DATA CENTRE COOLING SYSTEM

A) Differences between Precision and Comfort Cooling

Here are the major differences exist between precision air conditioning and comfort systems.

1. Cooling Optimized to Electronic System Requirements

Dense loads of electronics generate a dryer heat than typical comfort-cooling environments and this significantly changes the demands on the cooling system.

There are two types of cooling: latent and sensible. Latent cooling is the ability of the air conditioning system to remove moisture. This is important in typical comfort-cooling applications, such as office buildings, retail stores and other facilities with high human occupancy and use. The focus of latent cooling is to maintain a comfortable balance of temperature and humidity for people working in and visiting such a facility. These facilities often have doors leading directly to the outside and a considerable amount of entrance and egress by occupants. Sensible cooling is the ability of the air conditioning system to remove heat that can be measured by a thermometer. Data centers generate much higher heat per square foot than typical comfort-cooling building environments, and are typically not occupied by large numbers of people. In most cases, they have limited access and no direct means of egress to the outside of the building except for seldom-used emergency exits.

Comfort air conditioning systems have a sensible heat ratio of 0.60 to 0.70. This means that they are 60 to 70 percent dedicated to lowering temperature, and 30 to 40 percent dedicated to lowering humidity. Data center environments require a 0.80 to 0.90 sensible heat ratio for effective and efficient data center cooling. Precision air conditioning systems have been designed with a sensible heat ratio of 0.85 to 1.0. This means 80 to 100 percent of their effort is devoted to cooling and only 0 to 20 percent to removing humidity. So more “nominal” 20-ton comfort units will be required to handle the same sensible load as “nominal” 20-ton precision units

As the need for latent heat removal lessens, so too does the need for dehumidification.

Data centers have a minimal need for latent cooling and require minimal moisture removal. Because precision cooling systems are engineered with a focus on heat removal rather than moisture removal and have a higher sensible heat ratio, they are the most useful and appropriate choice for the data center.

2. Systems Sized to the Higher Densities of Data Centre Environments

Heat densities in electronics environments are three to five times higher than in a typical office setting, and increasing at a faster rate than ever before. To illustrate, one ton of comfort air conditioning capacity (12,000 BTU/hour or 3413 watts) is required per 250-300 square feet of office space. This translates into 15 watts per square foot. In contrast, one ton of precision air conditioning capacity is required per 50-100 square feet of data center space. That translates into a much larger 75 watts per square foot. And this is an average number that is increasing yearly. Some sites can have load densities as high as 200 – 300 watts per square foot.

From an airflow standpoint, precision air systems are designed differently than comfort air systems to manage the larger load densities in data centres. A precision air system achieves a higher sensible heat ratio, this helps maintain target temperature and humidity levels and contributes to better air filtration via the movement of significantly larger volumes of air. Precision air equipment typically supply 500 to 900 cubic feet per minute (cfm) per cooling ton. This contrasts with the much smaller range of 350 to 400 cfm typically delivered by comfort air equipment.

3. Precise Humidity Control

Ignoring the impact of humidity can result in serious long-term problems, including damage to equipment and other resources, and to the facility’s infrastructure. The optimal relative humidity range for a data center environment is 45-50 percent. An above-normal level of moisture can corrode switching circuitry, which can cause malfunctions and equipment failures. At the other end of the spectrum, low humidity can cause static discharges that interfere with normal equipment operation. This is a more likely scenario in a data center since it is typically cooled 24x7, creating lower levels of relative humidity.

Comfort air systems typically have no humidity control, which makes it difficult to maintain stable relative humidity levels. If the necessary controls and components are added, they have to be set-up to operate as a complete system. Precision air systems have multi-mode operation to provide a proper ratio of cooling, humidification and dehumidification. This makes them much more suitable for the low tolerance range of humidity levels in a data center.

4. **Protection against Airborne Contaminants**
   Even small amounts of dust or other particles can damage storage media and charged electronic components. Most comfort air systems use residential-type air filters that are 10 percent efficient, making them inadequate for a data center environment. Precision air system filters have higher quality internal filter chambers that are 20-30 percent efficient and ASHRAE compliant.

