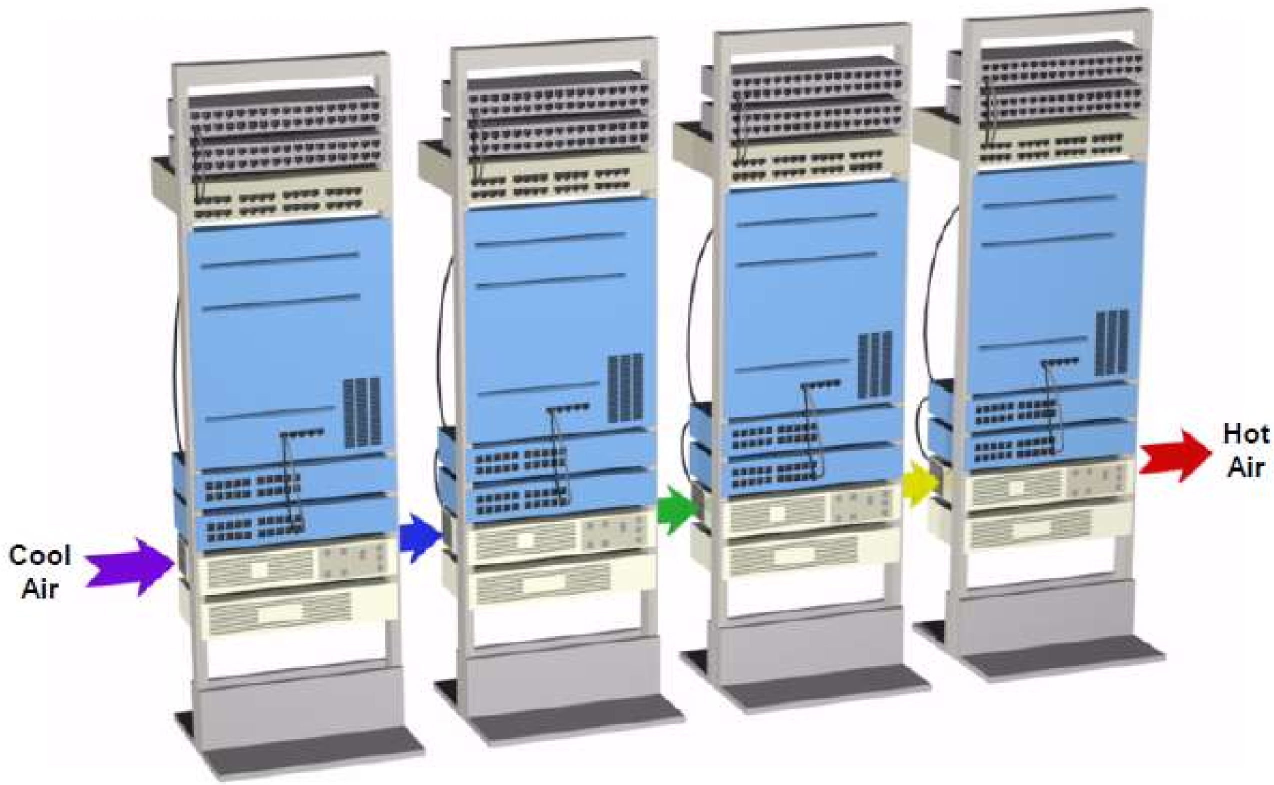

Open rack frames with increased inter-rack spacing

- Open frame racks are commonly used for equipment with side-to-side airflow. However, they do not prevent airflow from adjacent equipment from entering air intakes, and they do not provide any significant separation of the exhaust air from the intake.

Dispersion Effect

Figure 6

Dispersion effect of air on staggered equipment in open frame racks

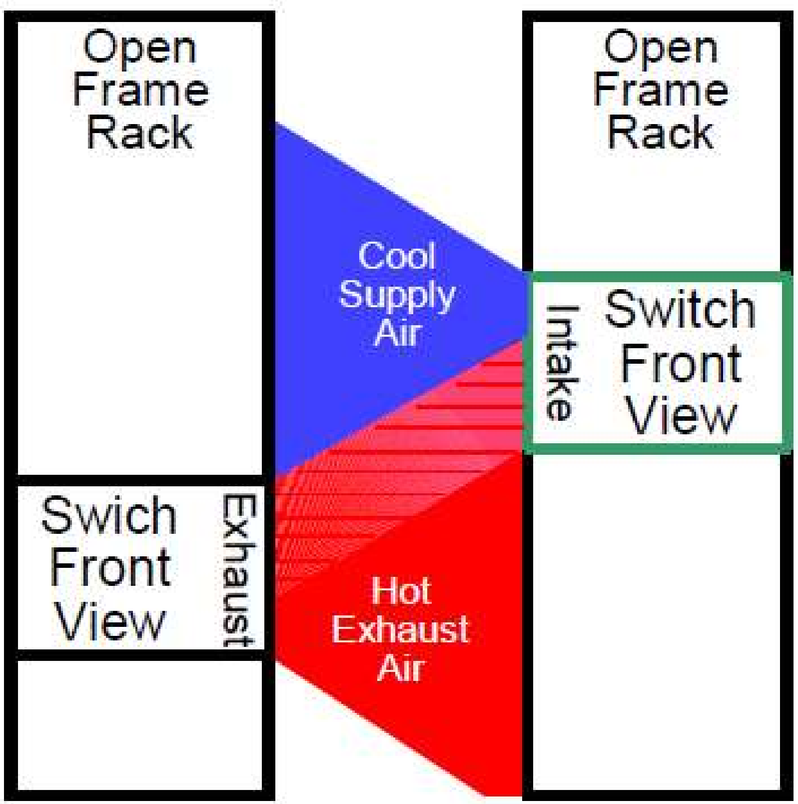

Typical side air distribution

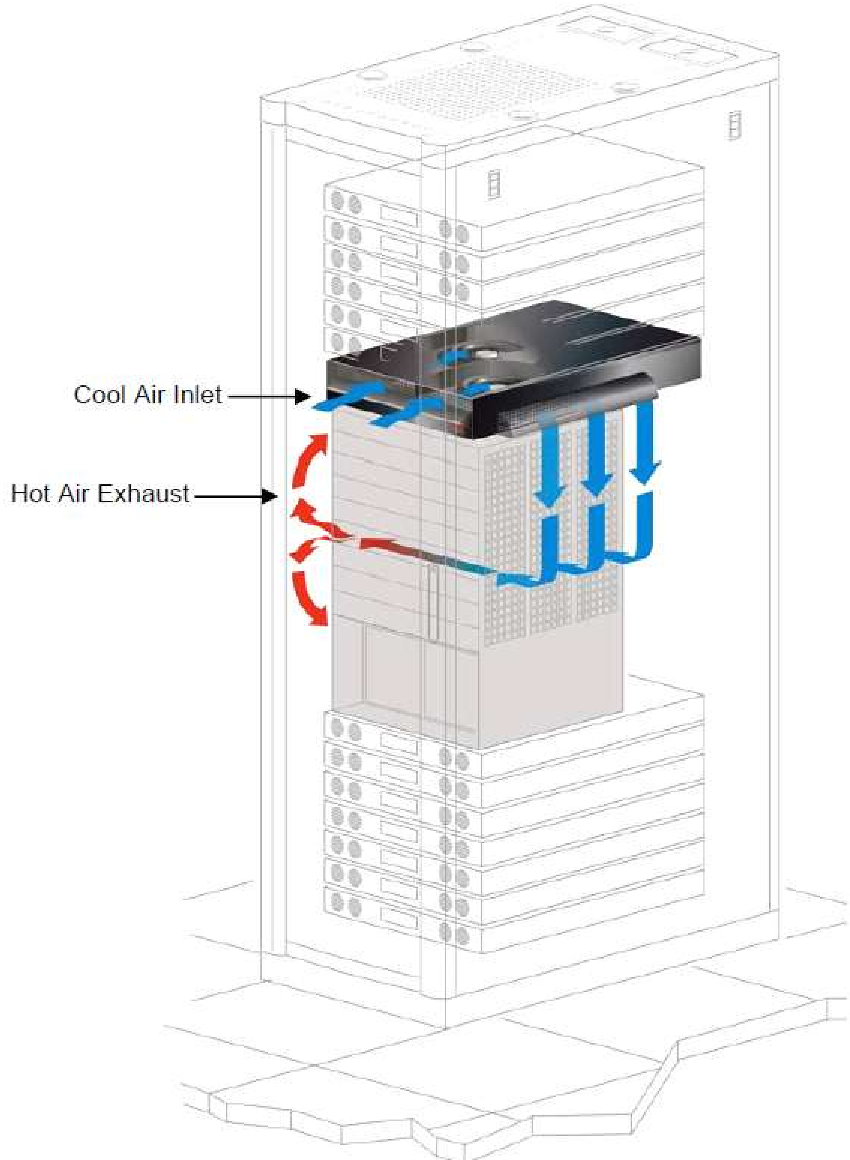

Figure 7

Top-front view perspective of side-to-side airflow supported in a single rack using an side-to-side airflow

國立交通大學
National Chiao Tung University

Table 1

Side-to-side airflow cooling methods

Side-to-side airflow variables	Side-to-side cooling methods: Open frame racks with increased inter-rack spacing	Low density enclosure	Supplemental fan	Side air distribution
Power density	•0-1 kW per rack •Heat load is distributed over larger floor area	•0-1 kW per rack •Heat load is distributed over larger floor area	•0-2 kW	•0-6 kW per rack
Floor space optimization	•Consumes more floor space	•Consumes more floor space	•Better utilization of floor space by increasing power density per rack	•Requires vertical rack space for air turning device •Floor space optimized by increasing power density per rack
Cooling system efficiency	•Mixing of cold and hot airflows results in lower efficiency •Reduction in cooling fault tolerance	•By directing the hot exhaust air into the hot aisle the cold and hot airflows are segregated, resulting in higher efficiency	•Higher efficiency if using roof fan in an enclosure •Lower efficiency if using free-standing fan and reduction in cooling fault tolerance	•Segregating cold and hot airflows results in higher efficiency
Total cost of ownership (TCO)	•Higher cooling costs •Low cost per rack •Less equipment per rack means more racks	•Lower cooling costs •High cost per rack •Less equipment per rack means more racks	•Higher cooling costs when using open frame racks	•Lower cooling costs •Less racks required results in cost savings •Less floor space required
Predictable supply of cool air at intake	•Air flow not impeded but inconsistent and hard to manage due to heat from adjacent systems	•Heating from equipment in same rack makes supply air inconsistent and hard to manage	•Can focus fans directly on equipment intake if using an open frame rack but inconsistent when using a roof mounted fan in an enclosure	•Consistent and predictable cool air delivery to the equipment intake
Side-to-side equipment reliability	•Decreases equipment reliability	•Decreases equipment reliability	•Potential for nominal conditions but unpredictable	•Provides predictable nominal conditions to maintain expected reliability
Data center planning	•Difficult to plan since open frame racks need to be spaced adequately	•Simplifies future planning by allowing standardized racks to be placed anywhere in the data center	•For rack enclosures, future planning by allowing standardized racks to be placed anywhere in the data center	•Simplifies future planning by allowing standardized racks to be placed anywhere in the data center
Rack level physical security	•No rack level physical security	•Rack level physical security	•Rack level physical security with enclosures only	•Rack level physical security
Recommended application	•Isolated racks in a low density environment	•Isolated racks in a low density environment •Data center with hot aisle / cold aisle design and front-to-back airflow environments	•Higher density environment •Hot aisle / cold aisle layout (for enclosures only)	•High density environment •Hot aisle / cold aisle layout •Converged voice / data / video networks

Note: shading indicates best performance for variable

Guidelines specific to cooling from various switch and router

- Have a minimum air space of 6 inches between walls and the chassis air vents.
- Have a minimum horizontal separation of 12 inches between two chassis is required.
- Avoid placing the chassis in an overly congested rack.
- Do not place equipment near the bottom of the rack, because it might generate excessive heat that is drawn upward and into the intake ports of the equipment above, leading to over temperature conditions in the chassis at the top or near the top of the rack.

Guidelines specific to cooling from various switch and router (Conti..)

- Never install the chassis in an enclosed rack that is not properly ventilated or air-conditioned.
- Install the chassis in an enclosed rack only if it has adequate ventilation or an exhaust fan; use an open rack whenever possible.
- Use baffles inside the enclosed rack to assist in cooling the chassis.
- Planning the proper location and layout of your equipment rack is essential for successful system operation.
- Unit is intended for installation in restricted access areas. A restricted access area can be special tool, lock and key or other means of security.

Some Thermal Measurements Guidelines

Environmental classifications

Class 1: Typically a data center with tightly controlled environmental parameters (dew point, temperature, and relative humidity) and mission critical operations; types of products typically designed for this environment are enterprise servers and storage products.

Class 2: Typically an information technology space or office or lab environment with some control of environmental parameters (dew point, temperature, and relative humidity); types of products typically designed for this environment are small servers, storage products, personal computers, and workstations.

Class 3: Typically an office, home, or transportable environment with little control of environmental parameters (temperature only); types of products typically designed for this environment are personal computers, workstations, laptops, and printers.

Class 4: Typically a point-of-sale or light industrial or factory environment with weather protection, sufficient winter heating, and ventilation; types of products typically designed for this environment are point-of-sale equipment, ruggedized controllers, or computers and PDAs.

Controlled Environment

- Control of humidity within the data center is essential to ensuring proper operation of the IT equipment.
- Humidity levels that are too low can cause a static electric discharge.
- Humidity levels that are too high can cause media failures in tape devices and other types of equipment failures

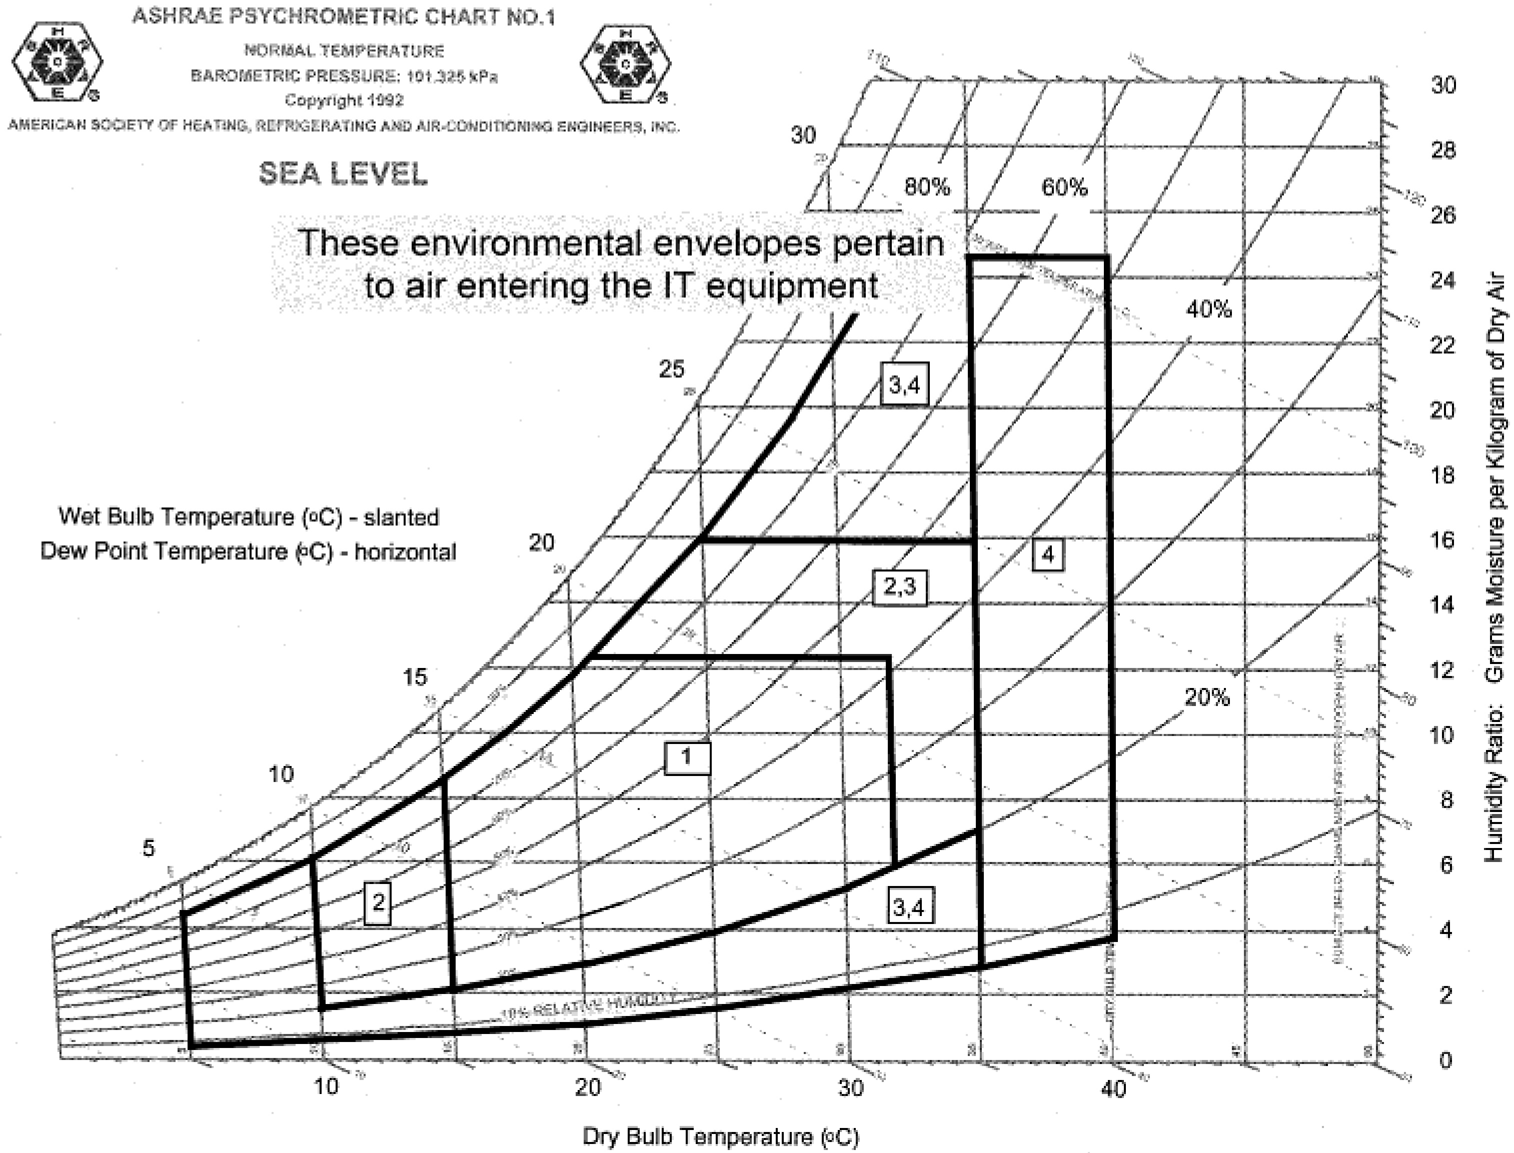

Allowable Class 1-4 operating conditions (SI units).

