

[Related Commercial Resources](#)

CHAPTER 20. DATA CENTERS AND TELECOMMUNICATION FACILITIES

1. OVERVIEW AND DEFINITIONS

DATA centers and telecommunication facilities are significantly different than most other facilities:

- Load density (10 to 100 times typical office building)
- 24 h load profile (relatively constant)
- Datacom equipment life cycle (3 to 5 years)
- Primary occupant (datacom equipment)
- Operation (24 × 7)
- Sensible load (nearly 100%)
- Commissioning level
- A1 Allowable Conditions: 59 to 89.6°F DB, 8 to 80% rh

The telecommunication industry, once predominantly regulated land lines, now uses wireless technology running the same communications protocol (Internet Protocol or IP) as the data center industry. As a result, data centers and telecommunications facilities continue to converge. TC 9.9 uses the term *datacom* to indicate both data centers and telecommunication facilities.

This chapter provides some basic information about datacom facilities, the fundamentals of their designs, and where to find additional information. Datacom facilities treat space, power, cooling, and networking (formally or informally). Each service can have a service-level agreement (SLA), but the services are highly interdependent. Therefore, overall reliability/availability is best achieved when all aspects of these services are designed together, with the same performance goals. In short, a datacom facility is a system, in every sense of the word, and all facets interrelate and affect each other in one way or another. Therefore, this chapter also includes information on facility elements such as power, lighting, and fire protection that are not strictly HVAC, but which can affect the cooling load, system performance, or both if not recognized as critical elements in the design phase and fully coordinated with each other.

Because of the high power and heat densities, it is becoming increasingly necessary to cool computing equipment with liquid approaches, either instead of or in combination with air. Further, it is becoming increasingly popular to provide metering for services at each service interface point, with centralized monitoring of the infrastructure. This **data center infrastructure monitoring (DCIM)** approach has become an almost mandatory additional requirement for modern datacom facilities.

DEFINITIONS

An **enterprise datacom facility** is owned and operated by an organization to support its data processing and data storage requirements. It includes the IT equipment like servers, networking systems, and data storage devices as well as the required power, cooling, and environmental monitoring and controls. Because of the high capital cost and short life cycles of many datacom facilities, and the ready availability of colocation centers, the trend is towards companies owning fewer of their own **enterprise** datacom facilities, and renting more resources from third-party facility owner/providers. Rented or leased services and facilities come in many varieties; a common general format is retail or wholesale colocation facilities.

A **colocation center** (also co-location, collocation, colo, or coloc) is a type of datacom facility where equipment, space, and bandwidth are available for rent. Colocation facilities provide space, power, cooling, and physical security services for server, storage, and networking equipment. Their fiber services are typically redundant and diverse, and connect the facilities to various telecommunications and network service providers, often referred to as a **point of presence (PoP)**. However, the power and cooling redundancies can be significantly different from one colocation center to another, and should be evaluated before signing a contract, which should include a carefully worded SLA. Failures in these facilities can have widespread effects.

A **cloud computing center** is fundamentally like any other datacom facility, except that it hosts users virtually rather than physically. Cloud centers tend to operate at a higher performance level than colocation or enterprise facilities because they can standardize their hardware and optimize both IT equipment (ITE) use and energy efficiency.

Supercomputing centers run highly specialized, extremely high-speed computers whose compute power is generally shared by multiple users. While there is some similarity in concept to cloud computing centers, supercomputing centers are often non-commercial, located at major universities, and interconnected via a dedicated Internet. Supercomputing centers tend to process scientific research that requires the processing of mammoth amounts of data. The rooms in which they are located are just as specialized, with liquid cooling being the predominant method of heat removal.

Edge computing facilities are usually much smaller datacom rooms, dedicated to specific computing functions where latency issues predominate. As such, they are located close to their users, with the results of their processing often linked back to larger hosting sites. Their name comes from the fact that they are at the edge of the computing domain when illustrated on a network diagram. Their requirements are just as rigorous as any other high-availability datacom facility, except that they must usually run unattended and still maintain their high level of reliability.

Datacom facilities provide space, power, cooling, and networking to datacom equipment (hardware), also known as information technology equipment (ITE) in the U.S. *National Electrical Code*[®] (NFPA *Standard 70*). The space within the datacom facility that actually houses the datacom hardware may be called the data hall, the ITE equipment room, the computer room, the white space, or the raised-floor area (regardless of whether a raised access floor is actually used). Although [Figure 1](#) shows the various elements that may make up a complete facility, the actual elements (and their arrangements) will vary considerably in each project. All these space components should at least be considered, however, since some are often overlooked, at the long-term jeopardy of the facility.

This chapter focuses on the most important facility requirements for the support of the datacom equipment, which include thermal, air quality, and power, as well as fire protection.

[Figure 1](#) shows an overview of the major spaces in a typical datacom facility, whether enterprise, colocation, or cloud.

2. DATACOM EQUIPMENT, POWER TRENDS, AND ENVIRONMENTAL GUIDELINES

2.1 DATACOM EQUIPMENT WORKLOAD

Datacom equipment (hardware), often called **information technology equipment (ITE)**, has various workload states ranging from essentially idle/static (not performing any actual useful work) to running at its maximum performance and central processing unit (CPU) utilization. The hardware workload is driven by software, which includes system or operating system (OS) software (including networking), and application software that yields calculation or data manipulation results (i.e., actual useful work).




Figure 1. Typical Datacom Facility Space Plan

The number of applications available across all hardware types is vast (in the multimillions at least). Software can often be added or upgraded in various ways, including remotely. On average, over a 24 h period and during the span of a year, power and heat loads are generally stable unless a significant amount of ITE is added; from moment to moment, however, workloads (and therefore power and cooling loads) can be very dynamic.

Datacom equipment life cycles are much shorter than power and cooling infrastructure life cycles. Application software life cycles are even shorter. It is critical that planning of power and cooling infrastructure considers the life cycles and refresh (churn) rates of hardware and software, which can be very short in industries such as financial, and quite long in other operations like education and government where funding is often limited and sporadic. Operations with long life cycles may also be using legacy hardware that takes more space and consumes more power than more modern equipment, often lowering the power and heat densities as well. But when this equipment is eventually replaced, newer hardware could require more power per rack and cabinet, and also have higher heat densities in less space. Planning must provide for these changes based on industry norms such as those published in the ASHRAE Datacom series of books.

Load Characterization

When planning the datacom power and cooling infrastructure, two common means of characterizing maximum load have been used: watts per square foot and kilowatts per datacom equipment rack or cabinet. The datacom industry sometimes uses **granularity** as a means of describing the unit size.

The datacom industry generally recognizes that kilowatts per rack or cabinet is superior to average watts per square foot in estimating loads. However, at the start of a project, there may be insufficient information about the quantity of racks or their expected contents to do this, making that metric too granular. Professional judgment and experience are critical to deciding which maximum load characterization to use, and when to refine the design based on a more accurate metric.

It is equally important to characterize the minimum load and load variation. If cooling equipment is sized for the ultimate potential load, it may not be able to function effectively or efficiently at the low day-1 load. This can particularly affect redundancy. If the initial load is too low, cooling units may need to be shut down rather than running simultaneously in a redundant load-sharing mode, which is generally more energy efficient as well as more operationally

secure. Further, since the time increment for load variation can be very short (e.g., seconds, minutes) or very long. It is important to obtain or develop a detailed load profile including future possibilities.

2.2 DATACOM EQUIPMENT RACKS

Most datacom equipment is rack or cabinet mounted (the terms are often used interchangeably), but others come in prepackaged configurations, including large, stand-alone cabinets that do not match the rest of the datacom facility standard. Regardless of size, rack and cabinet sizes and equipment mounting standards are defined by the Electronic Industries Alliance (EIA 2005).