5. **Efficient Continuous Year-Round Cooling**
   Comfort air conditioning systems for most buildings are designed to operate an average of 8 hours per day, 5 days per week. This translates into about 1,200 hours per year, assuming cooling is required only during the summer months.

   Most data centers require heat rejection 24 hours per day, 365 days per year regardless of outside weather conditions. Precision air conditioning systems and their components are engineered to meet this high cooling demand. A precision air unit's circulating fan operates continually, 8,760 hours per year, with other components turning on and off as needed.

6. **Lower Operating Costs**
   Because of basic engineering, design and equipment differences, a purchase price comparison of comfort versus precision air conditioning systems does not tell the complete story. A more accurate comparison will consider the difference in operating costs between the two systems.

7. **Enhanced Service and Support**
   Critical environments require high availability of critical support systems. Therefore, it is important that these systems operate reliably, and that their performance is tuned specifically to control the environment for the computer systems with which they are deployed. Precision air systems often feature greater internal redundancy of components than comfort cooling systems, allowing them to continue operating in the event of some failures. In addition, they are supported by professionally-trained, locally-based installation, service and support partners that are accustomed to the needs and sensitivities of working in the data center environment.

B) Traditional Cooling Systems (Room Oriented) For Data Centre

Traditional precision air conditioning units have provided effective cooling in thousands of data centers around the world; however, as system densities increase, they are being stretched to the limits of their practical capacity.

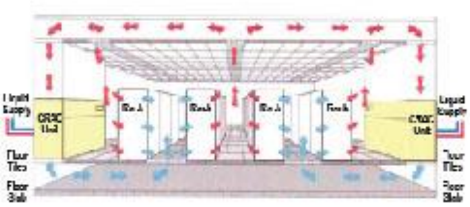

The key limitations involve the number of precision air conditioners that can be installed in a given room and the amount of air that can be pushed through perforated floor tiles in a raised floor. Floor-mounted precision air systems take up data center floor space, limiting how many systems can be installed in a data center facility. In addition, there is a physical limitation to how much air can be efficiently distributed through the raised floor. Trying to push too much air under the floor can create negative pressures that can actually draw cool air back down through the perforated tiles, rather than forcing it up into the room. In addition, the floor tiles themselves have physical limits as to how much air can actually pass through the perforations. Consequently, increasing cooling capacity will not necessarily result in a corresponding increase in cooling across the room. Computational fluid dynamics (CFD) Test showed wide variations in temperature across the room when raised floor cooling alone is used to cool hig h density racks. There are several steps that can be taken to optimize the efficiency of the raised floor system. The first is an examination of the cabling running under the floor to ensure it is not obstructing air flow. Floor height also plays a role. Doubling the floor height has been shown to increase capacity by as much as 50 percent. Data center managers planning new facilities should consider floors higher than the traditional 450mm. However, replacing the floor is usually not an option for existing data centers because of the disruption in operations it requires. First though, traditional cooling should be maximized to ensure it provides an efficient, flexible and reliable foundation for adaptive cooling.

Traditional floor-mounted cooling systems with under-floor air delivery will continue to play an essential role in data center environmental management. It is recommended that traditional systems be configured to deliver the required cooling for the first few kW/rack (about 150 W/sqft) of heat load as well as solve the room's full humidification and filtration requirements. Supplementary cooling is typically deployed for densities beyond 5 kW per rack.