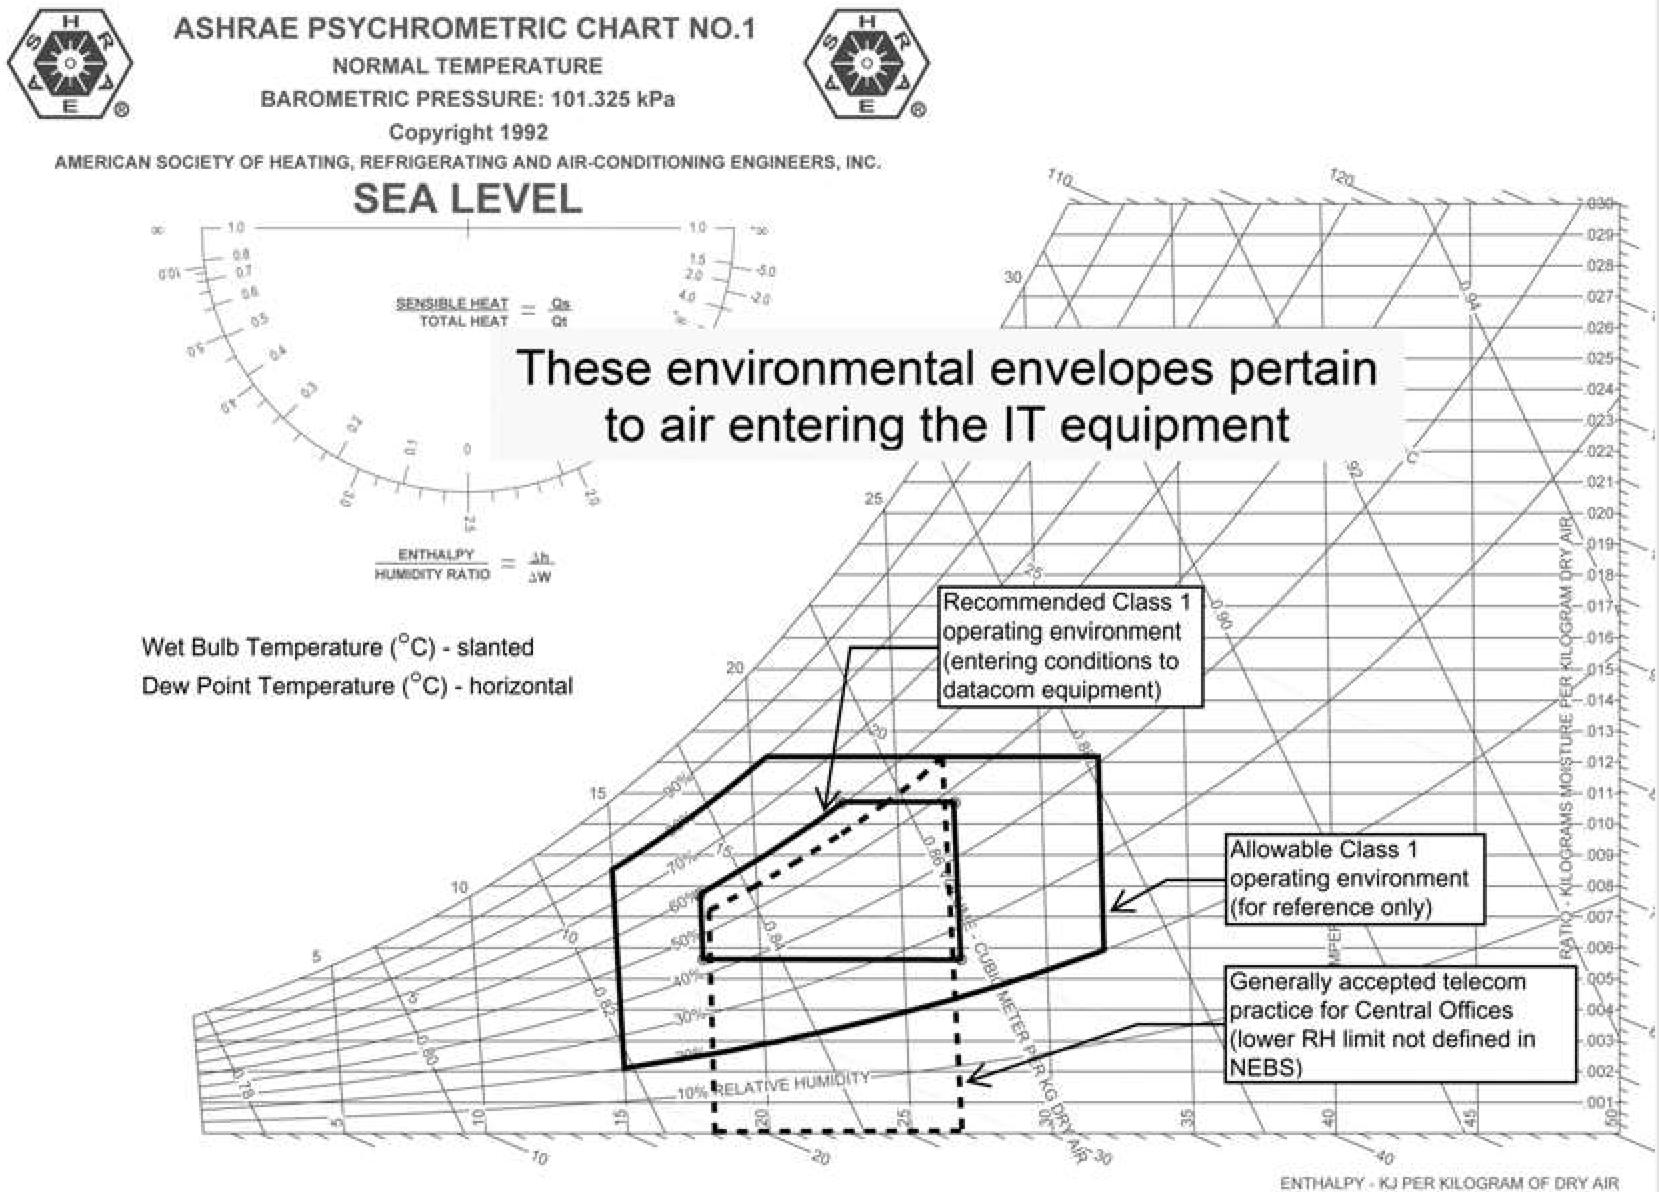

Recommended data center and central office operating conditions (SI units).

Why humidity control?

(APC white paper #58)

- The low humidity cooling air moving throughout the data center itself can be a source of static electricity every time it moves across an ungrounded insulated surface and must be guarded against by maintaining proper humidity levels.

***Table 1** – Comparison of static electrical buildups at different relative humidity levels*

Action	Static buildup at 80% RH	Static buildup at 20% RH
Walking across ungrounded raised floor tile	250 Volts	12,000 Volts
Walking across synthetic carpet	1500 Volts	35,000 Volts

HUMIDIFICATION CONTROLS

- Data centers often over-control humidity, which results in no real operational benefits and increases energy use, as humidification consumes large amounts of energy.

PRINCIPLES

- Humidification is very energy intensive. Dehumidification incurs even higher energy costs since the air is usually cooled to below 45°F to condense out water and then is reheated by an electric heater to ensure the supply air is not too cold.
- Modern servers do not require extremely tight humidity control and typical data centers cannot actually provide tight humidity control due to sensor drift.
- Centralized humidity control can keep all units serving the same space in the same humidification mode, eliminating simultaneous humidification/dehumidification common when using independent Computer Room Air Conditioners (CRACs) controls.

- Utilize adiabatic humidifiers and evaporative cooling for humidification whenever possible. Waste heat in the return air stream can be used to drive adiabatic humidification ‘for free’ when the outside air is too cold for adiabatic systems.
- Computers do not emit humidity, nor do they require ‘fresh’ outdoor ventilation air, so a well designed and controlled data center should have minimal humidity loads. Ensure outside air economizer, if present, is properly controlled to prevent unnecessary humidification loads. Provide a tight architectural envelope and optimized pressurization controls to minimize humid or dry outside air (OSA) infiltration.

Lower humidity limit

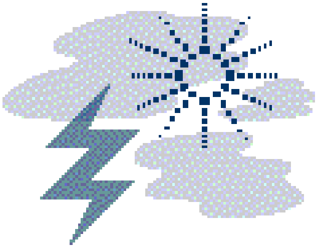

- Mitigate electrostatic discharge (ESD)
  - Recommended procedures
    - Personnel grounding
    - Cable grounding
  - Recommended equipment
    - Grounding wrist straps on racks
    - Grounded plate for cables
    - Grounded flooring
    - Servers rated for ESD resistance
  - Industry practices
    - Telecom industry has no lower limit
    - The Electrostatic Discharge Association has removed humidity control as a primary ESD control measure in their ESD/ANSI S20.20 standard
  - Humidity controls are a point of failure and are hard to maintain
  - Many data centers operate without humidification
  - This needs more research
- And for some physical media (tape storage, printing and bursting)
  - Old technology not found in most data centers
  - It is best to segregate these items rather than humidify the entire data center

Humidification

- The integration of humidification equipment into air conditioners as is commonly done is fundamentally flawed, and that humidification should be separate from air conditioning equipment and done at the room level. This is for three reasons:
  - Higher density installations may have a large number of CRAC units no matter which architecture is chosen; there is no technical need to have as many humidification units and there are many practical disadvantages, such as maintenance, of having large numbers of them.
  - When a room has a number of humidifiers it is difficult to coordinate their operation, resulting in a waste of water and electricity.
  - Cold air can accommodate less moisture and attempting to force moisture into the cold air output stream of an air conditioner is inefficient or not possible depending on saturation.

ESD control: floor grounding

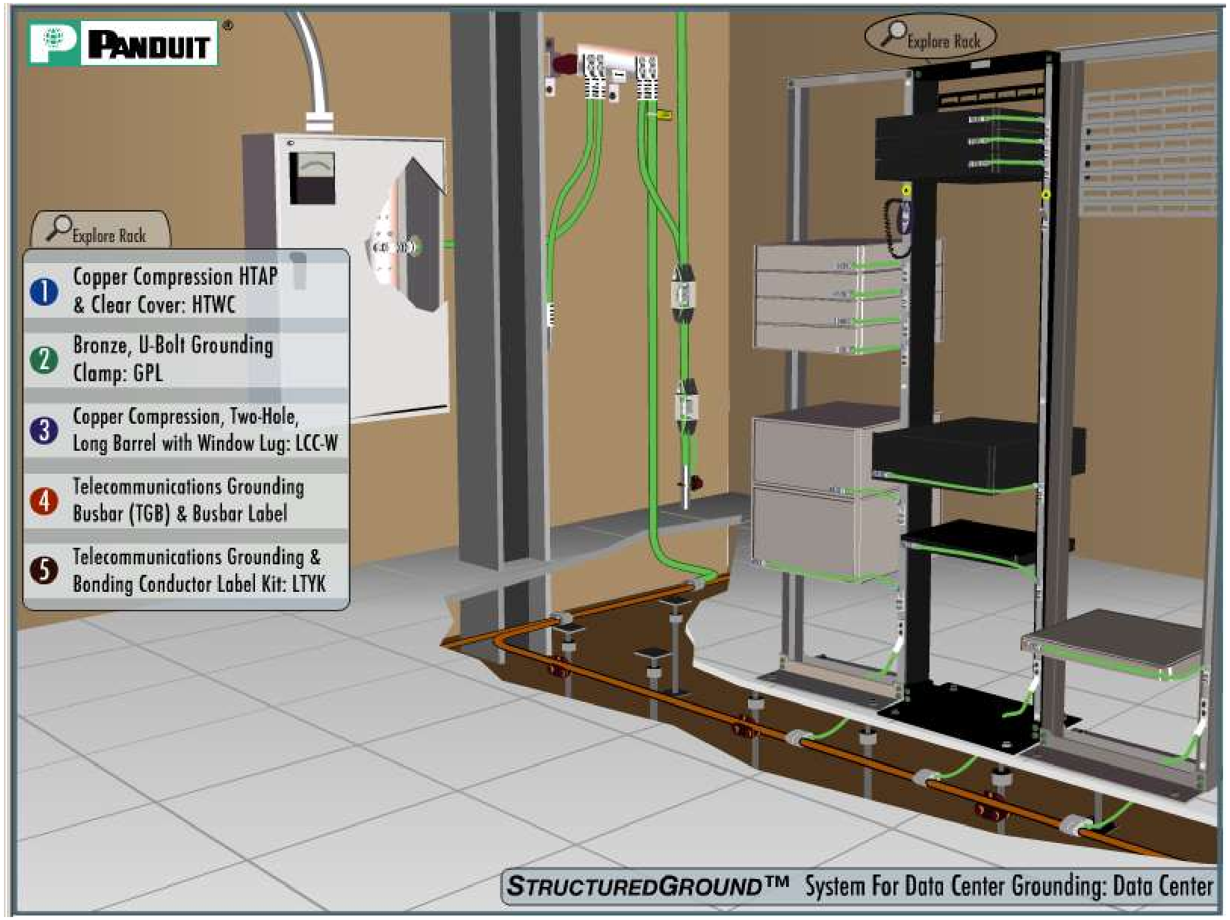

Image from Panduit, reprinted with permission

IT equipment and environment humidity guidelines

Table 2 – Humidity guidelines for IT equipment, computer rooms and data centers

Equipment or Room Type	Allowable RH range	Recommended RH range	Maximum Dew Point
Typical IT equipment	20-80%	Check product literature	Not Applicable
Wiring closets	20% - 80%	40% - 55%	70°F
Computer rooms and data centers	20% - 80%	40% - 55%	63°F

Maximum dew point temperatures are provided to establish criteria to reduce the chances of condensing humidity, especially when the IT environment is subject to rapid temperature change.

Humidity Control

- The best way to control humidity in the IT environment is to minimize the things that cause humidity levels to change and maximize the performance of the systems designed to regulate humidity.
  - Minimizing factors external to the IT environment that affect humidity
  - Minimizing factors internal to the IT environment that affect humidity

Three commonly installed types of humidification systems

- Steam canister humidifiers introduce liquid water into a canister containing electrodes. When the electrodes are powered water is boiled and steam (water vapor) is produced. The steam is introduced via a tube into the air stream to be humidified.
- Infrared humidifiers suspend quartz lamps over an open pool of water. The effect of the intense infrared light on the surface of the water is the release of water vapor that migrates into the air stream requiring humidification.
- Ultrasonic humidifiers rapidly vibrate water to create a fog or mist that is introduced into the air stream requiring humidification.

***Table 3** – Characteristics of humidifiers*

Type of Humidifier	Capital Cost	Operational Cost	Maintenance Cost
Steam Canister	Low	High	Low
Infrared	Low	High	Low
Ultrasonic	High	Low	High

Humidity Measurement

- The single most important place to maintain proper relative humidity is at the cooling air intake opening on IT equipment.