The vertical dimension is expressed in terms of **units (U)** (sometimes called **rack units [RU]**). One U or RU represents 1.75 in. of vertical height within a rack. A common height for a rack is 42 U of usable space, although some are taller and some are shorter.

Regardless of terminology, *racks* and *cabinets* are technically different. A **rack** is an open-frame two- or four-post mounting used more for telecom and patch panels than for servers. A **cabinet** is a similar four-post framework, but is equipped with sides, top, often front and rear doors, and sometimes a bottom plate (Figure 2).

Typical rack mounting rail widths are nominally 19 or 23 in., depending on their construction. The actual space between the mounting rails is approximately 2 in. less than the nominal panel widths, to allow room for screwing equipment flanges to the rails. Cabinets tend to be a nominal 24 to 32 in. wide and 24 to 48 in. deep. The 24 in. wide cabinets are not used much in datacom rooms any more. Wider cabinets provide space for the massive amounts of power and data cabling associated with full configurations of high-density hardware, and the often large power strips to which the ITE connects. Crowded cable can drape behind the ITE, blocking air discharge from the backs of computing hardware. The widest cabinets are also often used for large network switches that use non-conventional airflow patterns. The extra width provides space for the air diverters necessary to maintain uniform airflow within an equipment row (see the section on Datacom Equipment Airflow). Deeper cabinets have become necessary to accommodate the form factors of newer datacom hardware, which are often compressed to only 1 or 2 U high, but can be quite deep as a result.

Servers used for computing are available in rack mount and custom configurations. Most servers are full-rack width and are often identified as having 1U, 2U, 4U, etc., form factors (Figure 3). A half-width server mounts two separate boards side-by-side in a single-width 1U high chassis, or four separate boards in a single-width 2U chassis. Larger form factors may house multiple modular servers (**blade servers**) in its overall chassis.




Figure 2. Typical Rack and Cabinet Examples

2.3 DATACOM EQUIPMENT (HARDWARE)

Datacom components (e.g., processors, memory, storage, input/output [I/O], power supplies) may be packaged into single-chassis datacom equipment, or some components may be separated. This section is limited to total datacom equipment requirements and interfaces; components are addressed only to the level necessary to describe the requirements and interfaces, but are covered in more depth in the section on Datacom Equipment Components.




Figure 3. Typical Computer Server Packaging Form Factors

Datacom equipment predominantly consists of servers (volume, blade, etc.), communication equipment (switches, routers, etc.), and data storage devices (storage area network [SAN], network attached storage [NAS], and other formats that are beyond the scope of this chapter).

For air-cooled datacom equipment, the primary interfaces to the facility are the cool-air inlet to the datacom equipment and the hotter exhaust air from it. For liquid-cooled equipment, the interface to the facility is the liquid connection to the equipment or the rack. The datacom equipment interface focuses on

- Temperature
- Humidity
- Coolant quality (air or liquid)
- Coolant flow (air or liquid)

Server Classifications

Servers tend to be the most common equipment within a datacom space; there are many different types of servers and any type of server can go into any type of datacom facility. Although there are no set rules regarding what constitutes any specific server type, the following classifications help provide some guidance:

- **General purpose, SFF.** These servers are typically single or dual socket servers packaged in 1U, 2U, or half-width form factors. They generally have many features, enabling them to cover a wide variety of customer needs. These are often called **small-form-factor (SFF) servers** because of their minimal usage of rack height, but their depths can be significantly greater than will fit into legacy 24 in. deep racks, especially when their connectors and cabling are considered.
- **General purpose, LFF.** These servers are typically dual- or quad-socket servers packaged in 3U or 4U form factors. In addition to supporting more sockets than SFF, they also can hold more memory than SFF devices to support applications that require higher processing capability. These are often called **large-form-factor (LFF) servers** because of their increased rack heights. LFF servers can also be fairly deep, so cabinets should be selected with the increased depths of modern and future ITE in mind.
- **Cloud, volume.** These systems are typically single- or dual-socket boards packaged in 1U, 2U, or half-width form factors. These servers have a limited, targeted set of features selected to address specific workloads.
- **Special-purpose.** Mainframes and custom server designs fall into this category. Features and packaging vary widely, depending on the target customer. Chassis sizes also vary widely and include rack-level servers and multiframe systems.
- **Blade.** Typically, blade servers have a multi-U (7 to 14U) chassis supporting multiple individual servers constructed on independent circuit boards called blades. The blades plug into a common backplane, enabling interconnection of

boards. Cooling, power, and switch functionality are shared among the boards. Power supplies generally serve the entire chassis, and may require as many as six power connections.

Datacom Equipment Airflow

Standardized nomenclature defining the cooling airflow paths for datacom equipment have remained unaltered since 2004 (Figure 4). Most datacom equipment now uses the front-to-rear protocol. The exceptions are some legacy telecommunications equipment and some network switches, which may use a side-to-side protocol, or a mix of side-to-side-to-top and/or to-rear air flows, that are not shown. When airflow does not follow standardized protocol, special rack mountings and/or air deflectors may be necessary to achieve proper cooling in facilities designed predominantly for front-to-rear cooled equipment. This generally requires wider cabinets to accommodate the air diverters, or even sidecar add-ons to contain the baffles for larger devices.

Top-of-rack (TOR) consolidation switches have become common to reduce the cable counts interconnecting dense server cabinets to core network switches. TOR switches are usually small, stackable types that are designed for front-to-back airflow, with the connector side being the front. However, since server connections are in the rear, these switches are often mounted backwards in cabinets, which makes them draw in hot return air from the servers and exhaust even hotter air into the cold air intakes of other equipment. This is obviously not a desirable situation. Many of these switches have an option for reversing airflow, which may require a different fan module if the switches have not been ordered for TOR usage. Even then, switch cooling is likely to be inadequate since these switches are often mounted mid-cabinet to make connections more accessible. Devices are available from several manufacturers to extend air intakes to the cabinet fronts, thereby providing good cooling and true front-to-back airflow through switches with reversed airflow.




Figure 4. Equipment Airflow (ASHRAE 2015a)

Liquid-Cooled Datacom Equipment

The increasing heat densities of modern electronics are stretching the ability of air to adequately cool the electronic components within servers. The trend to higher recommended inlet air temperatures, done with the goal of saving energy, exacerbates the problem. Liquid cooling is therefore becoming more prevalent.

Liquid cooling is defined as the process where a liquid (rather than “fluid” air) is used to provide the heat removal (i.e., cooling) function. There are many different liquid-cooling solutions for datacom rooms. The most common implementations are

- **Liquid-cooled rack:** a circulated liquid provides heat removal (cooling) at a rack or cabinet level for operation. Examples include rear-door or in-rack heat exchangers that transfer a large percentage of the datacom equipment waste heat from air to liquid.
- **Liquid-cooled datacom equipment:** liquid is circulated within the datacom equipment for heat removal (cooling) operation.
- **Liquid-cooled electronics:** liquid is circulated directly to the electronics for cooling, with no other heat transfer mechanisms.

These definitions do not limit the cooling fluid to water. Various liquids can be considered and are in use, including some that could be in a vapor phase through part of the cooling loop.

Figure 5 depicts one example of liquid-cooled datacom equipment where a liquid loop internal to the rack is used to cool the components in the rack. In this case, the heat exchange is with a liquid-to-facility-water heat exchanger. Liquid circulating in the rack must be kept above dew point to eliminate any condensation concerns.

Contamination

Most datacom facilities are well designed and are geographically located in areas with relatively clean environments. Therefore, they do not have significant environmental contamination concerns. However, the overall cleanliness of the environment is only one factor in the long-term cleanliness and location decision of a datacom facility. While potential contamination is often not considered to be a major site selection factor (or is completely overlooked during that process), operational contamination can result in major reliability problems.