C) Hot-Aisle/ Cold-Aisle Configuration

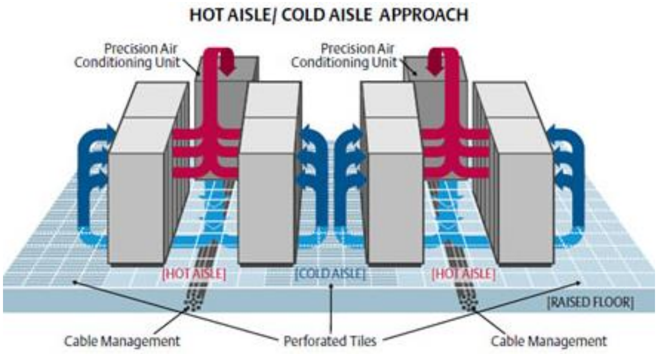

Most equipment manufactured today is designed to draw in air through the front and exhaust it from the rear. This allows equipment racks to be arranged to create hot aisles and cold aisles. This approach arranges racks front-to-front so the cooling air rising into the cold aisle is pulled through the front of the racks on both sides of the aisle and exhausted at the back of the racks into the hot aisle (see above Figure). Only cold aisles have perforated floor tiles, and floor mounted cooling is placed at the end of the hot aisles — not parallel to the row of racks. Parallel placement can cause air from the hot aisle to be drawn across the top of the racks and to mix with the cold air, causing insufficient cooling to equipment at the top of racks and reducing overall energy efficiency.

D) Equipment Trends and News Challenges for DC Cooling

New high-performance equipment, such as dual-processor servers and high-speed communications switches, are raising rack densities well above 30 kW. Adjacent Figure illustrates the rising heat load of electronics equipment and predicts future heat loads. With server power requirements exceeding projections, cooling strategies must adapt at a faster pace than anticipated to avoid downtime, equipment failure and reduced lifespan of electronics.

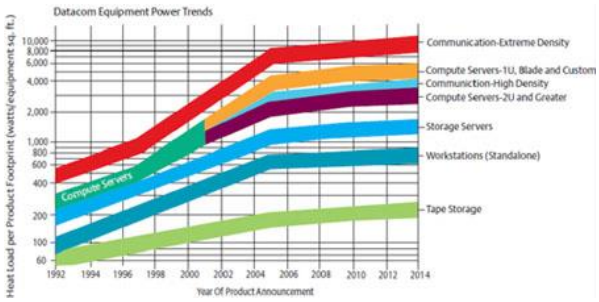

And with continued pressure to drive down data center operating costs, many organizations are attempting to pack as much equipment into as small a space as possible. As a result, rooms are heating up and organizations are feeling the effects. Already, many of today's data centers require more than 100 Watts of power per square foot. The latest generation of blade servers pushes power and heat levels even higher. A single rack loaded with four fully configured IBM BladeCentre™ H Chassis, each drawing 5.8kW, creates a load of almost 24 kW in an enclosure that takes just seven square feet of data center floor space. This shows a sharp contrast with the state of the industry in 2000 when the average rack consumed just 1 kW of power. Communications equipment is progressing in the same direction. Depending on its power supply configuration, the Cisco CRS-1 router creates a heat load of 15 to 16.6 kW per rack. To further complicate the challenge, the average server replacement cycle is three to four years, producing a diverse per-rack heat emission throughout the data center. While a rack of servers bought four years ago may use 2 kW of power per rack, the rack of IBM BladeCentres referenced previously has a power draw — and corresponding heat load — of almost 24 kW.

To approach cooling such a room by sizing a traditional precision air conditioning system to address the hot spots would mean vastly over-sizing your system and wasting a significant amount of energy, not to mention threats to availability posed by hot zones. If the heat from a rack is not effectively removed, the performance, availability and lifespan of the equipment in the rack will be reduced significantly. Increasingly, as organizations adopt the latest server technologies into their existing data centers, they are exposed to higher failure rates, especially in the top third of the rack. As cooling air is supplied from the raised floor, it is fully consumed by high-density equipment at the bottom of the rack, while the top of the rack is deprived of the cooling air it requires. This is compounded by the fact that high densities also cause hot air to be recirculated back through the top third of the rack. This is why failure rates are higher for equipment at the top of the rack. The increasing failure rate at the top of racks is occurring because current air delivery through a raised floor is generally limited to cooling an average room load of about 150 Watts per square foot, or racks with 2 to 3 kW load. Beyond that point, the volume of air that can be effectively delivered to the equipment in the upper part of the rack is insufficient.

E) Supplementary Cooling For High Density Loads

To effectively supplement traditional cooling and address high density areas, cooling must move closer to the source of heat. There are three major choices to be made when deciding what technology to employ in addressing high density cooling needs: cooling fluid, system architecture and future capabilities.