Figure 1 *– Server airflow and rack enclosure measuring point*

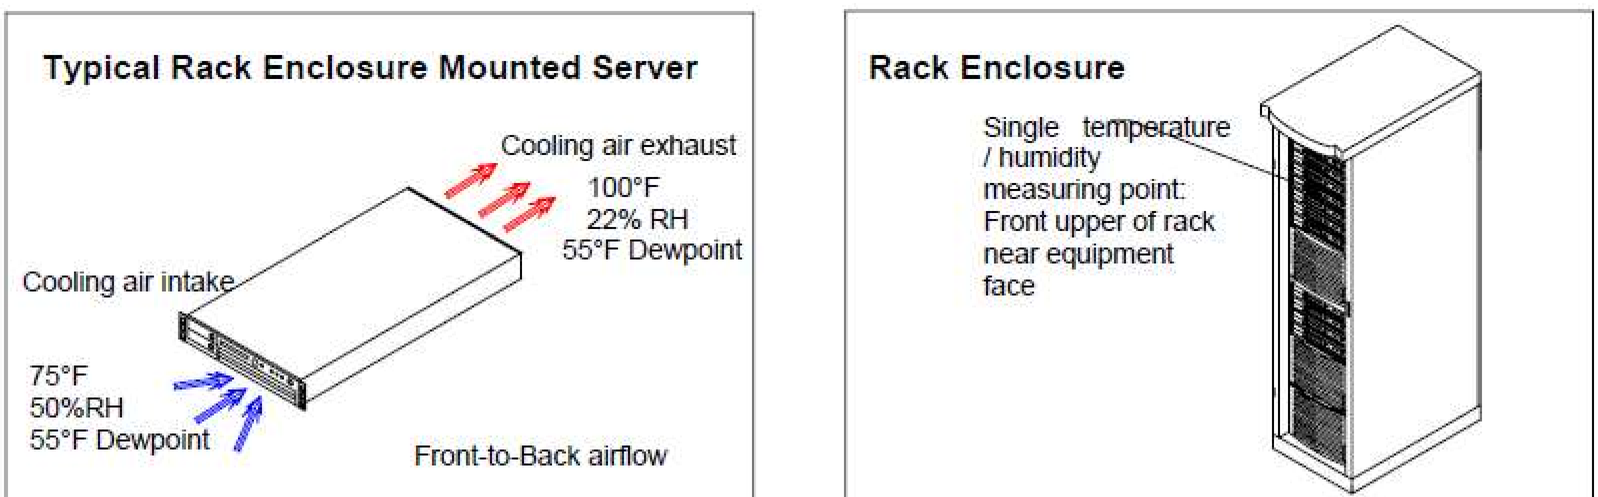

- Most computer room air conditioners and air handlers measure the humidity level of data center air as it returns into the unit from the IT environment. Figure 3 shows the monitoring points on a down flow unit.
- This data is used to control the operation of humidifiers if they are installed within the unit. The standard user interface on most precision cooling units provides readout of the relative humidity when requested.

Figure 3 – *Cooling system return air humidity monitoring*

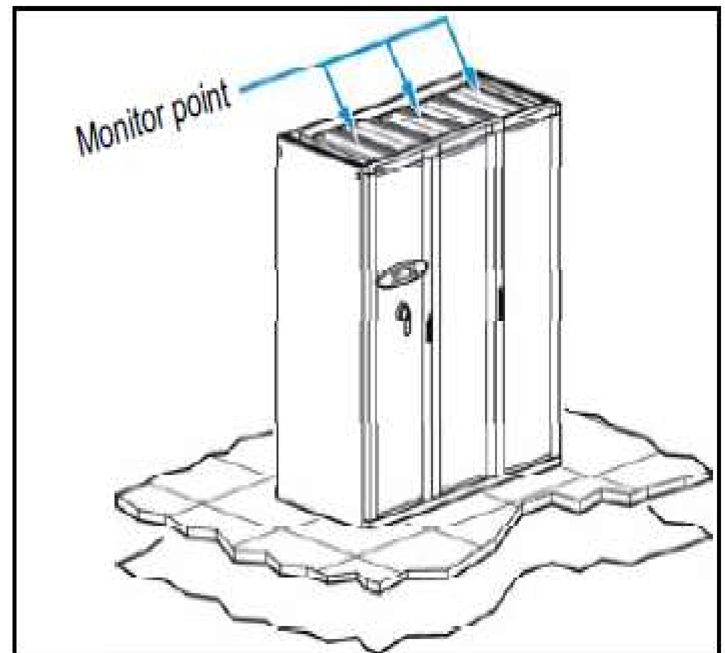

Table 3 – Characteristics of humidifiers

Type of Humidifier	Capital Cost	Operational Cost	Maintenance Cost
Steam Canister	Low	High	Low
Infrared	Low	High	Low
Ultrasonic	High	Low	High

CRAC Monitoring

APC white paper #40

Figure 2 – Supply and return temperature monitoring points

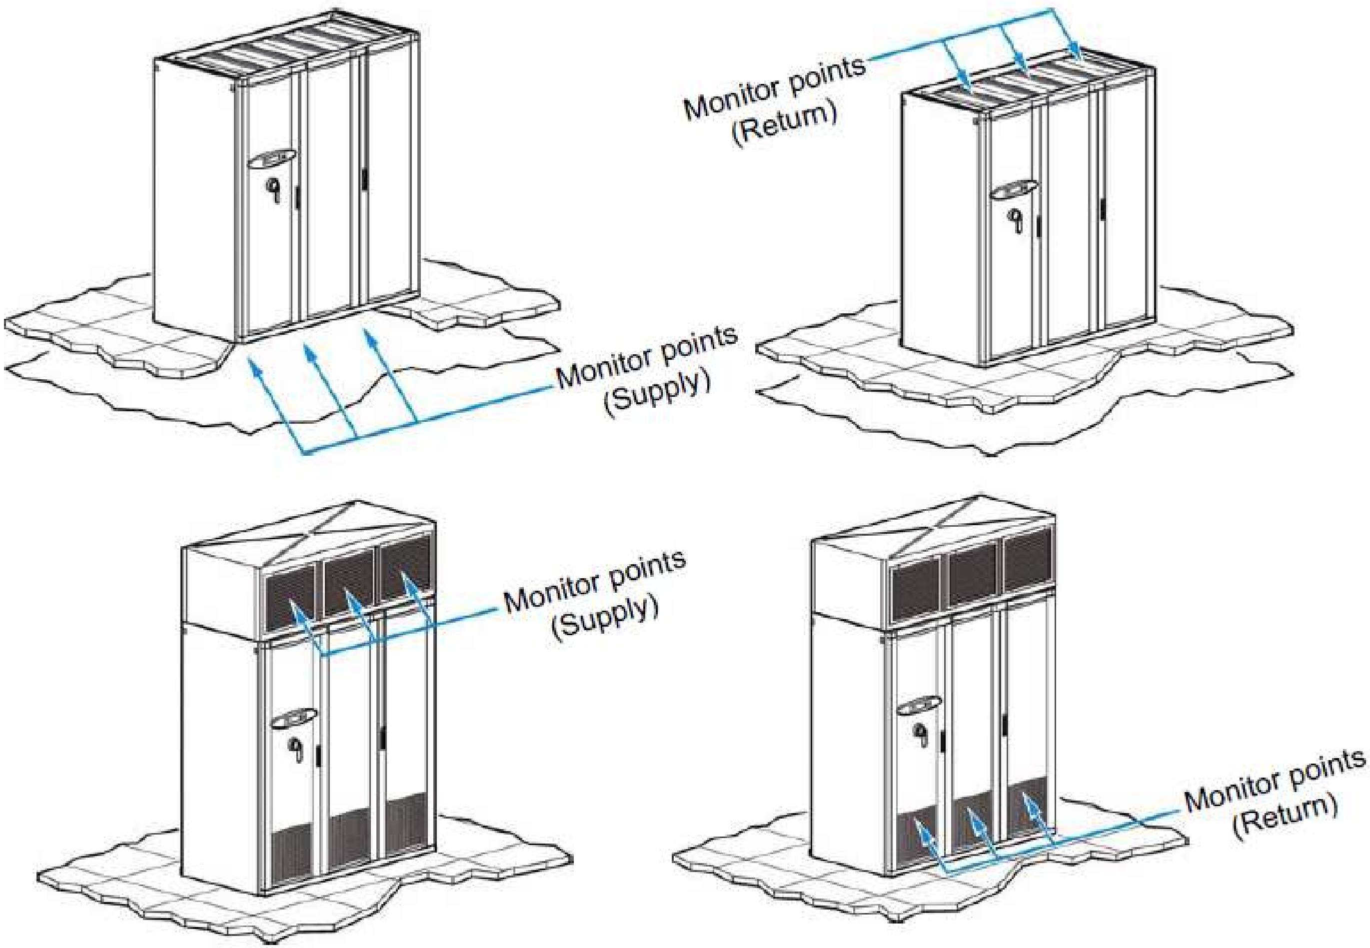

Testing the airflow of a vented floor tile

Figure 5 – Available rack enclosure cooling capacity of a floor tile as a function of per-tile airflow.

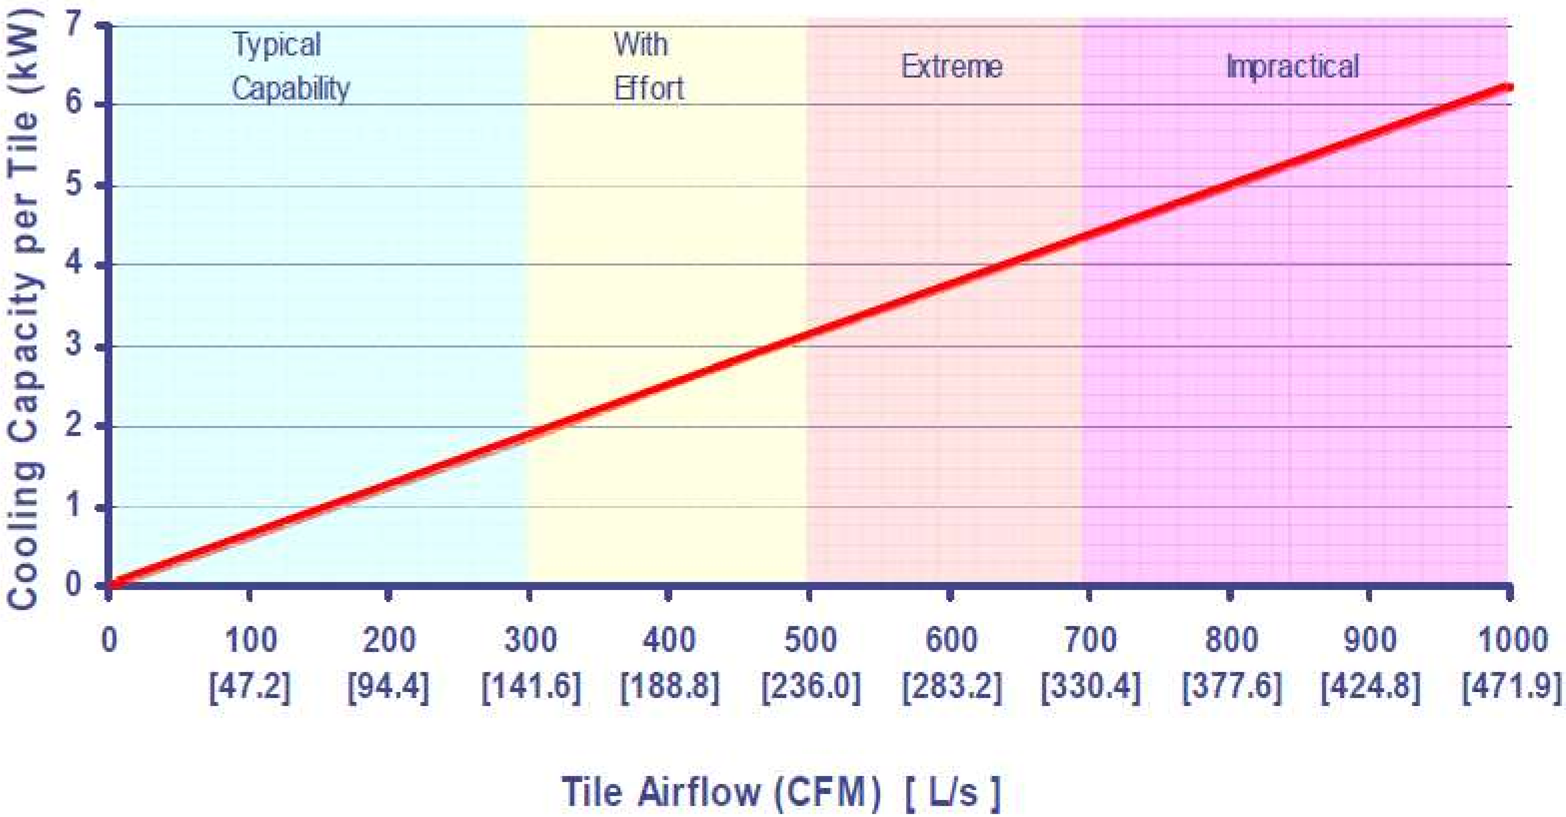

Thermal Measurements - Guidelines

Thermal guidelines for data processing environments 2nd. Ed., ASHRAE 2009

1. As shown in Figure 3.1, establish at least one point for every 10-30 ft (3-9 m) of aisle or every fourth rack position.
2. As shown in Figure 3.2, locate points midway along the aisle, centered between equipment rows.
3. As shown in Figure 3.3, where "hot aisle/cold aisle" configuration is employed, establish points in cold aisles only.[1]

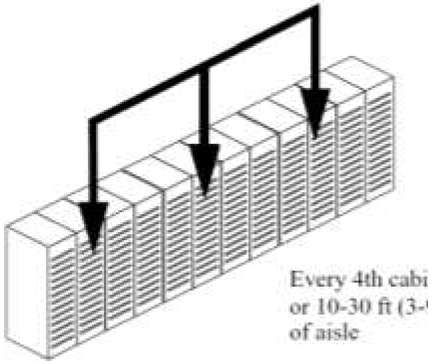

Figure 3.1 Measurement points in aisle.

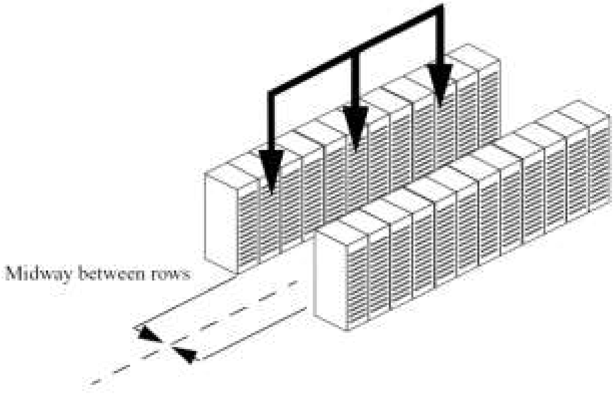

Figure 3.2 Measurement points between rows.

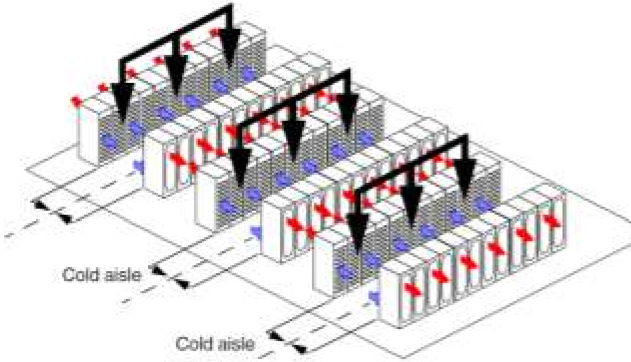

Equipment Monitoring positions

1. Measure and record temperature and humidity at the geometric center of the air intake of the top, middle, and bottom racked equipment at 50 mm (2 in.) from the front of the equipment. For example, if there are 20 servers in a rack, measure the temperature and humidity at the center of the first, tenth or eleventh, and twentieth server, as shown in Figure 3.4. For configurations with three pieces of equipment or less per cabinet, measure the inlet temperature and humidity of each piece of equipment at 50 mm (2 in.) from the front at the geometric center of each piece of equipment, also as shown in Figure 3.4.
2. All temperature and humidity levels should fall within the specifications for the class environment specified in Table 2.1. If any measurement falls outside of the recommended operating condition as specified by Chapter 2, the facility operations personnel may wish to consult with the equipment manufacturer regarding the risks involved.

Case A: For equipment that is 1U to 3U in height, arrange the monitor points as shown in Figure 3.5.

Case B: For equipment that is 4U to 6U in height, arrange the monitor points as shown in Figure 3.6.