Figure 5. Internal Liquid-Cooling Loop Exchanging Heat with Liquid-Cooling Loop External to Racks

There are two types of contaminants: particulate and gaseous. Some datacom facilities may have harmful environments arising from the ingress of outdoor contamination. In many instances, however, contamination is brought into the datacom facility itself. Boxes and dirty equipment stored inside the datacom room accelerate filter clogging in both air conditioners and ITE, resulting in reduced cooling effectiveness, shortened ITE lifespans, and increased energy use. Equipment unpacking and uncrating should never be done inside a datacom room, but that often occurs if a space is not provided for that purpose. (See [Figure 1](#) for a generic illustration.) Internal or nearby printing facilities are also major sources of paper dust, particularly if relative pressures are not considered to avoid pulling contaminants into the datacom space. Facilities without foot wipe pads at entrances are also magnets for contamination.

- **Particulate matter** refers to airborne solid and liquid particles. For the purposes of this chapter, the terms *particle*, *particulate*, *aerosol*, and *dust* are considered equivalent and are represented by the term *particulate matter*. The size of airborne particulate matter can span a vast range from about 0.001 μm to more than 100 μm . Agencies that monitor particulate matter from a health point of view categorize particle mass concentration as $\text{PM}_{2.5}$ and PM_{10} , representing particles smaller than 2.5 μm and 10 μm , respectively. Particulate matter may also be categorized in three size modes: fine (0.001 to 0.1 μm), accumulation (0.1 to 2.5 μm), and coarse (2.5 to 10 μm). Coarse mode is generally limited to particles smaller than 10 μm , but can include much larger airborne fibers and particles. Particulate matter in each of these size categories may be composed of various materials from many different sources.
- **Gaseous contaminants** relevant to information technology equipment (ITE) and datacom equipment reliability include hydrogen sulfide, sulfur dioxide, mercaptans, and oxides of nitrogen, chlorine, and ozone, each of which can produce adverse effects on computer hardware. These harmful gases are by-products of geological, biological, agricultural, industrial, and manufacturing activities. They can, even at low ppb (parts per billion) levels, act alone or in synergy with each other or with particulate matter to corrode metallic materials, causing irreversible damage to circuit boards, connectors, integrated circuits, and other electronic components.

If a datacom facility serving a critical application happens to be susceptible to gaseous or particulate contamination, the consequences could be severe. As a result, it is important to address the potential for contamination and mitigate the risk as much as is practical.

A number of factors can result in an increased failure rate. The change in solder type from lead based to lead free under RoHS, combined with an ongoing miniaturizing of datacom equipment components, has increased the risk. Changes in datacom room temperature and humidity operating conditions combined with a lower priority consideration for the surrounding air quality are other factors of concern.

Contaminants can cause either electrical or cooling failures within datacom equipment. Electrical circuits typically fail in either an open or a shorted condition. Datacom equipment circuits are much finer than normal power circuits, often with conductors smaller than a human hair. They are, therefore, more susceptible to damage, but they fail in a similar manner. The two most common datacom equipment circuit failures are copper creep corrosion and silver creep

corrosion. Printed circuit boards use tiny copper ribbons (lands) with components attached with silver solder, making them highly susceptible to corrosion damage in polluted environments.

Airborne dust contaminants can be detected from detailed visual inspections of filters in the air-handling systems. Gaseous contaminant presence may require seasonal or periodical monitoring and measurement through the use of copper and silver coupon testing in the datacom rooms. The coupons react when exposed to various gases, with the typical exposure period being around one month. A subsequent lab analysis of the coupons can quantify the level of contaminants present.

Filtration systems (particulate filtration or gas filtration units) can be used to mitigate the risk of contaminants in the datacom facility. More information on this topic can be found in ASHRAE (2015a).

Environmental Guidelines for Air-Cooled Equipment

The first edition of ASHRAE's Thermal Guidelines for Datacom Processing Environments in 2004 created a common design point: the inlet temperature for datacom equipment. The 2008 edition expanded the recommended thermal envelope, and the 2011 edition increased the datacom class definitions from two to four, with wider thermal ranges. The fourth edition (ASHRAE 2015a) made significant changes to the humidity ranges as well. All of these changes were made after a great deal of industry study and, in the case of the humidity changes, from the results of a major ASHRAE research study (Pommerenke and Swenson 2014). The fifth edition of *Thermal Guidelines* (ASHRAE 2021) made additional enhancements to the recommend envelope that aid users in making datacom facility reliability and efficiency improvements, particularly when operating with high humidity and gaseous pollutants. These changes were primarily the result of the ASHRAE-funded research project (RP-1755) on the effects of high relative humidity and gaseous pollutants on the corrosion of IT equipment, conducted by Syracuse University from 2015 to 2018 (Zhang et al. 2019). The fifth edition added a new environmental class (H1) for high-density equipment. High-density products that use high-powered components such as central processing units (CPUs), graphics processing units (GPUs), and high density memory require increased cooling. If necessary, these high-density products should have separate environmental controls and may use separate cooling systems to facilitate optimization of cooling efficiency. Important considerations include the following.

Recommended Environmental Range. To achieve both energy efficiency and equipment operating reliability and longevity, facilities must be designed to achieve ambient equipment inlet conditions that fall within the ASHRAE recommended temperature and humidity ranges under normal circumstances. See [Table 1](#) for this range, or use the process defined by ASHRAE (2021).

Table 1 2021 Thermal Guidelines: Equipment Environment Specifications for Air Cooling

Class ^a	Product Operation ^{b,c}					Product Power Off ^{c,d}	
	Dry-Bulb Temperature, ^{e,g} °F	Humidity Range, Noncondensing ^{b,i,j,k,l,n}	Maximum Dew Point, ^{k,o} °F	Maximum Elevation, ^{e,j,m} ft	Maximum Rate of Change, ^f °F/h	Dry-Bulb Temperature, °F	Relative Humidity, ^{k,o} %
Recommended (suitable for all classes; explore datacom facility metrics in ASHRAE [2021] for conditions outside this range)							
A1 to A4	64.4 to 80.6	15.8 to 59°F dp and 70% rh ⁿ or 50% rh ⁿ					
H1 ^o	64.4 to 71.6	15.8 to 59°F dp and 70% rh ⁿ or 50% rh ⁿ					
Allowable							
A1	59 to 89.6	10.4°F dp and 8% rh to 62.6°F dp and 80% rh	62.6	10,000	9/36	41 to 113	8 to 80
A2	50 to 95	10.4°F dp and 8% rh to 69.8°F dp and 80% rh	69.8	10,000	9/36	41 to 113	8 to 80
A3	41 to 104	10.4°F dp and 8% rh to 75.2°F dp and 85% rh	75.2	10,000	9/36	41 to 113	8 to 80
A4	41 to 113	10.4°F dp and 8% rh to 75.2°F dp and 90% rh	75.2	10,000	9/36	41 to 113	8 to 80
H1 ^o	41 to 77	10.4°F dp and 8% rh to 62.6°F dp and 80% rh	62.6	10,000	9/36	41 to 113	8 to 80

Note: For potentially greater energy savings, refer to Appendix C of ASHRAE (2021) for the process needed to account for multiple server metrics that impact overall total cost of ownership (TCO).