Cooling Fluid: Water vs. Refrigerant

While water is regularly used in floor-mounted cooling, when the source of cooling moves close to sensitive electrical equipment, safety becomes a key concern. That's why pumped R134a refrigerant is an ideal choice for high-density applications. Because refrigerant turns into a gas when it reaches the air, a leak would not damage IT equipment or pose a safety hazard. Pumped refrigerant solutions also provide an incremental energy efficiency savings of between 25 percent and 35 percent, based on kilowatts of cooling capacity per kW of heat load.

System Architecture: Open (Row oriented) vs Closed (Rack oriented)

Cooling can be brought close to the load through either an open or closed cooling system architecture. In a closed architecture, the electronics and cooling equipment are located together in a sealed environment.

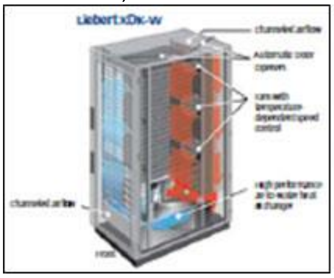

Example of high heat density cooling system with open architecture

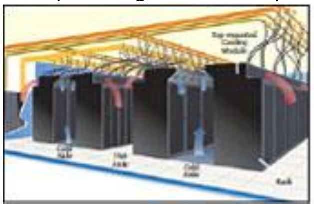

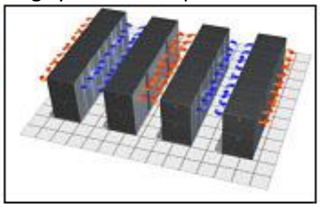

This approach provides high-capacity cooling at the expense of flexibility and fault tolerance if failure-mode precautions are not built in. While this approach is appropriate for small implementations, assuming failure mode ventilation, an open architecture approach is preferred for cooling in a data center environment. In a data center, closed cooling offers limited flexibility of rack combinations and often no backup emergency cooling. If the cooling fails, racks are isolated from room cooling and the temperature in the enclosure can reach the server over-temperature limit condition in less than 15 seconds. This is a risk that's unnecessary in a room environment. In an open architecture, where modules are on or near racks, but not part of an enclosure, room air is used as a buffer in the event of a failure, making it a safer alternative in many cases. Additionally, an open architecture allows greater flexibility to reconfigure as needs change and additional cooling capacity is needed.

Mixed Architecture

Nothing prevents room, rack and row architectures from being used together in the same installation. In fact, there are many cases where mixed use is beneficial. Specifically, a data center operating with a broad spectrum of power densities could benefit from a mix of all three types as shown in Figure here:

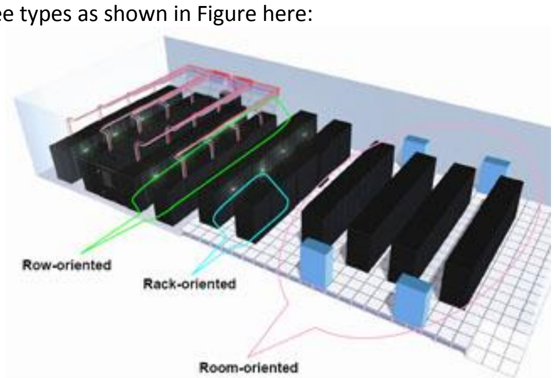

Supplemental cooling has a number of additional benefits, including:

- Ability to target zones and hot spots with overhead and next-to-the-rack cooling modules
- Increased cooling system scalability with plug-and-play cooling modules
- Greater floor space savings due to available cooling modules positioned above the racks
- Applicability in raised floor and non-raised floor environments
- Improved energy efficiency that lowers operating costs

  - **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/m2 (30- 80 W/ft2)
  - **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.

Another effective use of row and rack-oriented architecture is for density upgrades within an existing low density room-oriented design. In this case, small groups of racks within an existing data center are outfitted with row or rack-oriented cooling systems. The row or rack cooling equipment effectively isolates the new high density racks, making them essentially “thermally neutral” to the existing room-oriented cooling system.

In this way, high density loads can be added to an existing low density data center without modifying the existing room-oriented cooling system. When deployed, this approach results in the same mixed architecture depicted by Figure above.