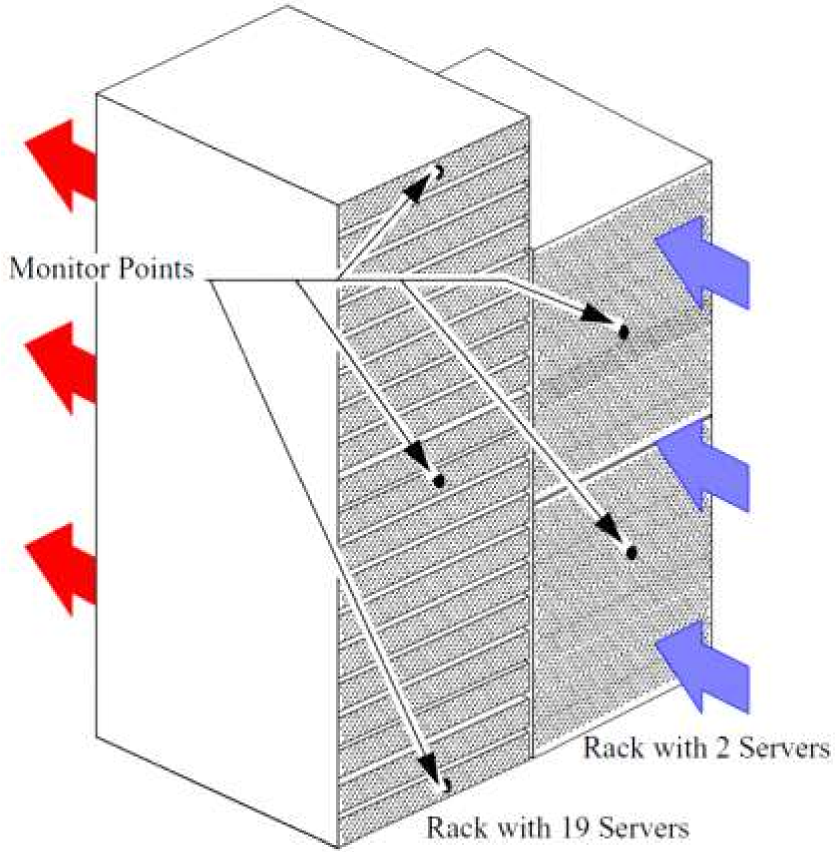

Figure 3.4 **Monitor points for configured racks.**

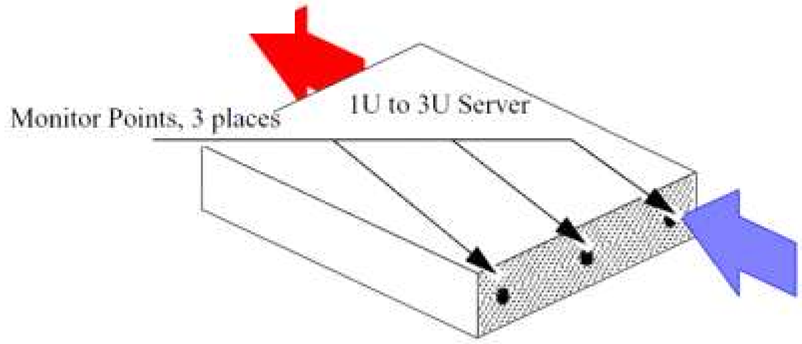

Figure 3.5 **Monitor points for 1U to 3U equipment.**

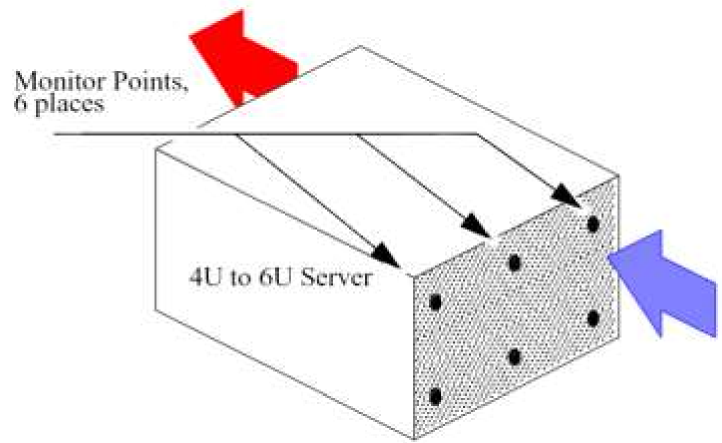

Figure 3.6 **Monitor points for 4U to 6U equipment.**

- Case C: For equipment that is 7U and larger in height, arrange the monitor points as shown in Figure 3.7.
- Case D: For equipment that has a localized area for inlet air, arrange the monitor points in a grid pattern on the inlet as shown in Figure 3.8.
- Case E: For equipment cabinets with external doors, monitor the temperature and humidity with the cabinet in its normal operational mode, which typically will be with the doors closed.

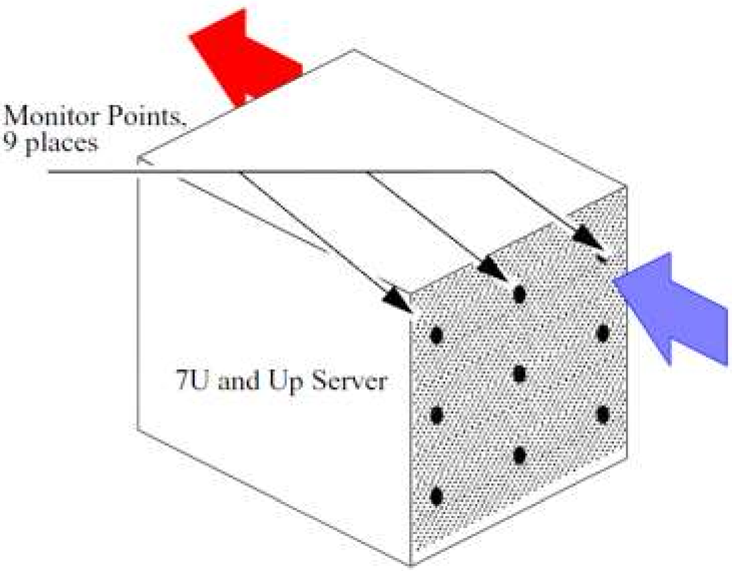

Figure 3.7 Monitor points for 7U and larger equipment.

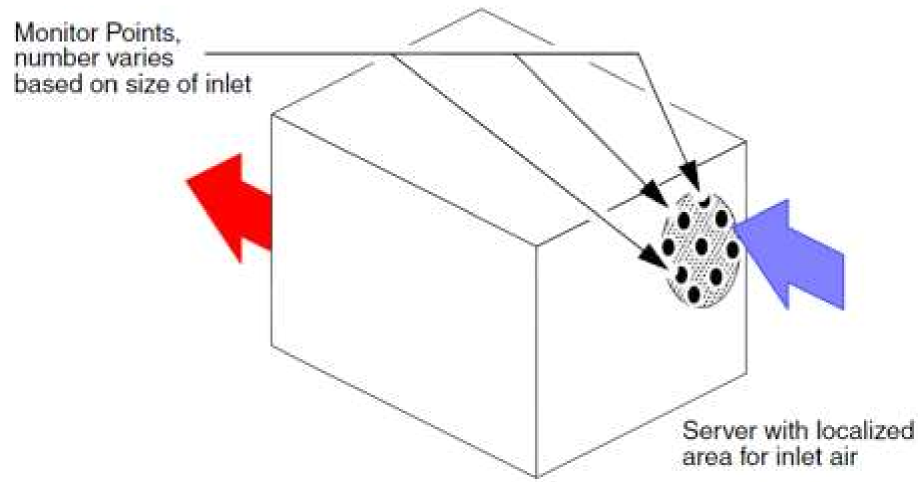

Figure 3.8 Monitor points for equipment with localized cooling.

Equipment room airflow

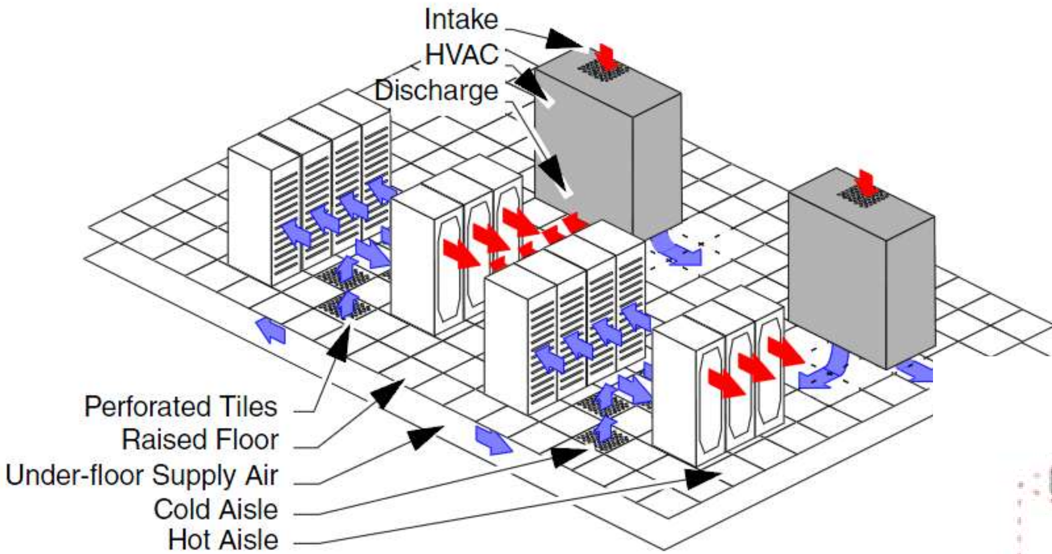

Example of hot and cold aisles with underfloor cooli

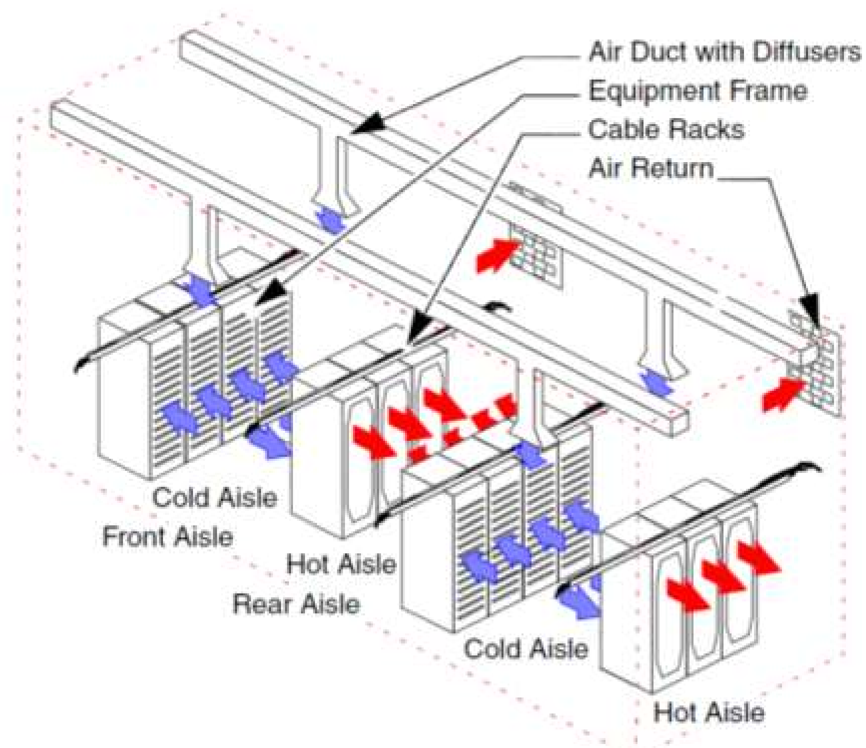

Example of hot and cold aisles for non-raised floor.

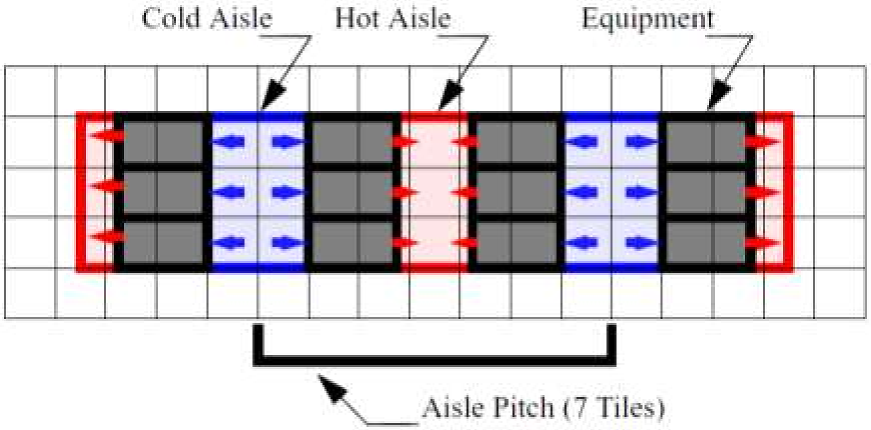

Seven-tile aisle pitch, equipment aligned on cold aisle.

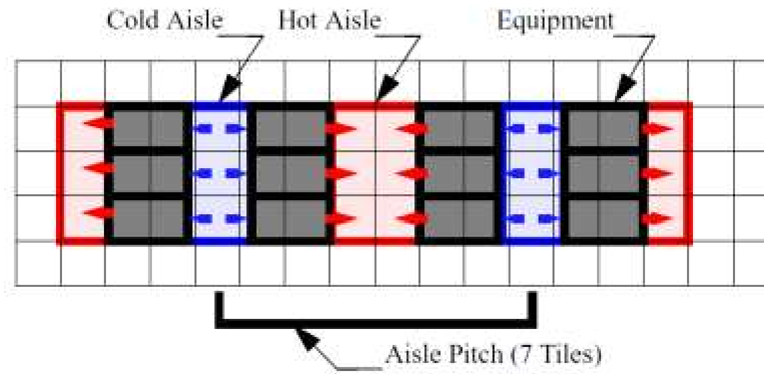

Seven-tile aisle pitch, equipment aligned on hot aisle.

Aisle Pitch Allocation

	Tile Size	Aisle Pitch (cold aisle to cold aisle)[a]	Nominal Cold Aisle Size[b]	Maximum Space Allocated for Equipment with No Overhang[c]	Hot Aisle Size
U.S.	**2 ft** (610 mm)	**14 ft** (4267 mm)	**4 ft** (1220 mm)	**42 in.** (1067 mm)	**3 ft** (914 mm)
Global	**600 mm** (23.6 in.)	**4200 mm** (13.78 ft)	**1200 mm** (3.94 ft)	**1043 mm** (41 in.)	**914 mm** (3 ft)

a. If considering a pitch other than seven floor tiles, it is advised to increase or decrease the pitch in whole tile increments. Any overhang into the cold aisle should take into account the specific design of the front of the rack and how it affects access to the tile and flow through the tile.
b. Nominal dimension assumes no overhang; less if front door overhang exists.
c. Typically a one meter rack is 1070 mm deep with the door and would overhang the front tile 3 mm for a U.S. configuration and 27 mm for global configuration.

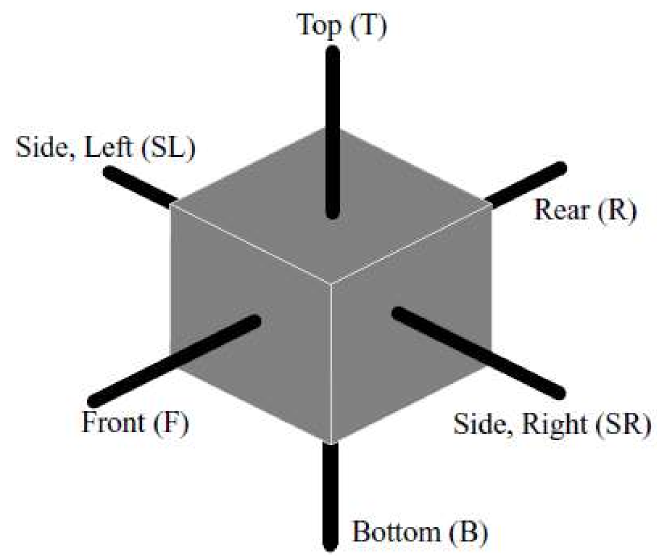

Figure 4.1 **Syntax of face definitions.**

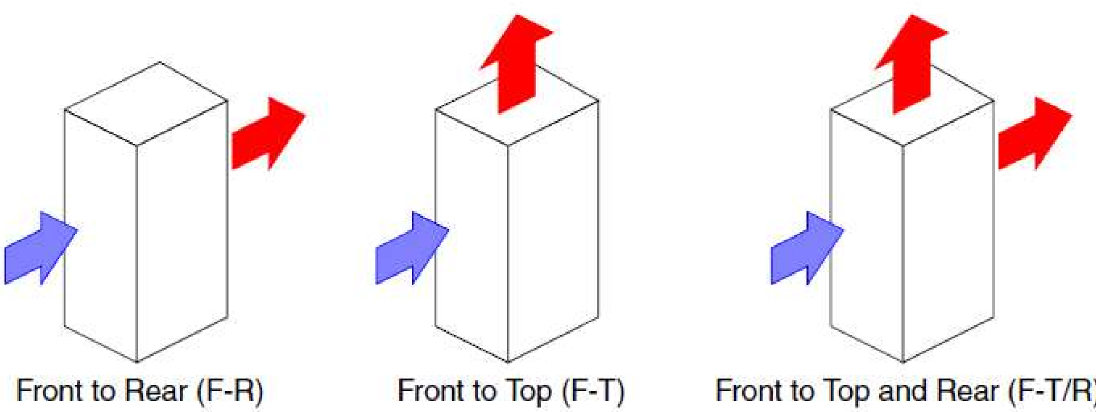

Figure 4.2 **Recommended airflow protocol.**

Figure 2.22 CRACs with return-air duct extensions.

Figure 2.23 Under raised-access floor blockage.