- a** Classes A3, A4, are identical to those in the 2011 edition of *Thermal Guidelines for Data Processing Environments*. The 2015 version of the A1 and A2 classes have expanded relative humidity levels compared to the 2011 version.
- b** Product equipment is powered on.
- c** Tape products require a stable and more restrictive environment (similar to Class A1). Typical requirements: temperature between 59°F and 89.6°F, relative humidity between 20 and 80%, maximum dew point 71.6°F, rate of change of temperature less than 9°F/h, rate of change of humidity less than 5% rh per hour, and no condensation.
- d** Product equipment is removed from original shipping container and installed but not in use (e.g., during repair, maintenance, or upgrade).
- e** Classes A1, A2: Derate maximum allowable dry-bulb temperature 1.8°F/984 ft above 2953 ft. Above 7874 ft altitude, derated dry-bulb temperature takes precedence over recommended temperature. Class A3: Derate maximum allowable dry-bulb temperature 1.8°F/574 ft above 2953 ft. Class A4: Derate maximum allowable dry-bulb temperature 1.8°F/410 ft above 2953 ft.
- f** **For tape storage:** 9°F in an hour. **For all other ITE:** 36°F in an hour and no more than 9°F in any 15 min period of time. Temperature change of ITE must meet limits shown in table, and is calculated as maximum air inlet temperature minus minimum air inlet temperature within specified time window. The 9 and 36°F temperature change is considered to be a temperature change within a specified period of time and not a rate of change. See Appendix K of ASHRAE (2021) for additional information and examples.
- g** With diskette in drive, minimum temperature is 50°F (not applicable to Classes A1 or A2).
- h** Minimum humidity level for Classes A1, A2, A3, and A4 is the higher (more moisture) of the 10.4°F dew point and the 8% rh. These intersect at approximately 77°F. Below this intersection, the dew point represents the minimum moisture level, whereas above it, the relative humidity is the minimum.
- i** Based on ASHRAE research and performed at low relative humidity, minimum requirements are
1. Datacom facilities with non-ESD floors and where people are allowed to wear non-ESD shoes may want to consider increasing humidity, given that the risk of generating 8 kV increases slightly from 0.27% at 25% rh to 0.43% at 8% (see Appendix D of ASHRAE [2015a] for details).
 2. All mobile furnishing/equipment must be made of conductive or static dissipative materials and bonded to ground.
 3. During maintenance on any hardware, a properly functioning and grounded wrist strap must be used by any personnel who contacts ITE.
- j** To accommodate rounding when converting between SI and I-P units, maximum elevation is considered to have a variation of $\pm 0.1\%$. The effect on ITE thermal performance in this variation range is negligible and allows use of rounded values of 10,000 ft. Operation above 10,000 ft requires consultation with IT supplier for each specific piece of equipment.
- k** See Appendix L of ASHRAE (2021) for graphs showing how maximum and minimum dew-point limits restrict the stated relative humidity range for each class for both product operations and product power off.
- l** For the upper moisture limit, the limit is the minimum absolute humidity of dew point and relative humidity stated. For lower moisture limit, the limit is the maximum absolute humidity of dew point and relative humidity stated.
- m** Operation above 10,000 ft requires consultation with IT supplier for each specific piece of equipment.
- n** If testing with silver and copper coupons results in values less than 200 and 300 Å/month, respectively, then operating up to 70% rh is acceptable. If testing shows corrosion levels that exceed these limits, then catalyst-type pollutants are probably present and RH should be driven to 50% or lower. See note 3 of Section 2.2 in ASHRAE (2021) for more details.
- o** This is a new class specific to high-density servers. It is at the discretion of the ITE manufacturer to determine the need for a product to use this high-density server class. Classes A1 through A4 are separate.

Allowable Environmental Range. The allowable envelope is where datacom equipment manufacturers test their equipment to verify that it will function within those environmental boundaries. Typically, datacom equipment manufacturers perform tests before product announcement, to verify that products meet all functional requirements within this environmental envelope. This is not a statement of reliability, but rather one of functionality of the datacom equipment. In addition to the allowable dry-bulb temperature and relative humidity ranges, the maximum dew point and maximum elevation values are part of the allowable operating environment definitions.

The allowable envelope has two major purposes: (1) to enable more hours of free cooling by recognizing that outdoor air temperature may exceed the recommended maximum for as long as a few hours during the day, thereby not requiring short-term reversion to mechanical refrigeration; and (2) to relieve operators of concern when cooling equipment is being serviced or experiences partial failure and temperatures rise until service is complete. Temperatures may also rise above the A1 Recommended range during the brief interval between utility interruption and the restoration of full cooling under generator power. Designing to avoid these short-term excursions into the allowable temperature range can require excessive amounts of cooling equipment and reduced energy efficiency, which should be carefully weighed by datacom facility managers before requiring this level of design.

Practical Application. Prolonged exposure of operating equipment to conditions outside its recommended range, especially approaching the extremes of the allowable operating environment for extended time periods, can result in decreased equipment reliability and longevity (server reliability values versus inlet air temperatures are provided in ASHRAE [2021] to provide some guidance on operating outside the recommended range). Long-term exposure of operating equipment to conditions outside the allowable operating environment risks catastrophic equipment failure. With equipment at high power density, it may be difficult to maintain the air entering the equipment within the recommended range, particularly over the entire face of the equipment. Reasonable efforts should always be made to

achieve conditions within the recommended range. However, if these efforts prove unsuccessful, operation outside the recommended range, but within the allowable environmental range, is likely to be adequate, but facility operators may wish to consult with the equipment manufacturers regarding the risks involved.

Environmental Class Definitions for Air-Cooled Equipment. For any piece of datacom equipment to comply with a particular environmental class (ASHRAE 2021), it must be able to reliably provide its full operational capabilities over the entire allowable environmental range, based on nonfailure conditions. The recommended and allowable ranges for each datacom equipment class are given in [Table 1](#). In the fifth edition (ASHRAE 2021), the expanded recommended envelope relies on the datacom facility measuring the corrosion rate from gaseous contamination at least twice a year using silver and copper coupons. The recommended and allowable ranges for each datacom equipment class are given in [Table 1](#). For environments that are shown by testing with silver and copper coupons to have corrosion levels less than 300 Å/month for copper and 200 Å/month for silver, operating at the expanded 70% rh is acceptable since these low corrosion levels suggest that only pervasive pollutants (SO₂, NO₂, and O₂) may be present. Datacom facility environments that are shown by coupon testing to have levels exceeding 300 Å/month for copper and 200 Å/month for silver, or environments that are not routinely tested, should consider maintaining lower humidity limits as shown in [Table 1](#). The allowable environmental ranges for the four datacom equipment classes are illustrated in psychrometric format in [Figure 6](#):

- Class A1: Typically, a datacom room with tightly controlled environmental parameters (dew point, temperature, relative humidity) and mission critical operations; types of products typically designed for this environment are enterprise servers and storage products.
- Class A2/A3/A4: Typically, an information technology space with some control of environmental parameters (dew point, temperature, relative humidity); types of products typically designed for this environment are volume servers, storage products, personal computers, and workstations. Among these three classes, A2 has the narrowest temperature and moisture requirements. A4 has the widest environmental requirements.
- Class H1: Typically, a zone within a datacom facility that is cooled to lower temperatures to accommodate high-density air-cooled products.

Note: Dry-bulb temperatures must be derated based on altitude for all classes. See ASHRAE (2021) for more information on derating methodology.