Some calculations/verifications from existing publications

Journal of Electronic Packaging, MARCH 2011, Vol. 133 / 011004-1

國立交通大學 National Chiao Tung University

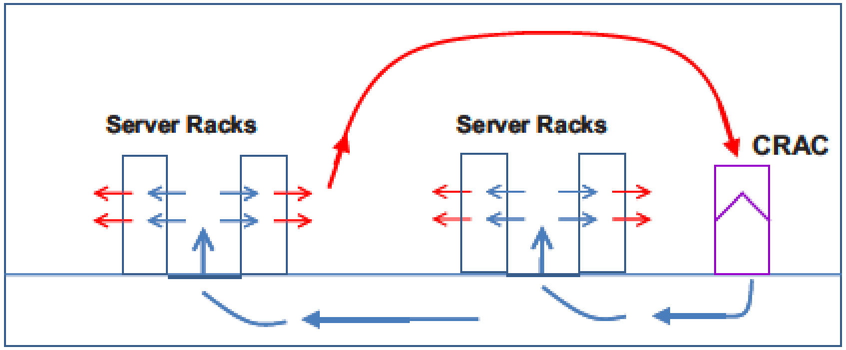

Fig. 1 A raised-floor data center

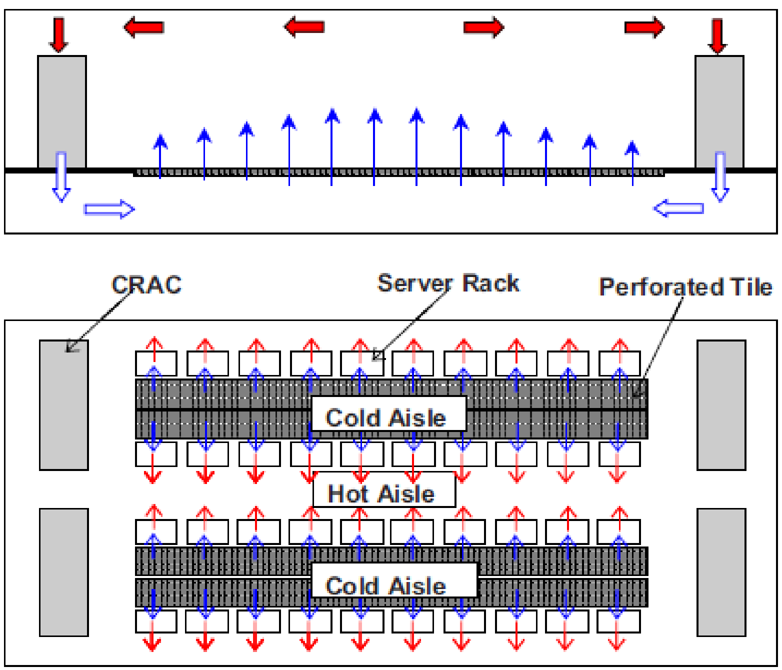

Fig. 2 The hot aisle/cold aisle arrangement

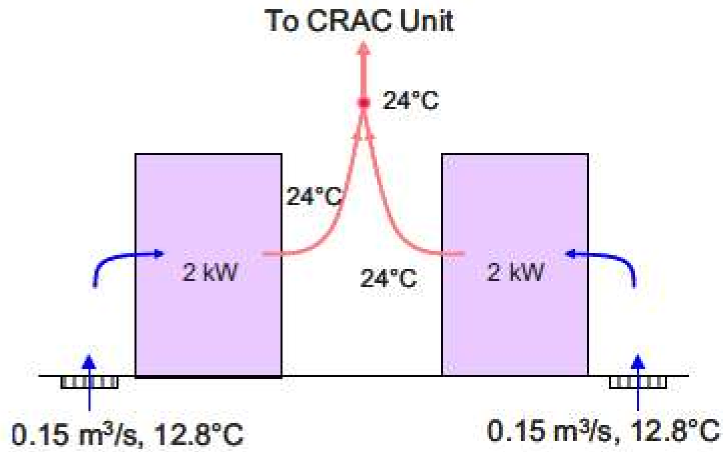

Fig. 4 Required airflow supplied

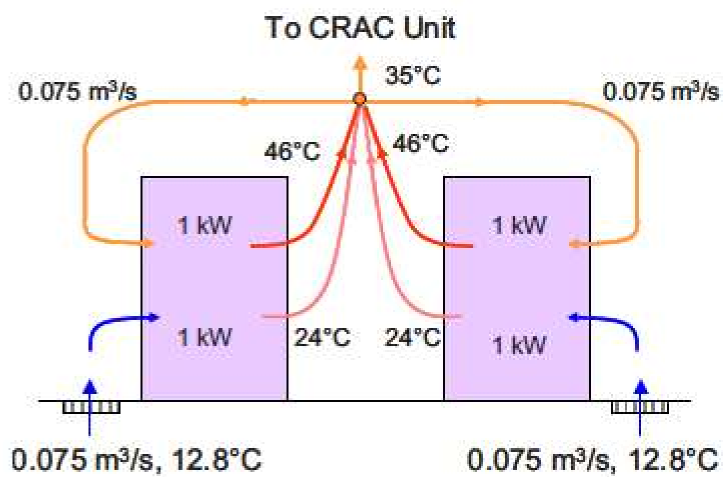

Fig. 5 Insufficient airflow

Factors Affecting the Airflow Distribution

Journal of Electronic Packaging, MARCH 2011, Vol. 133 / 011004-1

- Relationship Between the Flow Field in the Plenum and Flow Rates Through Perforated Tiles.
  - the flow rate through a perforated tile depends on the pressure drop across the tile, that is, the difference between the plenum pressure just below the tile and the ambient pressure above the raised floor.
  - plenum height
  - open area of perforated tiles

Airflow and Cooling in a Data Center

Journal of Heat Transfer, JULY 2010, Vol. 132 / 073001-1

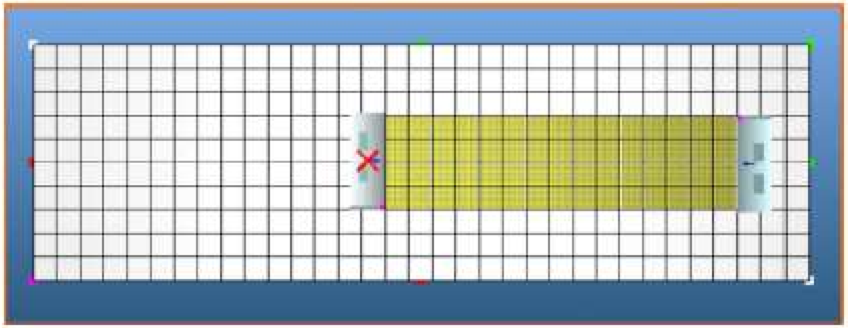

Fig. 8 Layout for a small test data center

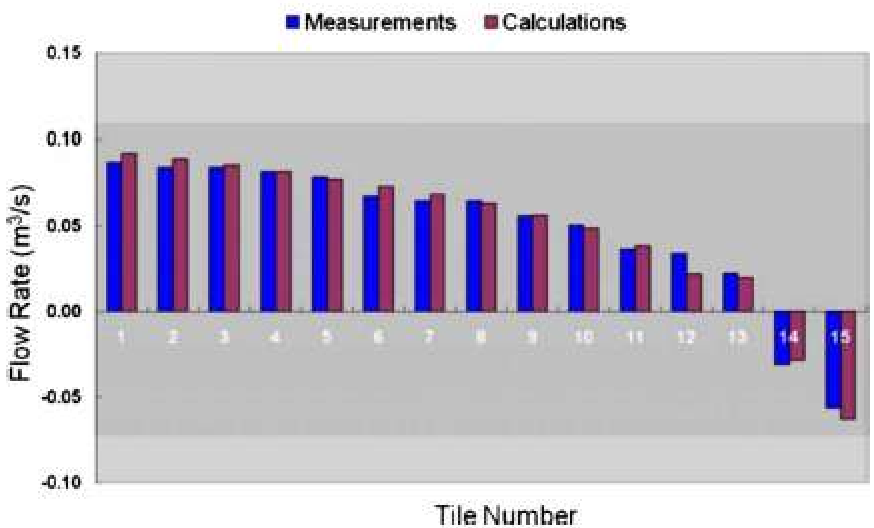

Fig. 9 Comparison of measured and calculated airflow rates (the tiles are numbered from left to right)

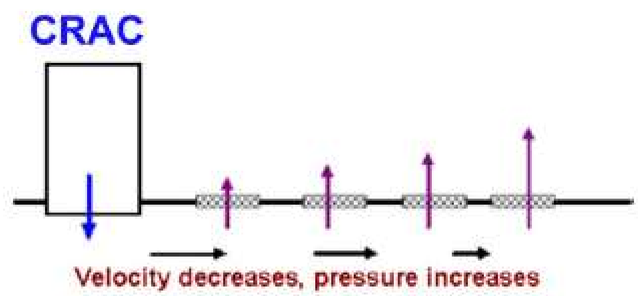

Fig. 7 Maldistribution of airflow and its cause

Base Case

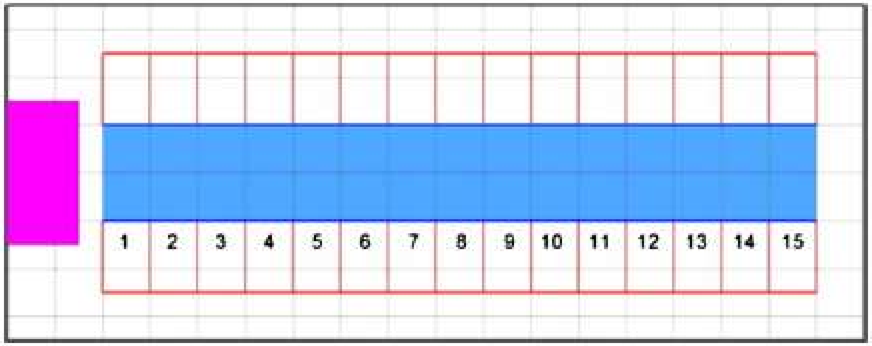

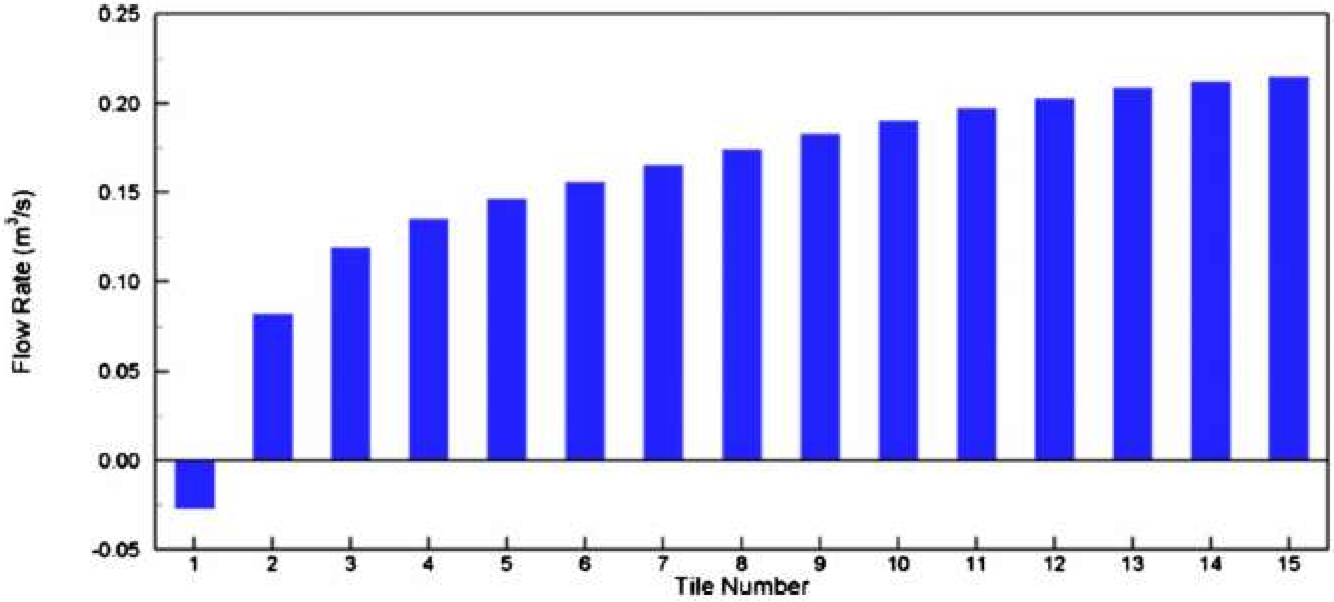

Fig. 13 Flow rates through perforated tiles for the base case

Effect of the Plenum Height

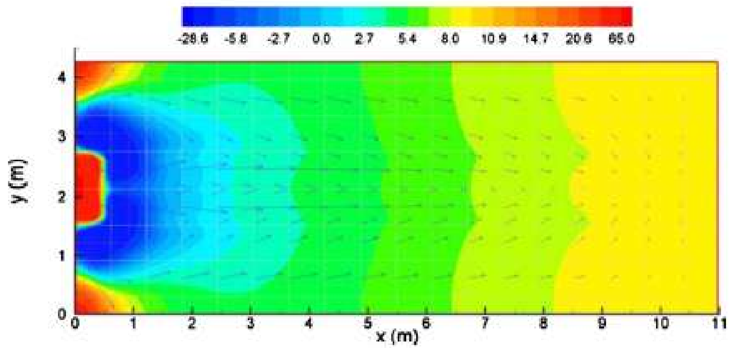

Fig. 14 Pressure distribution and velocity vectors under the raised floor for the base case (plenum height=12 in. (0.3048 m); the pressure values are in Pa)

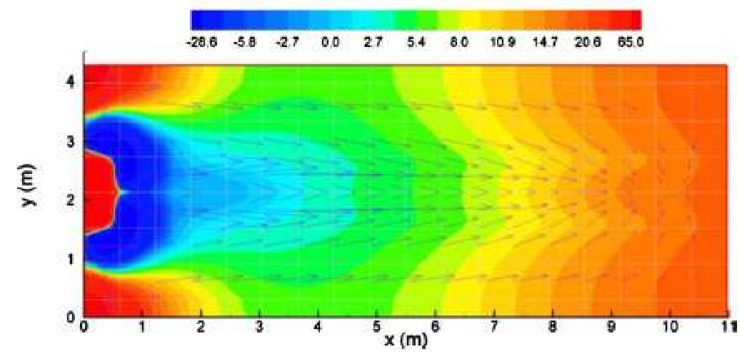

Fig. 16 Pressure distribution and velocity vectors under the raised floor for plenum height=6 in. (0.1524 m) (the pressure values are in Pa)

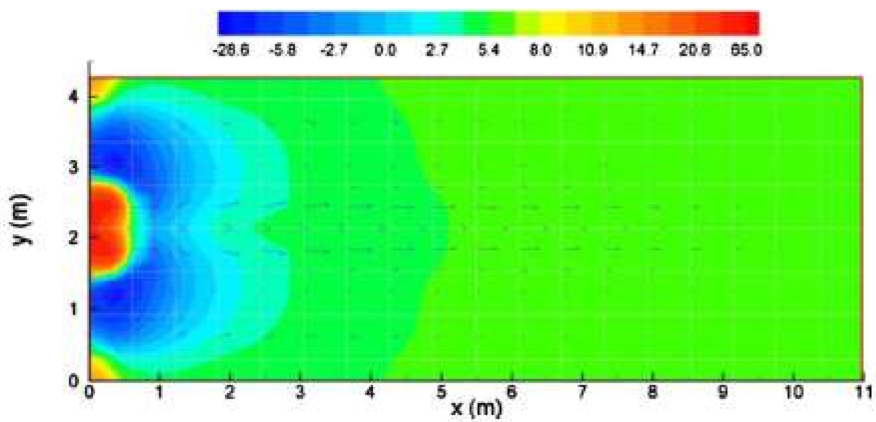

Fig. 17 Pressure distribution and velocity vectors under the raised floor for plenum height=24 in. (0.6096 m) (the pressure values are in Pa)

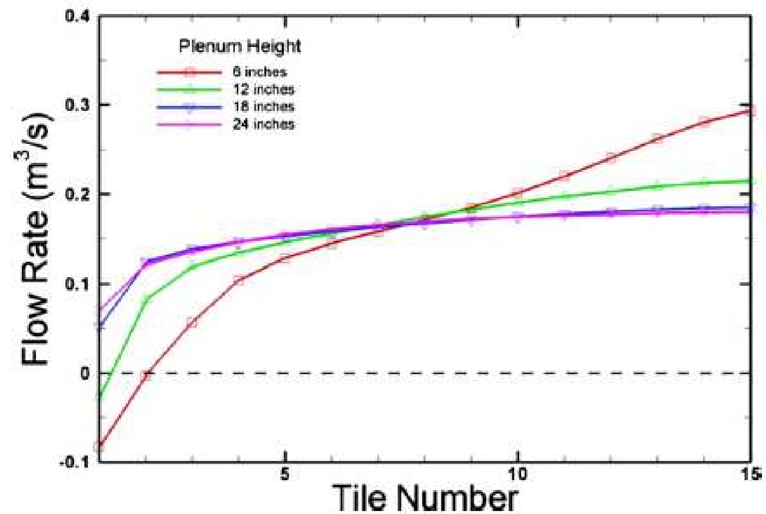

Fig. 15 Effect of plenum height on the airflow distribution

Effect of open area

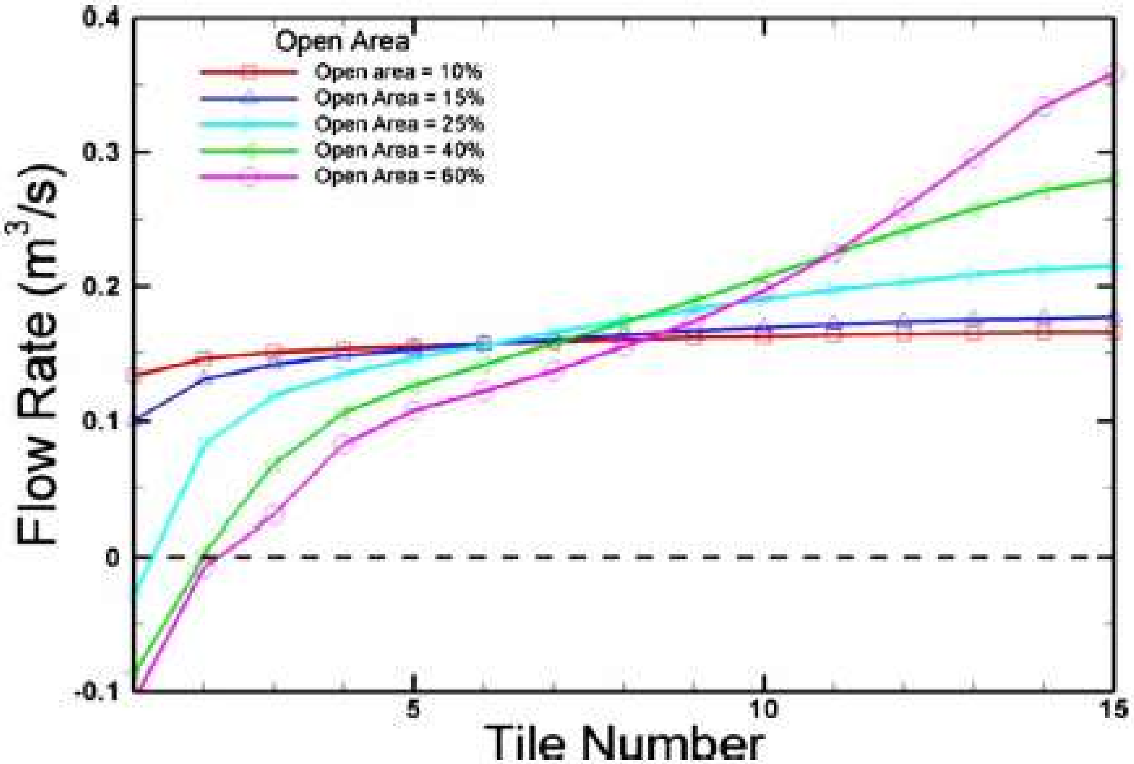

Fig. 18 Effect of open area of perforated tiles on the airflow distribution

Effect of Under-Floor Obstructions

Effect of a Circular Pipe as an Obstruction

– 6 in. pipe @ 12 in. plenum

Fig. 19 A circular pipe as an under-floor obstruction (only the centerline of the pipe is shown)

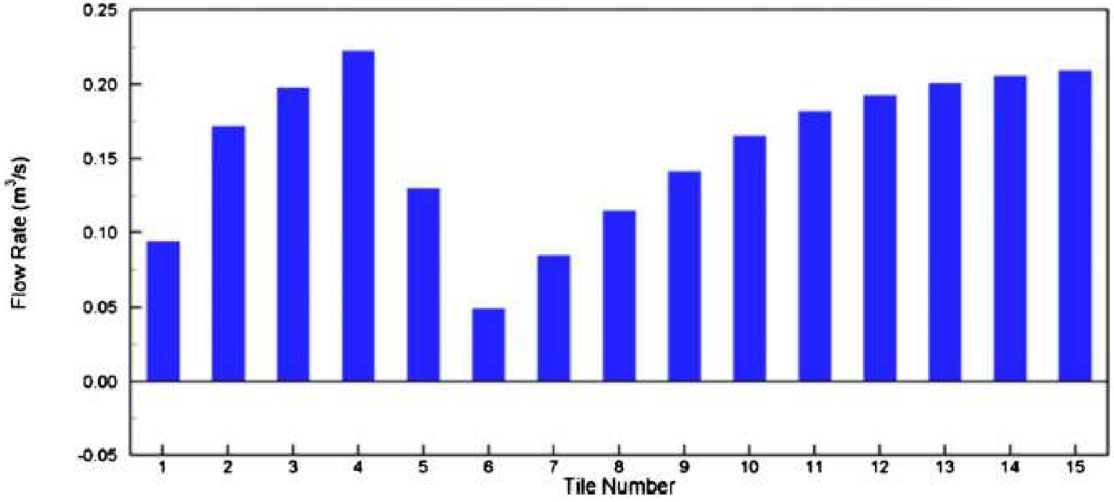

Fig. 20 Flow rates through perforated tiles as affected by the circular-pipe obstruction

Effect of Under-Floor Obstructions

Inclined partition

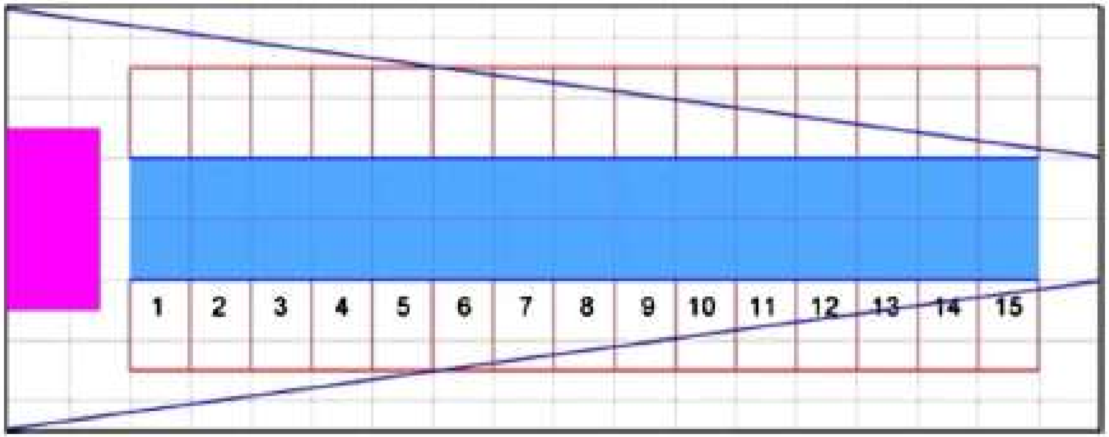

Fig. 22 Use of inclined partitions in the under-floor space

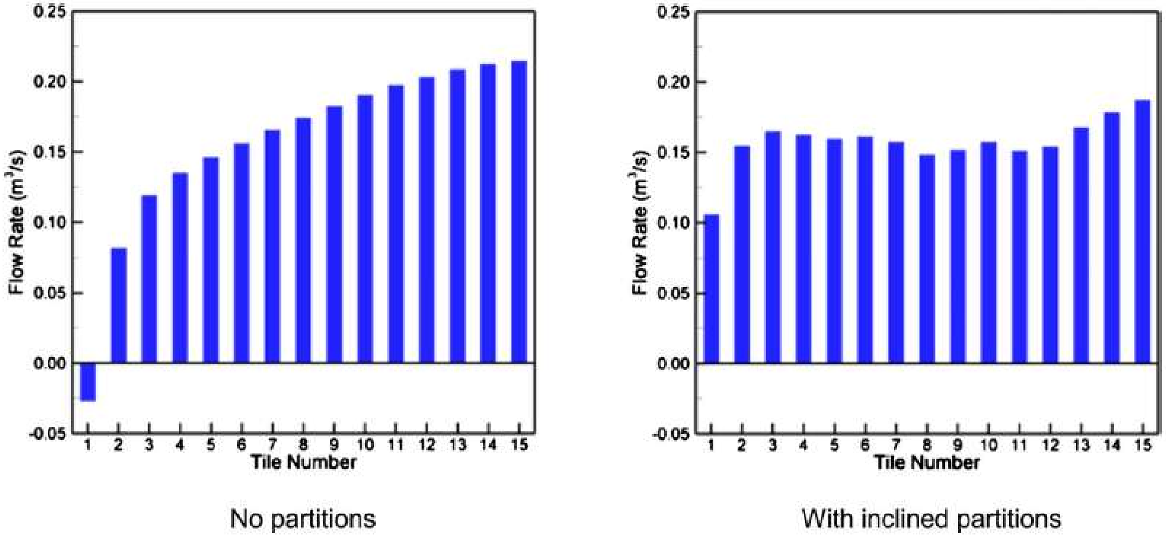

Fig. 23 Airflow distribution with and without inclined partitions

Effect of Under-Floor Obstructions

perforated partitions

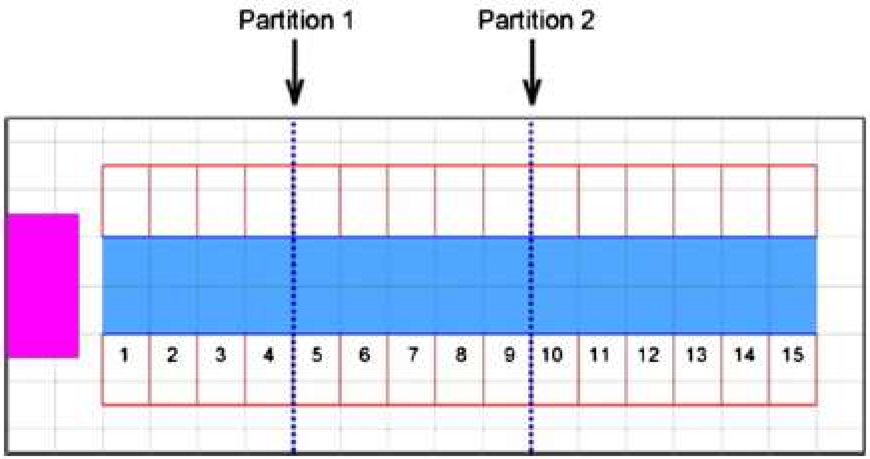

Fig. 24 Use of perforated partitions in the under-floor space

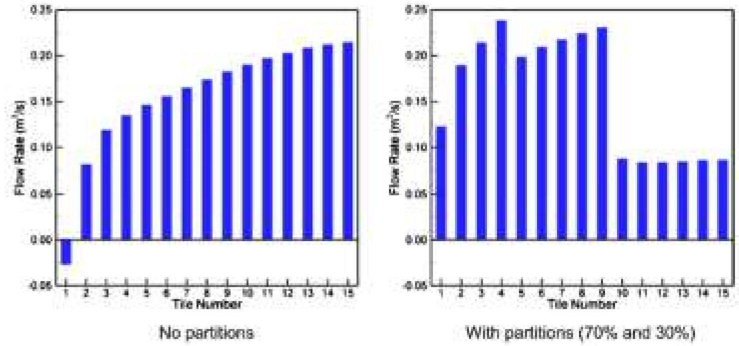

Fig. 25 Airflow distribution with and without perforated partitions (70% and 30%)

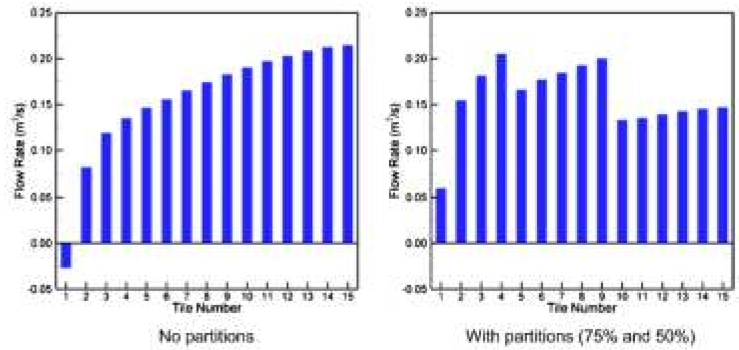

Fig. 26 Airflow distribution with and without perforated partitions (75% and 50%)

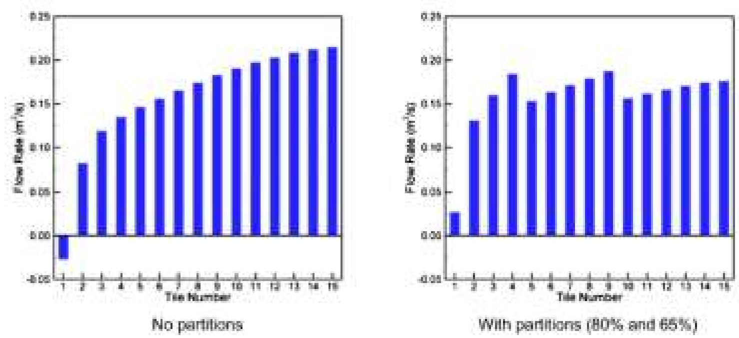

Fig. 27 Airflow distribution with and without perforated partitions (80% and 65%)

Our experiences to improve mal-distributions

Table 1. The geometric sizes of the test sections.

D mm	H (mm)	W (mm)	AR (mn)	t (mm)	b (mm)	L (mm)
3	12	12	0.442	18.5	6	120
	9	9	0.785	18.5	5	120
				13.5		110
				3.5		90
				18.5		120
	7	7	1.298	18.5	5	120
2	12	12	0.196	19	6	120
	9	9	0.349	19	5	120
	7	7	0.557	19	5	120

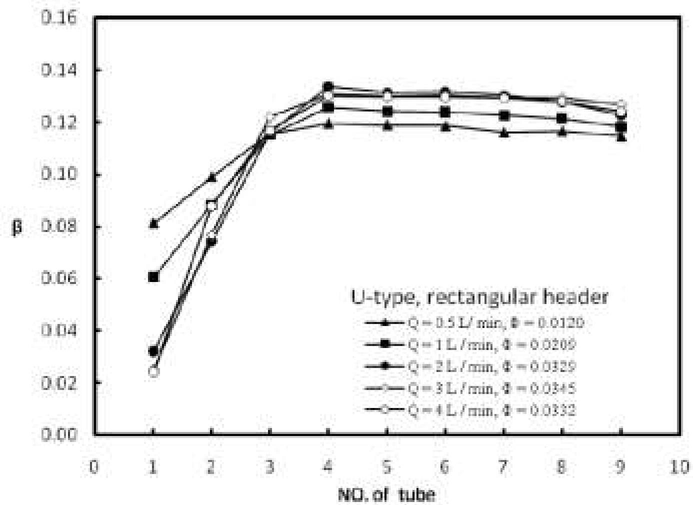

(a) Flow ratio of the typical rectangular header for U-type flow.

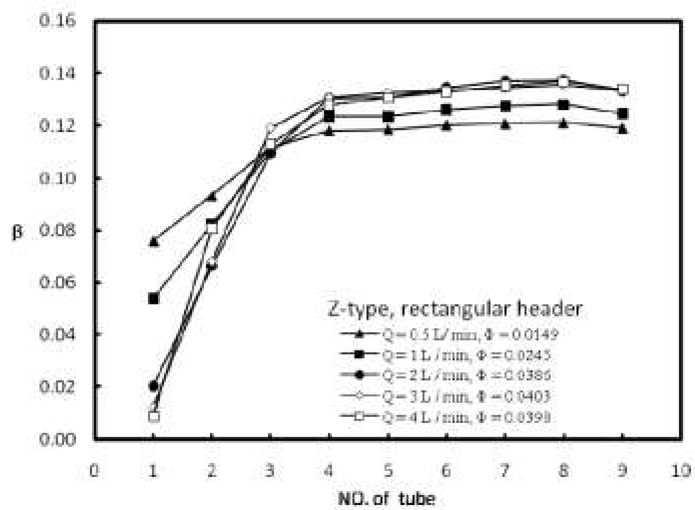

(b) Flow ratio of the typical rectangular header for Z-type flow.

Fig. 2. Flow ratio of the rectangular header for U-type and Z-type flows.

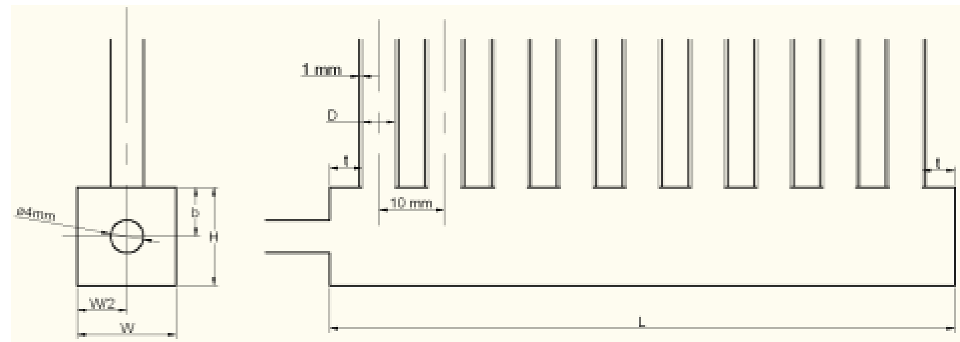

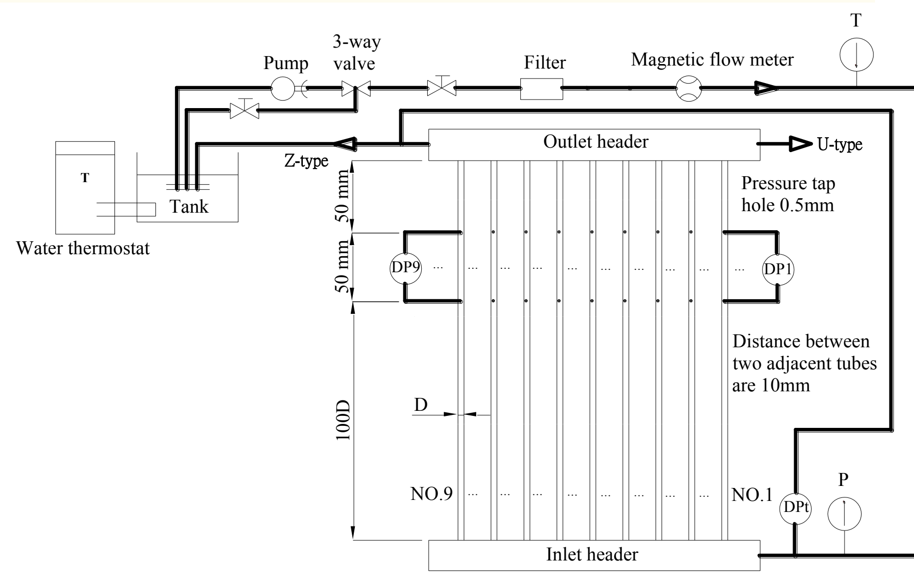

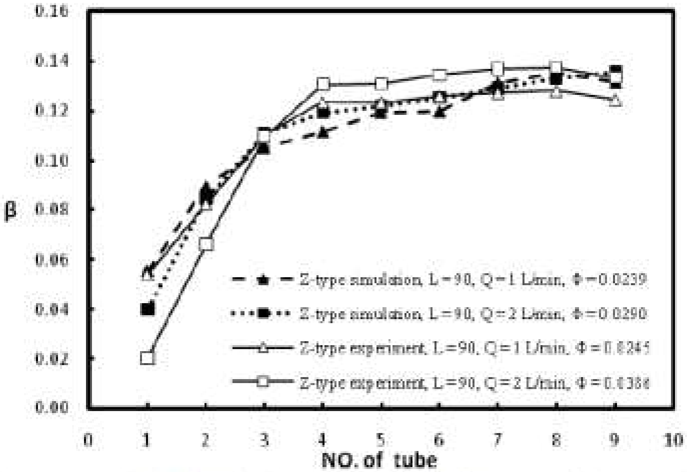

Fig. 7. Numerical simulation of flow ratio for Z-type flow vs. data.