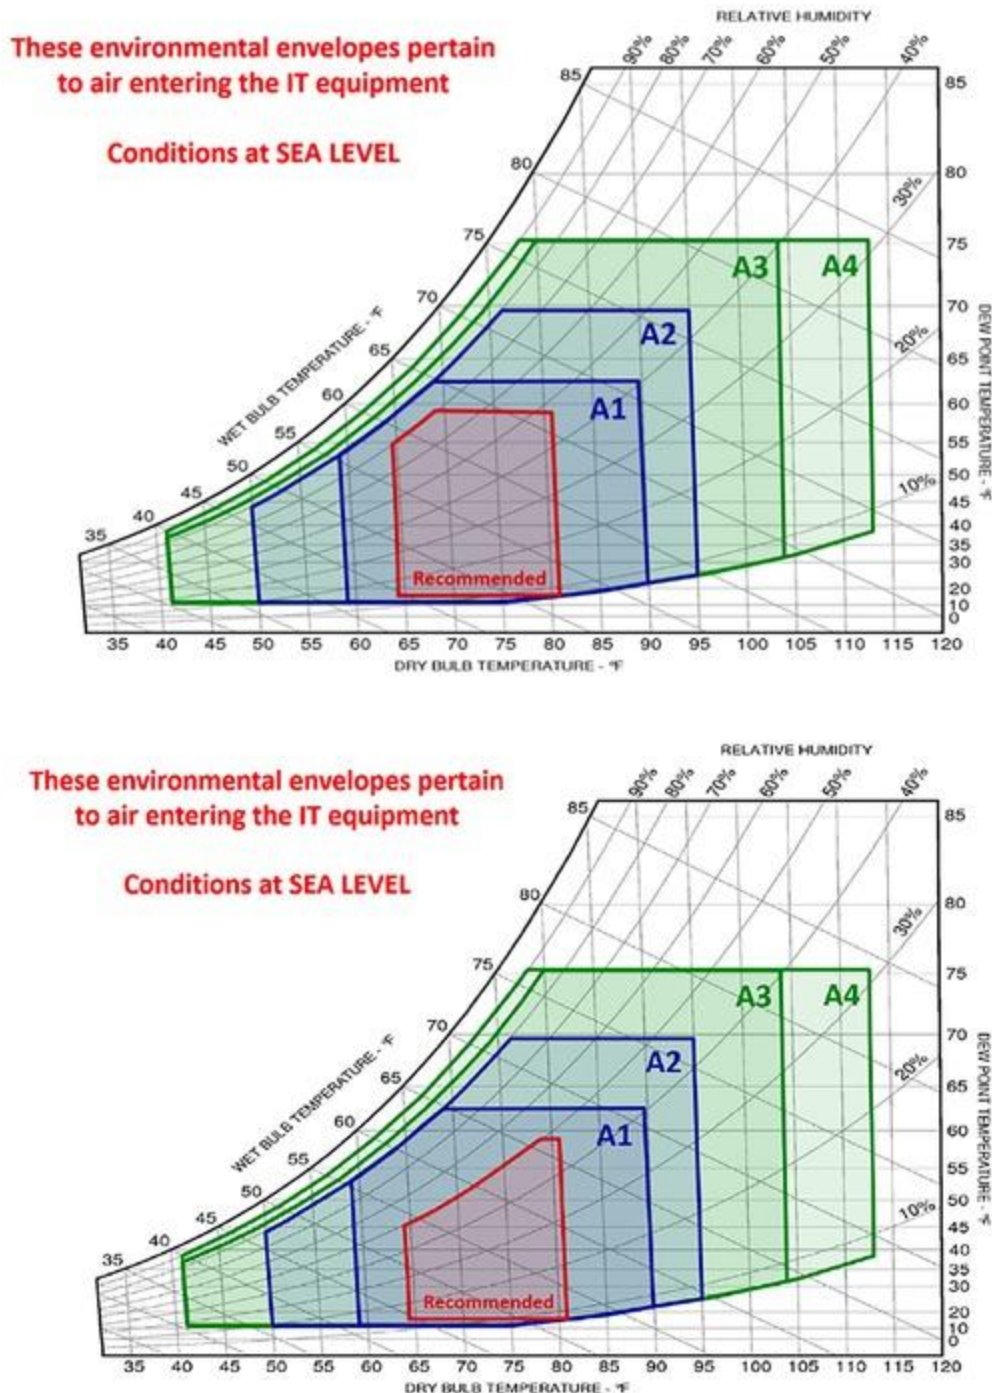


Figure 6. Environmental Classes for Datacom Equipment Classes with Low (Top) and High (Bottom) Pollutant Levels (ASHRAE 2021)

The latest guidelines were developed with a focus on providing as much information as possible, so that datacom facility operators can maximize energy efficiency without sacrificing the reliability required by their businesses. This assumes that the designs enable them to take advantage of reduced energy operation.

Controlling Both Temperature and Moisture in a Datacom Environment to Maintain High Reliability

If both moisture and temperature are properly controlled, the result will be more reliable long-term operation, plus significant energy savings in the operation of the datacom facility. The effects of improper moisture control on a data center operation are twofold:

- **High relative humidity** has been shown to affect failure rates of electronic components. Examples include conductive anodic failures, hygroscopic dust failures, tape media errors, and excessive wear and corrosion. The recommended upper moisture limit is set to limit these effects. New ASHRAE research (Zhang et al. 2019) tested five gaseous pollutants under a variety of temperature and relative humidity conditions. The results suggested that

environments with high levels of gaseous pollutants, or those operating above 70% rh, could be susceptible to increased levels of corrosion.

- **Low relative humidity** has been historically considered a factor in the susceptibility of electronic devices to damage by electrostatic discharge (ESD). However, ASHRAE research (Pommerenke and Swenson 2014) suggests that ESD in low-relative-humidity environments is of far less concern to ITE than once thought.

Based on Pommerenke and Swenson’s (2014) results, the recommended moisture limits were greatly expanded in the fourth edition of the *Thermal Guidelines for Data Processing Environments* (ASHRAE 2015a). The recommended lower moisture limit has been significantly reduced, as shown in [Table 1](#) and [Figure 6](#) for the ASHRAE environmental classes. Note that the recommended upper and lower moisture limits are represented by dew point (dp) limits rather than relative humidity. Although static concerns are actually related to relative humidity, dew point is used because it is fairly constant throughout the datacom facility, whereas relative humidity varies widely due to the range of dry-bulb temperatures across a facility. Because dew point can be easily monitored and consistently controlled, best practice is to monitor moisture content in a datacom room using dew point rather than relative humidity.

The results of Zhang et al. (2019) were incorporated into the fifth edition of the *Thermal Guidelines for Data Processing Environments* (ASHRAE 2021), as shown in [Table 1](#) and [Figure 6](#). To gain the full advantage of the fifth edition, datacom facility operators should use silver and copper coupons inside their datacom rooms at least twice a year (once in the winter and once in the summer) to detect the level of corrosion in the environment. See ASHRAE’s (2014b) *Particulate and Gaseous Contamination in Datacom Environments* for more details on these measurements. This research found that copper corrosion is strongly dependent on relative humidity and pollutant mixture. Having Cl₂ or H₂S caused significant corrosion even with humidity as low as 50% rh. The corrosion rate of silver was found to have no obvious dependence on relative humidity; however, having H₂S present caused significant corrosion on both the copper and silver coupons.

Environmental Guidelines for Liquid-Cooled Equipment

Similar to the guidelines for air-cooled ITE, classes are provided for liquid-cooled ITE. For any piece of datacom equipment to comply with a particular environmental class, it must be able to reliably deliver its full performance over the entire classification temperature range based on nonfailure conditions. The latest version of the liquid-cooling classes (ASHRAE 2021) uses a simple numerical value in the class name to represent the maximum facility supply liquid temperature in SI units (e.g., W17 represents a 17°C maximum facility liquid temperature).

- • Class W17/W27: A datacom facility that is traditionally cooled using chillers and a cooling tower, but with an optional water-side economizer to improve energy efficiency, depending on the facility’s location.
- • Class W32/W40: For most locations, these datacom facilities may be operated without chillers, although some locations may require chillers.
- • Class W45/W+: These datacom facilities are operated without chillers to take advantage of energy efficiency and reduce capital expense.

For datacom equipment that meets the higher supply temperatures as referenced by the ASHRAE classes in [Table 2](#), enhanced thermal designs are required to keep liquid-cooled components within the desired temperature limits. Generally, the higher the supply water temperature, the higher the cost of the cooling solution; however, the lower the cost of the datacom facility cooling solution.