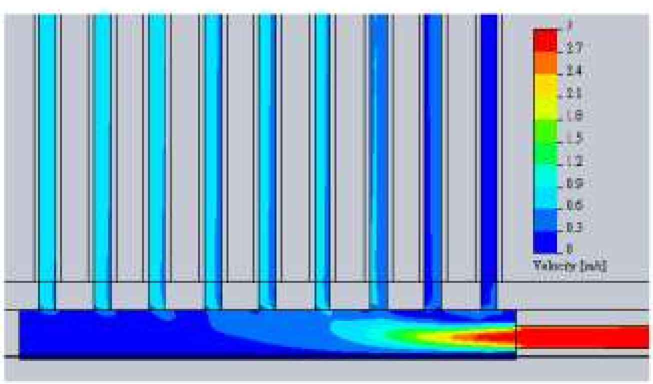

Fig. 8. Velocity simulation in 9×9 mm header with 90 mm length and D = 3 mm tubes.

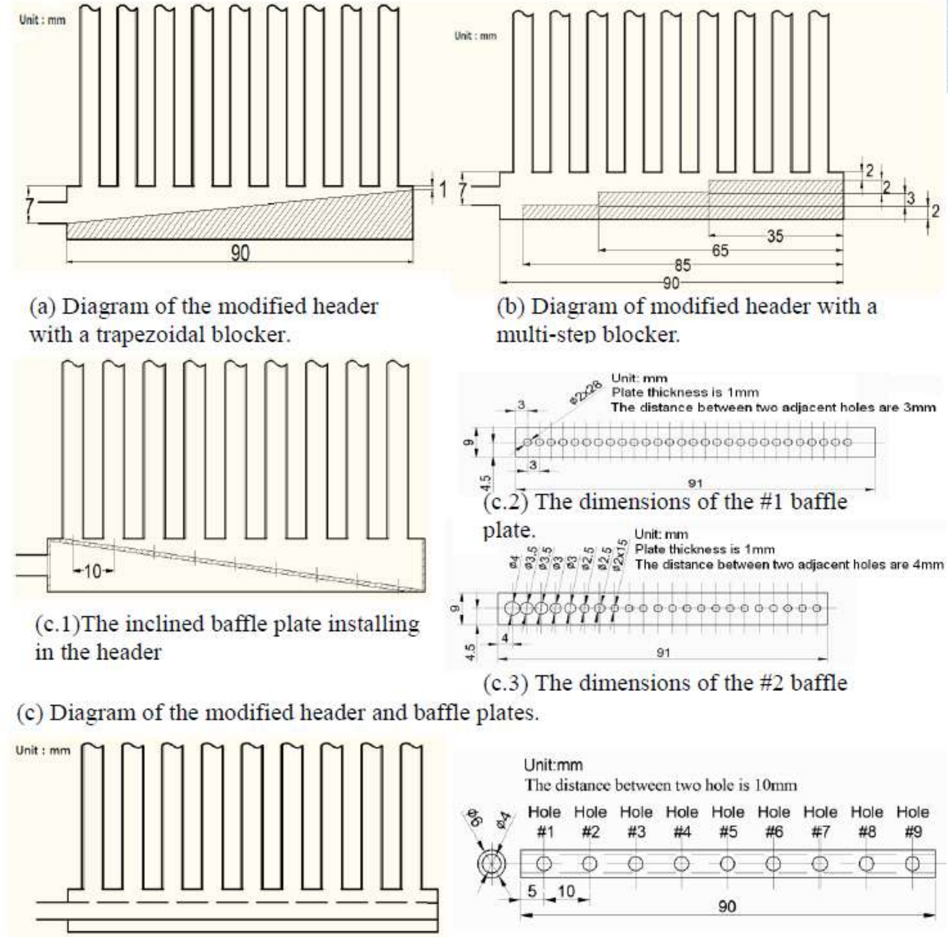

(a) Diagram of the modified header with a trapezoidal blocker.

(b) Diagram of modified header with a multi-step blocker.

(c.1)The inclined baffle plate installing in the header

(c.2) The dimensions of the #1 baffle plate.

(c.3) The dimensions of the #2 baffle

(c) Diagram of the modified header and baffle plates.

(d) Modified header with installed baffle

(d.1) Dimensions of baffle tubes.

Baffle tube	Hole #1 (mm)	Hole #2 (mm)	Hole #3 (mm)	Hole #4 (mm)	Hole #5 (mm)	Hole #6 (mm)	Hole #7 (mm)	Hole #8 (mm)	Hole #9 (mm)
#1	4	3.7	3.2	3.2	3.2	3.2	3.2	3.2	3.2
#2	4	3.5	3	3	3	3	3	3	3
#3	3.7	3.2	2.8	2.8	2.8	2.8	2.8	2.8	2.8
#4	3.5	3	2	2	2	2	2	2	2
#5	3	2.2	1.5	1.5	1.5	1.5	1.5	1.5	1.5
#6	2.8	2	1.2	1.2	1.2	1.2	1.2	1.2	1.2
#7	3.8	2.2	1.5	1.5	1.5	1.5	1.5	1.5	1.5

(e) The hole diameters for 7 baffle tubes

Fig. 1. Diagram of the modified headers.

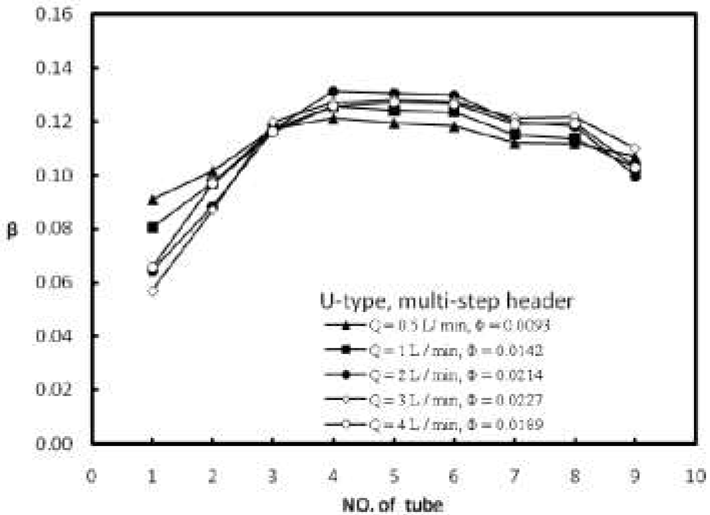

(a) Flow ratio of the multi-step header for U-type flow.

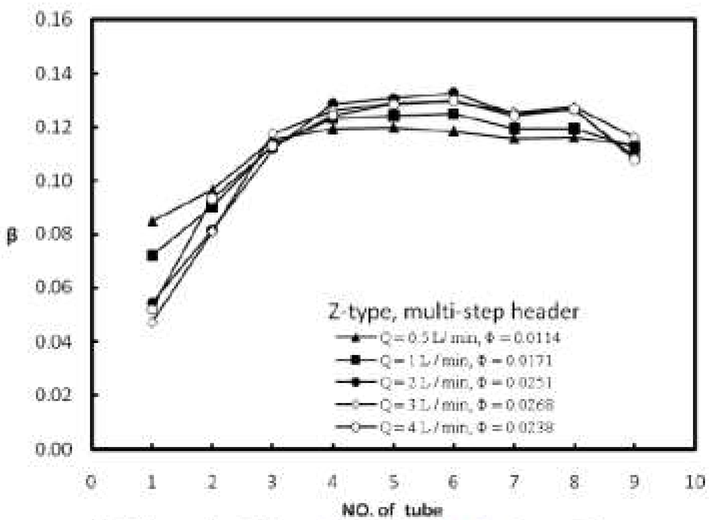

(b) Flow ratio of the multi-step header for Z-type flow.

Fig. 4. Flow ratio of the modified header with multi-step blocker for U-type and Z-type flows.

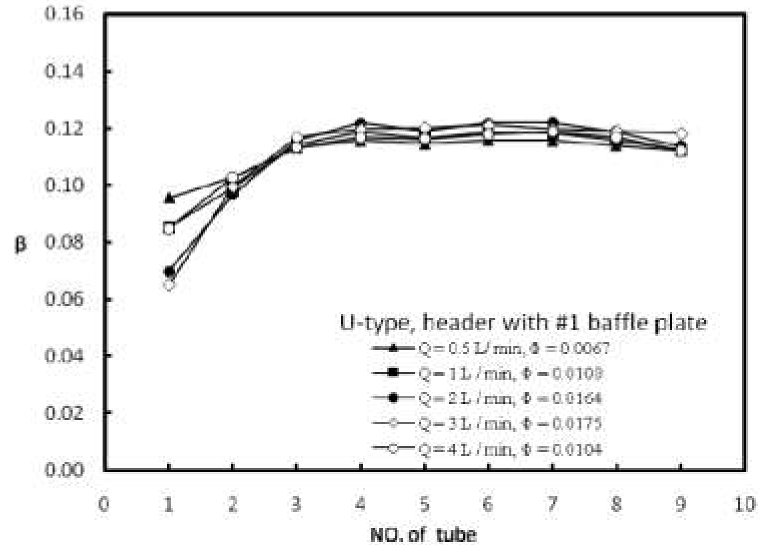

(a) Flow ratio of the modified header with #1 baffle plate for U-type flow.

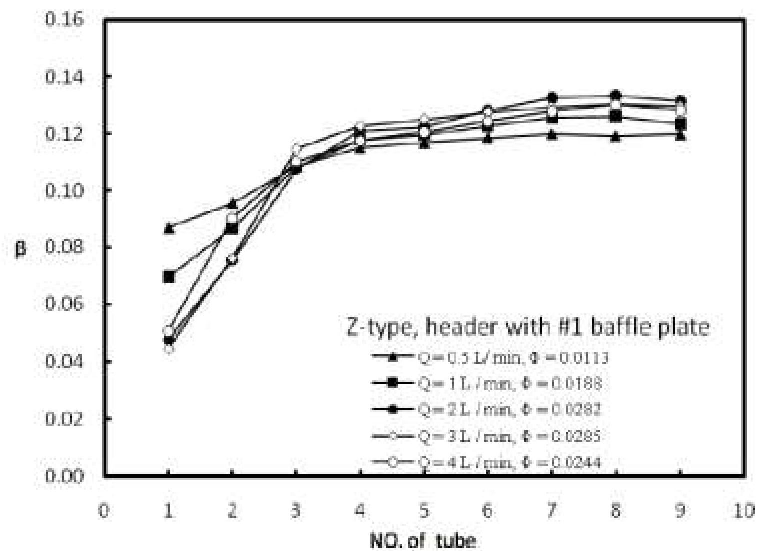

(b) Flow ratio of the modified header with #1 baffle plate for Z-type flow.

Fig. 5. Flow ratio of the modified header with #1 baffle plate for U-type and Z-type flows.

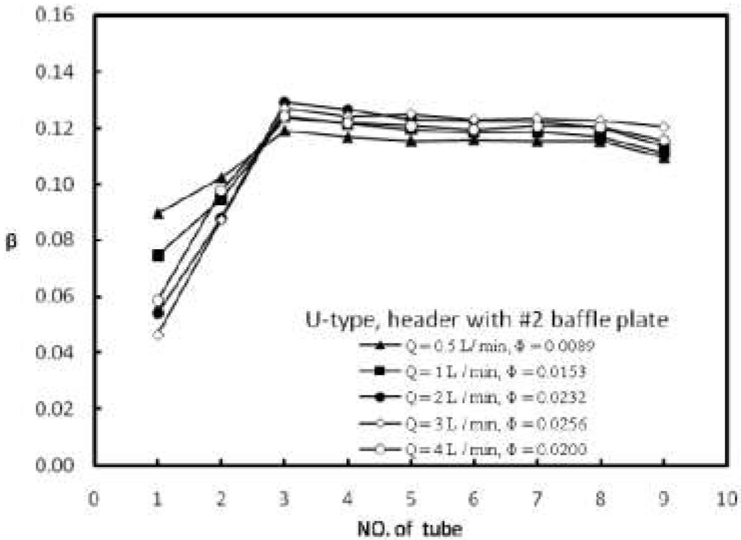

(a) Flow ratio of the modified header with #2 baffle plate for U-type flow.

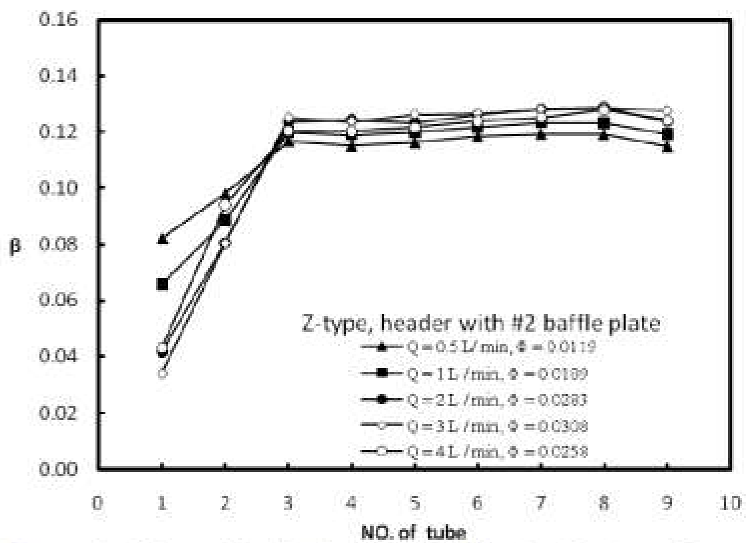

(b) Flow ratio of the modified header with #2 baffle plate for Z-type flow.

Fig. 6. Flow ratio of the modified header with #2 baffle plate for U-type and Z-type flows.

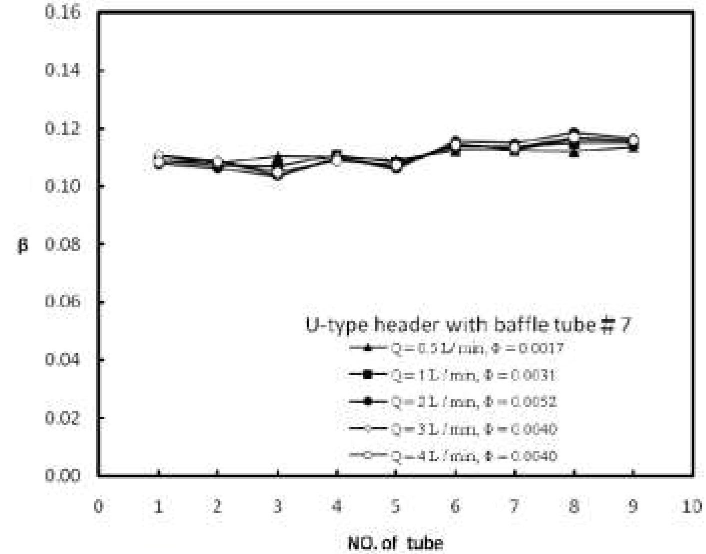

(a) Flow ratio of the modified header with #7 baffle tube for U-type flow.

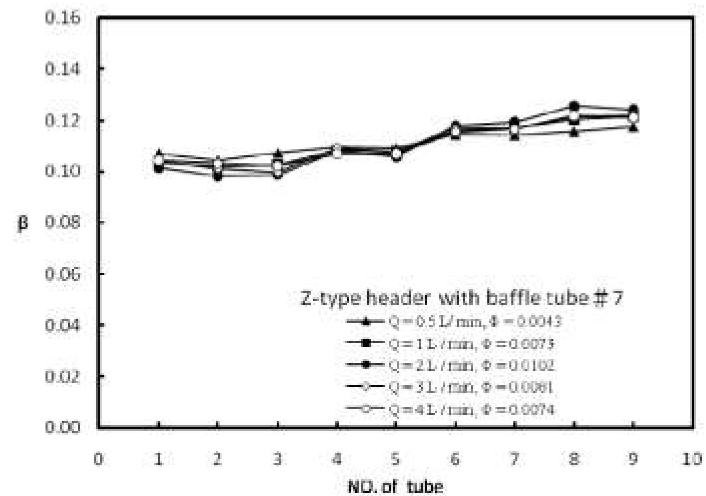

(b) Flow ratio of the modified header with #7 baffle tube for Z-type flow.

Fig. 9. Flow ratio of the modified header with #7 baffle tube for U-type and Z-type flows.