The supply water temperatures in [Table 2](#) are requirements to be met by the IT equipment, It is incumbent on the facility owner/designer to ensure the approach temperature for any planned CDU is taken into account to ensure the proper temperature for the IT equipment.

For classes W17 and W27, the datacom equipment should accommodate facility water supply temperatures that may be set by a campus-wide operational requirement. In these cases, condensation prevention is a must.

Table 2 Liquid Cooled Datacom Facility Classes (Product Operation)

Typical Infrastructure Design			
Class	Main Cooling Equipment	Supplemental Cooling Equipment	Facility Supply Water Temperature, °F
W1	Chiller/cooling tower	Water-side economizer	36 to 63
W2			36 to 81
W3	Cooling tower	Chiller	36 to 90
W4	Water-side economizer (with dry-cooler or cooling tower)	N/A	36 to 113
W5	Building heating system	Cooling tower	>113

Source: ASHRAE (2013).

Liquid-cooled equipment is available from most manufacturers, some even capable of handling W+ environments. The industry is, however, on the verge of large chip power increases that may drive equipment to a lower liquid temperature to support future power density requirements. Facility water flow rate requirements and pressure drop values of the datacom equipment vary. Manufacturers typically provide configuration-specific flow rate and pressure differential requirements that are based on a given facility water supply temperature and rack heat dissipation to the water. Conformance with the water quality requirements for each cooling solution is important to long-term reliability.

Datacom Equipment Nameplate Ratings and Manufacturers’ Heat Release

A power supply nameplate rating indicates the maximum power draw for the datacom equipment’s safety and regulatory approval. A nameplate rating does *not* represent actual power draw during usage and should *not* be used as a means of computing datacom equipment heat release.

Manufacturers that follow ASHRAE guidelines use a template for each product that tabulates heat release based on configuration and use. In addition, most major datacom equipment manufacturers have online tools that can provide even more specific and detailed heat release and airflow information. Obtaining realistic heat release and airflow information is critical to the datacom facility and datacom equipment communities for use in datacom facility planning and designs that are both ample and energy efficient.

Power Trends

Datacom equipment manufacturers compete to create equipment that balances power and performance based on the markets and workloads they are targeting. Datacom equipment is no longer “one size and one configuration fits all.” More IT departments are shifting to purpose-built servers in order to meet customer-specific business needs. These purpose-built servers include specific features and components sized to meet each customer’s workload requirements. This requires facility power projections to comprehend the software workload being deployed. ASHRAE (2018) captures these power trends by select workload ([Table 3](#)) for typical and maximum equipment configurations.

A workload-based methodology provides a much more accurate estimate of actual power consumption in a modern datacom facility, compared to using the maximum power for a given server family from a datacom hardware provider. An example of this trend methodology is shown in [Figure 7](#) (for a 2U 2-socket [i.e., two central processing units] server configuration).

Table 3 Workload Types

Workload Type	Definition/Examples
Scientific	Includes biological sciences, geosciences, weather forecasting, engineering, simulation, design, defense, security, and training of deep machine learning applications (versus run-time)
Analytics	Discrete data warehousing, data analysis, big data analytics, and run-time deep machine learning applications
Business processing	Enterprise-wide line of business applications that manage transactional, operational, and customer databases
Cloud/Internet portal datacom facility (IPDC)	Wikis, portals, social media, video-sharing websites, search engines, and online auction websites
Visualization and audio	Datacom facility visualization applications including video processing, remote visualization, and audio processing
Communications/telco	Wired and wireless networking applications: application, control, packet, and signal processing
Storage	Dedicated storage infrastructure and services including back-up, tiering, and deduplication

Source: ASHRAE (2008).

An important addition to the 2018 edition of ASHRAE’s *IT Equipment Power Trends* was the introduction of the power compound annual growth rate (CAGR) for the years 2016 to 2025. The CAGR allows the measured power of a server running a business’s specific application(s) to be used to project the future power for a similar workload demand. The most striking growth rates ([Figure 8](#)) occur in the scientific (4.6%) and analytics (5.9%) workloads at maximum expected configurations. These higher growth rates can be attributed to higher-power CPUs, maximizing the number of components, and the potential use of graphical processing units (GPUs) or another application-specific processor technology.




Figure 7. ASHRAE Projected Power Trends for 2U 2-Datacom Hardware by Workload Type (ASHRAE 2018)

When appropriately applied, knowledge of datacom equipment power trends can be a powerful tool in considering what future loads might be in a facility or space. Future load is a critical component in the planning, design, construction, and operation of facilities to avoid ineffective expenditures, premature obsolescence, stranded cost or assets, energy waste, etc.




Figure 8. ASHRAE Power Compound Annual Growth Rate for Datacom Hardware by Workload Type and Size (ASHRAE 2018)

Refer to ASHRAE (2018) for details on how the trends were created, along with how to apply them.

2.4 DATACOM EQUIPMENT COMPONENTS

Thermal Design Overview

The goal of a good datacom facility cooling design is to match cooling capacity to actual heat load. This requires a correct and realistic assessment of the heat release of the projected datacom equipment. Even when actual datacom equipment is known, this can be challenging and is often done incorrectly. A basic understanding of datacom equipment

thermal design is therefore valuable to comprehending how the datacom equipment interacts with the datacom facility, and vice versa.

The thermal design must ensure that the temperatures of all datacom equipment components (e.g., processors, memory, storage, I/O, power supplies) are maintained between the high and low limits of their specifications. Datacom equipment components have functional, reliability, and damage temperature specifications. Maximum functional temperature limits for silicon components are generally in the 185 to 221°F range.

The thermal management system (Figure 9) in the datacom equipment must take the appropriate actions to ensure compliance with these specifications. This ensures data integrity and maximizes equipment service life.

A well-designed thermal management implementation balances component temperatures, datacom equipment performance, humidity, and acoustics, to achieve reliable equipment performance with minimal power consumption.

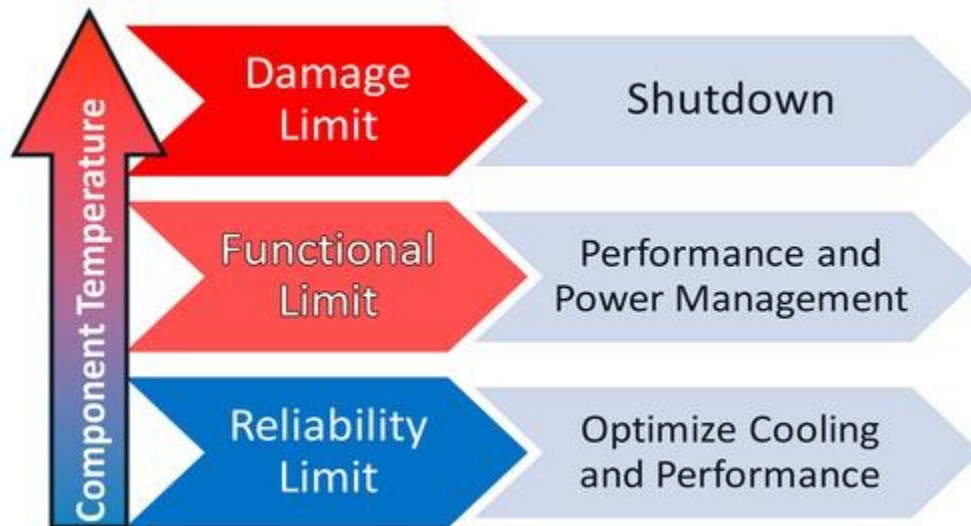


Figure 9. System Thermal Management (ASHRAE 2016a)

Air-Cooled Datacom Equipment Components

Air-cooled solutions remain the most common approach for datacom equipment. The information described here is applicable to most mainstream, air-cooled volume servers; however, the principles apply to most types of datacom equipment.

Typical datacom equipment relies on variable-speed, forced-convection cooling to maintain the required temperatures. Component temperature is driven by one of three factors in an air-cooled system (Figure 10):

- **System ambient**, or inlet temperature to the datacom equipment
- **Air heating**, or increase in air temperature caused by upstream heat sources among the datacom equipment
- **Self heating**, or increase in component temperature above local ambient caused by heat dissipated by the component itself; driven by component packaging, power dissipation, and thermal solution (e.g., heat sink)

Datacom equipment manufacturers develop component- and equipment-level cooling solutions that balance cost, performance, and energy consumption, including trade-offs between air movers and heat sink designs. Cooling performance, power consumption, acoustic signature, fan reliability, and redundancy features are also important characteristics that must factor into the overall solution.

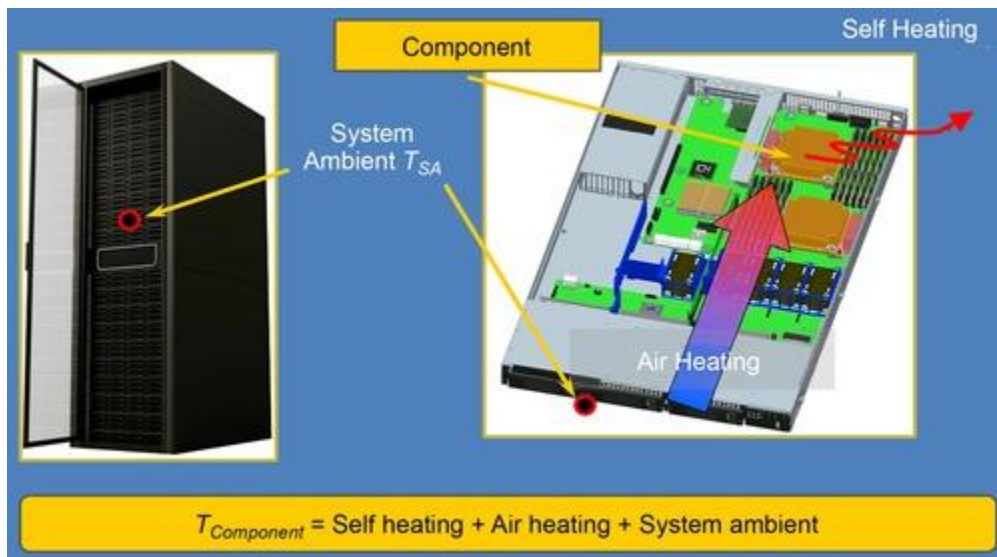


Figure 10. Example Component in System and Rack (ASHRAE 2016a)

Fan speed control and/or cooling zones are often used to precisely adjust specific fans to the needs of the components most coupled with those fans. Cooling zones can be proximity based or physically separated. By using a fan zone approach, total fan power and acoustic output can be minimized. Fans in a non-stressed zone can run at lower speeds than those in a more highly stressed zone.

Power and Thermal Management

Thermal control enables optimization of datacom equipment system performance as a function of usage or workload, configuration, cooling capability, and environment. Underlying this optimization is the use of fan speed control and power management operating in parallel. Optimization for differential air temperature ΔT through the datacom equipment is generally not a significant design consideration because of the more critical requirement of ensuring that functional limits of components are maintained.

Components and their specifications are the primary drivers in a server's thermal design (e.g., heat sink, fan selection, airflow management).

Power management features enable all components to stay within temperature limits while minimizing overall power consumption during periods of low activity.

Highly advanced control algorithms vary the datacom equipment fan speeds and airflows, and tune the fan speeds based on the datacom equipment's usage model. Multiple algorithms can be used simultaneously, with the final fan speeds determined by comparing the results of these algorithms.

Sensors (and proxy sensors) create the data necessary to trigger power management, and are the basis of a cohesive thermal management implementation.

As previously stated, the temperature rise through the server is a byproduct of the thermal management solution. Over time, IT manufacturers have tried to reduce airflow for energy efficiency, which has resulted in higher temperature rise for a given power. For servers operating with an inlet temperature in the recommended range, it is common to have a temperature rise between 36 to 54°F. These thermal efficiencies asymptotically approach a point of maximum optimization as the exhaust-side temperature reaches 140°F. Above this temperature level, different materials might be needed for standard cables. PDU sockets and other power components might also need to be derated, and surface temperatures could become too hot to touch. A consequence of this trend is that future air-cooled IT systems could require larger increases in airflow to keep exhaust temperatures within limits.

It is incumbent on datacom facility designers and operators to ensure that the temperature rise of the cooling solution closely matches that of the IT equipment being deployed.

Liquid-Cooled Datacom Equipment Components

With increasingly dense datacom equipment packaging, some components may require liquid cooling to maintain the environmental specifications dictated by the manufacturer.

Liquids considered for cooling electronic equipment are dielectric, engineered fluids, water with corrosion inhibitors and biocide, oils, or refrigerants. Heat transfer from the liquid-cooled datacom equipment or components to the datacom facility generally takes place through a liquid-to-liquid heat exchanger.

Some liquid-cooling solutions include immersion of the datacom equipment components directly in a dielectric fluid, in either single- or two-phase applications. Dielectric fluids include mineral oil and fluoroketones.

Further details on the implementation of liquid-cooled datacom equipment can be found in *IT Equipment Design Impact on Data Center Solutions* (ASHRAE 2016a).

3. DATACOM FACILITIES

3.1 GENERAL CONSIDERATIONS

Spatial and Envelope Considerations

A datacom facility can be a dedicated building, or be part of a general purpose building that houses other business functions or tenants. Regardless of the type of building in which it is housed, or its location within a building structure, a datacom facility is comprised of a number of spaces having different but interrelated functions (see [Figure 1](#)).

The main computing area is identified by different names: data hall, ITE equipment room, computer room, machine room, raised floor area, or white space. This chapter uses the term *datacom room* to differentiate it from the other areas that support it, and which comprise the complete datacom facility.

Determining the appropriate size of a datacom room is more challenging than it has ever been. There are three major reasons: the ever-increasing demand for computing services; the consolidation, virtualization, and increasing density of datacom equipment; and the transition of many computing services to the cloud, leased colocation facilities, or computing as a service (CAS) facilitators.

It is all too common to underestimate the amount of electrical/mechanical space required to support a datacom facility. The need for reliability in these critical facilities dictates a requirement not just for systems that are often redundant, but also for adequate maintenance space. Overcrowding, even if minimum legal or manufacturer-dictated service clearances are maintained, can lead to inadvertent interruption of one system while servicing another. A rule of thumb, to be used as a starting point only, is to base minimum electrical/mechanical support space requirements on a percentage of the datacom room area:

- At least 50% for non-redundant facilities
- From 75 to 100% for $N + 1$ redundant facilities
- From 100 to 150% for $2N$ redundant facilities

See the section on Redundancy, Reliability, and Concurrent Maintainability for definitions of $N + 1$ and $2N$).

Every increase in reliability requirements also increases the need for additional pieces of redundant equipment, which in turn requires even more support space. Further, highly redundant facilities require physical compartmentalization of duplicate or parallel systems by fire-rated walls, further increasing support space requirements.

The structure enclosing a datacom room should provide good thermal separation from the surrounding areas, whether those are exterior or interior spaces. In addition to thermal insulation, a major concern with the overhead structure is that it not be a source of particulate contamination or water leakage.

The overhead structure must be cleanly finished and sealed to avoid concrete or insulation flake-off. If it is a roof structure, extra precautions must be taken to preclude leakage. In highly critical spaces, a double roof structure is often used for added insurance. Gaps and joints should be caulked.