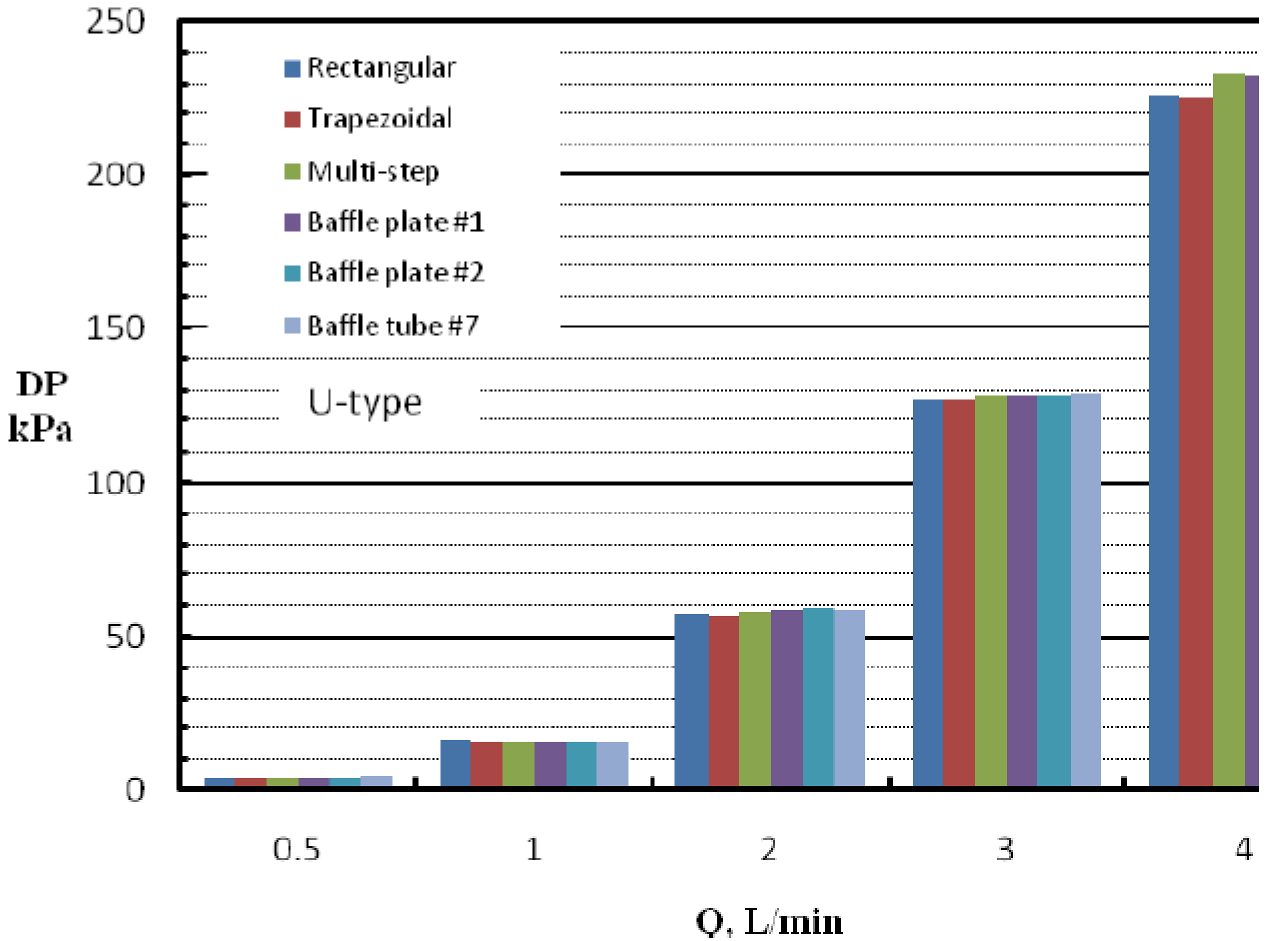

Fig.11. Comparison of the total pressure drop across inlet and outlet headers.

End Effects

- Insufficient cooling occurs near and far away from the CRAC.

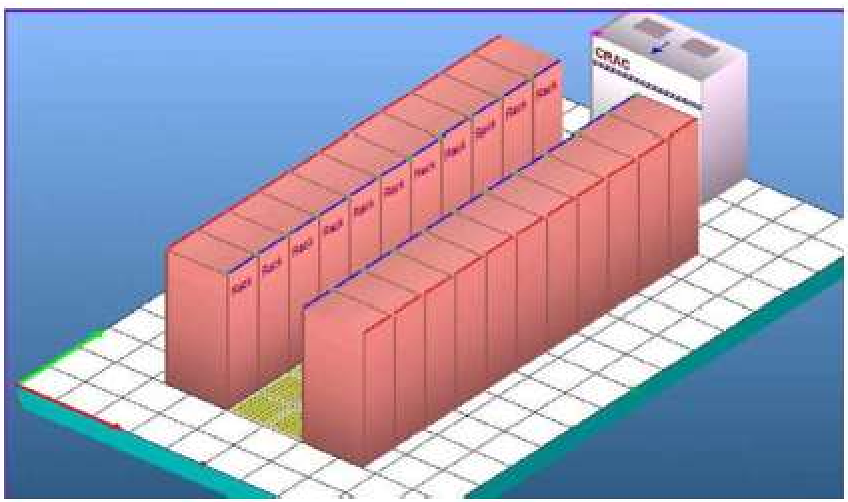

Fig. 28 A simple data center model with one CRAC, several racks, and perforated tiles (insufficient cooling airflow)

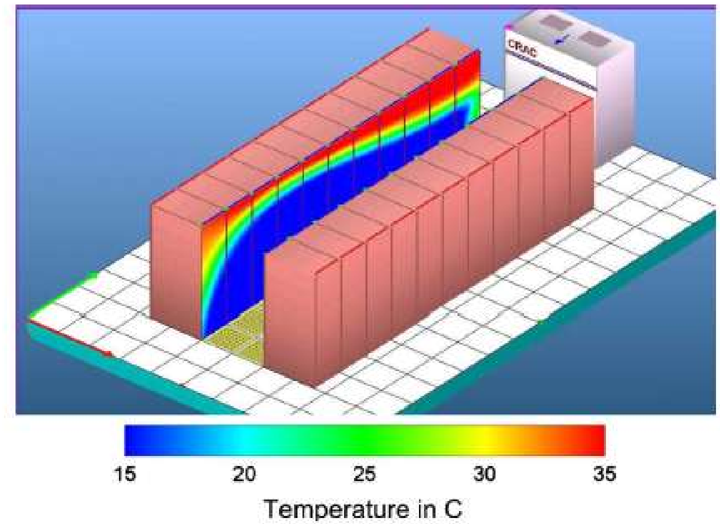

Fig. 29 Temperature distribution on the inlet faces of the racks (insufficient cooling airflow)

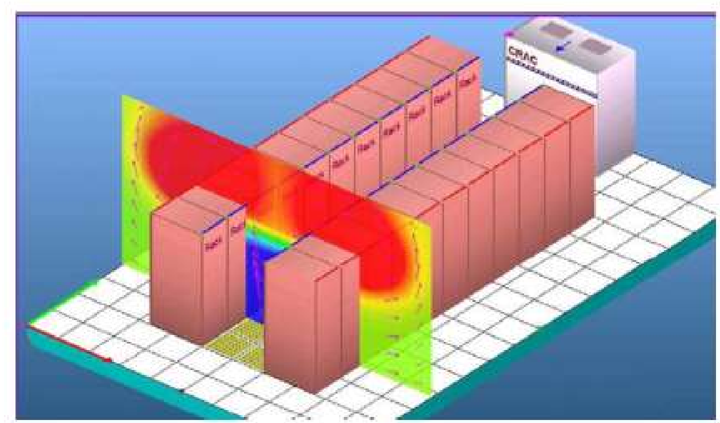

Fig. 30 Temperature distribution and velocity vectors on a plane (insufficient cooling airflow)

End effects (Conti..)

Fig. 1. Demonstration of heat recirculation: heated air in the hot aisle loops around the equipment to enter the air inlets.

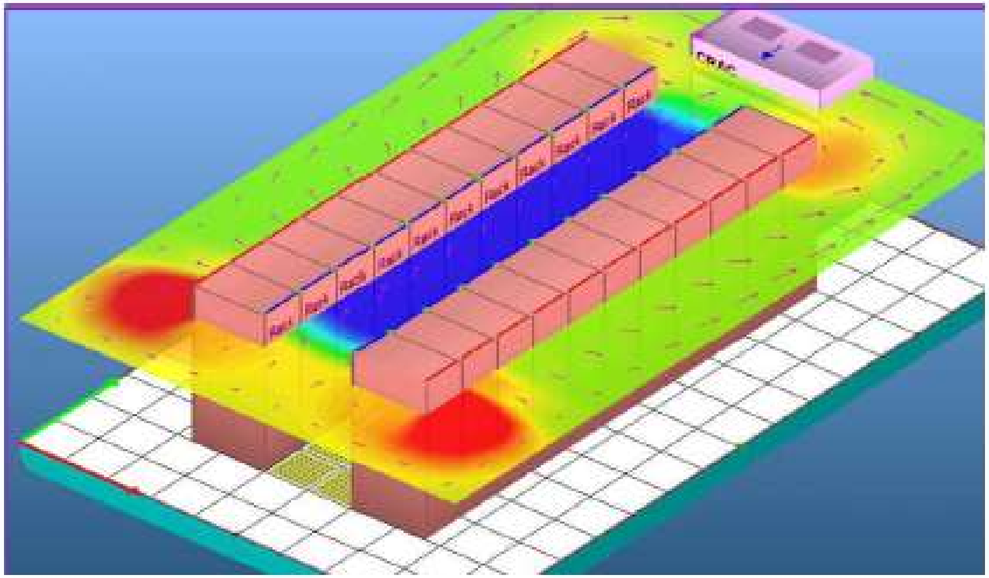

Fig. 32 Temperature distribution and velocity vectors on a horizontal plane

- The former is due to mal-distribution near the CRAC location.
- The latter is due to the hot exhaust circulation (heat recirculation).

IEEE TRANSACTIONS ON PARALLEL AND DISTRIBUTED SYSTEMS, VOL. 19, pp. 1458-1472, 2008

Gaps Between the Racks

- Normally, the racks are placed in a row in a contiguous manner. However, occasionally, there may be gaps between them. For example, in practice, gaps are created by removing a rack from a row. It is easy to see that the gaps provide additional places where the “end effects” can be observed. Hot air from the back of the racks can enter the cold aisle through the gaps and influence the inlet temperatures of the racks. An obvious remedy is to close the gaps by using impermeable plates or partitions.

High-Velocity Flow Through the Perforated Tiles

- The heat loads of modern server racks can be very high 10–20 kW
  - Corresponding airflow demand may be of the order of 1.0 m^3/s.
  - Air emerges from the perforated tile at a velocity of 3 m/s.
  - With this high-velocity stream flows over the inlet face of the rack, would the cooling air enter the rack or simply flow past it?
  - The high-velocity airflow does create a low-pressure region at the bottom of the rack. The server fans in the bottom region deliver a lower flow rate compared with the uniform-pressure environment.
  - Fortunately, this flow reduction is not large. For realistic values of the flow resistance inside a server rack and for common fan curves, the flow reduction at the bottom of the rack is less than 15%.

Use of Above-Floor Partitions

- The success of datacenter cooling
  - keeping the hot air away from the inlets of the server racks.

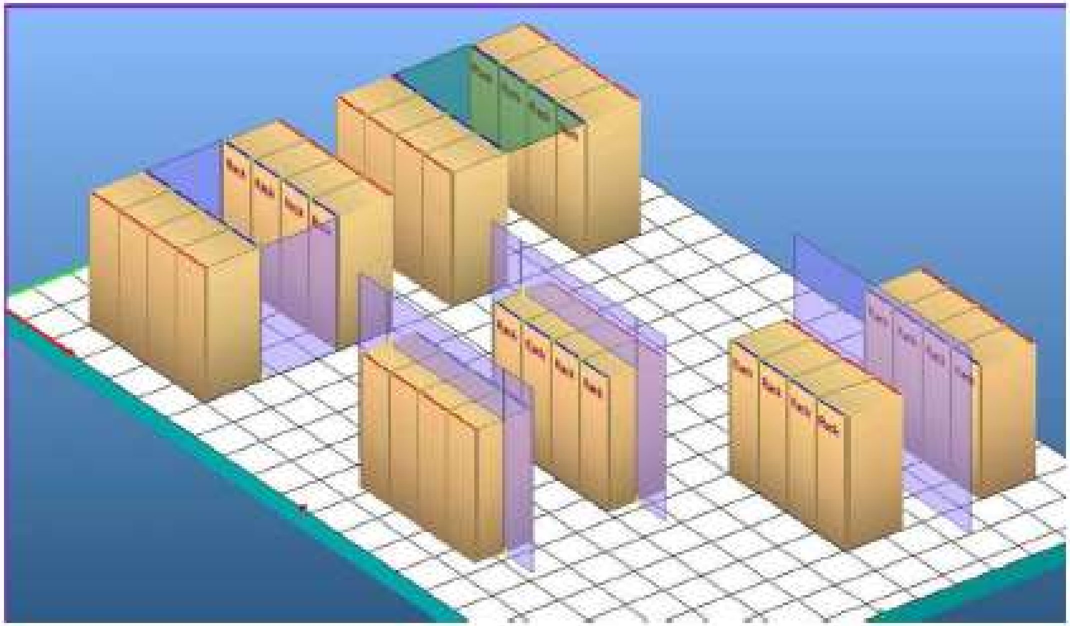

Fig. 35 Use of partitions to prevent hot air from entering the inlets of the racks

Dropped ceiling & ducted rack

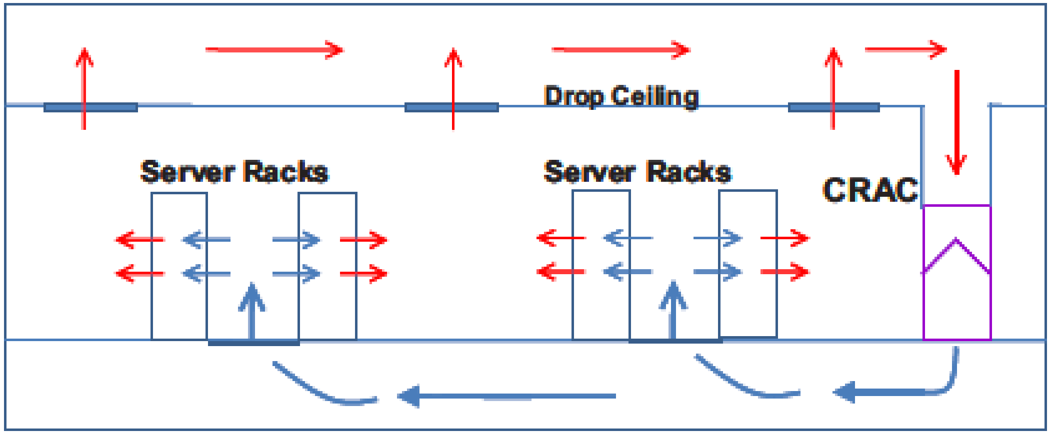

Fig. 36 A schematic of the drop-ceiling arrangement

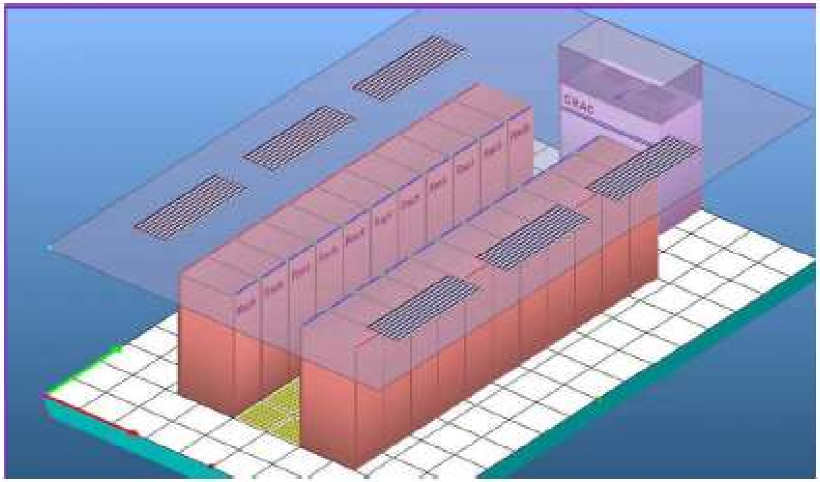

Fig. 37 Use of the drop ceiling

Concerns About the Ducted Solutions

- Originally, the attractions of the raised-floor design were its simplicity and flexibility. One could easily move server racks to new locations and simply place perforated tiles next to them to provide cooling air. No ducting was involved. The use of a drop ceiling requires ducting of the return airflow to the CRAC units. This makes it difficult to relocate the CRACs.
- Moving any racks to new locations requires the relocation of the vents on the drop ceiling in addition to the movement of perforated tiles.
- The use of ducted racks presents additional problems. Relocating the racks is now even more difficult. Further, the flow rate provided by the CRAC blowers should match the flow rate demanded by the internal fans in the server racks.
- The most serious concern is what happens when one or more CRAC units fail. Then the above-mentioned matching of the flow rates does not hold any more. Since the airflow exhausted by the server fans cannot be handled by the remaining CRAC units, hot air will flow from the space above the drop ceiling into the room below through the extra vents provided. This hot air can directly enter the server racks causing a catastrophic failure.

Short Summary

- Significant progress had been made for datacenter thermal managements during the past decades.
- Some guidelines are available but normal designs are still in a "case by case" basis for complex interactions amid building configuration, heat transfer, air flow, IT equipment and zoom layout.
- Recirculation/bypass/negative pressure elimination is regarded as the most important issue for air flow management.
- Flow distribution is the most crucial problems in thermal managements, it depends on various parameters. This is especially critical for a air cooling system.
- CFD is a viable tool in thermal management of datacenter.

Thank You