Suspended ceiling tiles must be either metal pan or plastic encapsulated on both sides to prevent particulate flake-off. This is particularly important when the above-ceiling plenum is used to convey return air. Cut edges must be sealed with spray paint or similar. It is best to suspend overhead infrastructure (lights, cable trays, power busway, air containment panels, etc.) from a structural ceiling grid designed for use in datacom room suspended ceiling installations. This not only provides virtually infinite flexibility in adjusting suspended equipment locations, but also avoids penetrating the ceiling tiles with threaded suspension rods and drilling into overhead concrete to make future changes. Any suspension rods that do penetrate the tiles should be sealed at the penetrations. Metals used above a return air plenum ceiling should be either hot-dip galvanized or of a type that will not grow zinc whiskers.

Walls surrounding a datacom room should be well insulated to avoid both cooling loss and heat infiltration. All cracks should be sealed, which is mandatory if the room is also protected by a gas-based fire protection system.

Although most datacom equipment can accept a broad range of allowable humidity levels, consider installing vapor barriers for datacom spaces. Avoiding condensation anywhere in the room is very important.

Windows should generally be avoided in a datacom room, but if they exist or are somehow necessary, they should be double-glazed and sealed. If covering the windows is allowed, but replacement is only possible from the inside, it will be necessary to make the coverings removable and to avoid blocking access. This may require locating large pieces of mechanical/electrical equipment to the sides of, or well in front of, the windows, which can prevent placing cooling units in ideal locations or waste large amounts of perimeter floor space.

Datacom Rooms

Although raised access floors are still used in many datacom rooms, they are no longer a standard requirement. It is not only possible, but now relatively common, to put the entire power, cooling, and network infrastructure overhead, particularly when close-coupled cooling is used. In these designs, the raised floor is not necessary to convey air.

However, with the amount of piping often used to service liquid-cooling systems (in-row, rear-door, and direct liquid cooling), raised floors are often used anyway to avoid concerns about overhead water, as well as to minimize congestion above cabinets. When power, cable tray, and lighting are all run overhead, the vertical space can become congested and difficult to coordinate. Three-dimensional modeling of the space is highly recommended when overhead infrastructure is used, to avoid both installation conflicts and long-term operational difficulties.

There are several advantages and disadvantages to using raised access floors, regardless of their purpose. The two most common uses are to convey under-floor air to the ITE, and to provide a space for permanent infrastructure such as power, piping, and cabling. An additional advantage can be to level slab variations too large to be leveled with patching and unrealistic for self-leveling cement. Once it is determined that there is a preference for a raised access floor, several factors should be considered in its selection and design.

When used for cooling air delivery, CRAC/CRAH units, which are generally located around the room perimeter, blow cold air under the raised floor, pressurizing the floor plenum. Openings in the floor (airflow tiles) allow the air to be pushed upward in front of cabinets by the under-floor pressure and delivered to the server air intakes. The amount of air delivered through each opening can be regulated by the choice of airflow tiles, which may have different opening percentages (generally 25 or 63%), adjustable air dampers, or both. The supply of air is not infinite and its static pressure is not entirely uniform throughout the under-floor plenum, so delivering more air through one tile can reduce the air delivered through others. Air balancing is therefore important to maintaining good cooling throughout the space. More information is provided in the section on Air Cooling.

A raised access floor also adds total weight to the structure. The floor must also be maintained, which includes re-leveling every few years, particularly if technicians do not take care in replacing tiles where they were removed, or open too many tiles in a row at the same time and destabilize the floor. If not well managed, the plenum space can become a tangle of wire and cable if care is not used in installing new cable and removing old. When the floor is used to convey air, masses of unmanaged cable can reduce or totally block airflow. When cables are located in a raised floor plenum that is also used to convey cooling air, best practice is to run cables parallel to airflow, and to provide overhead cable pathways for ad hoc cable installation to minimize the need to work under the floor. It is even better to not locate cables in an air-plenum floor space at all.

The height of a raised access floor is determined by its purpose. If it is used to convey cooling air, it must be high enough to deliver the required air quantity while maintaining the necessary static pressure as evenly as possible throughout the floor plenum,

Piping, power systems, or cable tray that will also occupy the space must be taken into consideration in determining the floor plenum height and its effect on air flow. It is generally accepted today that a raised access floor used to convey cooling air needs to be at least 24 to 30 in. high to be effective, and that even higher is better.

After height, the biggest consideration is floor structural strength. Raised floors for datacom rooms should use bolted stringer substructures to increase load capacity and to make it easy to remove and replace tiles without destabilizing the floor.

Newer computing equipment cabinets are usually rated for 2500 to 3000 lb., which is significantly higher than most legacy cabinets. Even if they are not full of the heaviest available equipment, floor loading must be planned as if they will be to address potential maximum loading in the future.

However, cabinets with these load ratings also tend to be larger than legacy 24 by 24 in. cabinets, so the load is spread over more than one floor tile. It is not unrealistic to specify raised-floor systems designed for 100 psi or 3000 lb rated tile capacity. Any abnormally heavy equipment can be supported with supplemental floor pedestals under the tiles, so long as the slab structural strength is sufficient for both the total and point loads.

In selecting floor strengths, it is particularly important to examine the rolling load characteristics along with the static structural ratings, because equipment must often be moved into position on small integral wheels. The rolling load tests are generally performed for 10 passes and 10,000 passes with weight on test wheels of particular sizes in accordance with testing methods established by the Ceilings & Interior Systems Construction Association (CISCA) (CISCA 2007). However, not all tests are performed with the same wheel sizes, and the results can be misleading.

It is always safest to put thick hardboard or plywood over the floor when moving particularly heavy loads. This is especially important when rolling equipment through cold aisles with perforated airflow tiles, many of which do not have a rolling load specification at all and are easily deformed.

The most common surface material for raised access floor panels is high-pressure laminate (HPL). This material holds up well to heavy rolling loads without deforming or cracking, has good static dissipative characteristics when properly bonded to a grounded surface, is available in light colors to maximize lighting effectiveness, and is easy to maintain with damp mopping.

Heavy scrubbing, buffing, or waxing should *never* occur in a datacom room. This precludes the use of vinyl composite tile (VCT), pure vinyl tile (which can also be easily deformed under rolling loads), or linoleum (which is also too easily damaged). Carpeting, of course, should never be used in a datacom room, even if it is antistatic, because it both accumulates and generates particulate contaminants. (*Note:* it is generally accepted that the ground resistance of an installed raised-floor panel, when properly connected to a robust grounding system, should be in the range of 10^4 to $10^6 \Omega$ to minimize any potential for static generation.)

One of the most challenging decisions in selecting materials for air plenum raised access floors is the airflow panels. As noted previously, a range of types is available, including legacy perforated tiles (25% open), grate-style cast

aluminum tiles (56 to 63% open), and tiles incorporating directional vanes, air boost fans, and automatic air flow control. (See further discussion of airflow tiles in the section on Underfloor Air Delivery).

All air plenum raised floors leak air. Because cool air is expensive to produce and requires considerable fan energy to distribute through the plenum, this wastes energy. A good-quality raised-floor installation should be tight enough to leak no more than 2% air. This requires not only expert installation, but through-sealing of all cuts around equipment and penetrations, and completely around the floor perimeter.

For datacom rooms designed without a raised access floor, the primary concern is that all power, cooling, and network infrastructure must be routed overhead. Depending mainly on the cooling method used, this can create a congested overhead space that requires careful design coordination and exacting installation. As previously noted, 3D modeling techniques are highly recommended when designing complex overhead systems. It should also be recognized that the cost of overhead infrastructure, particularly if extensive ductwork is necessary, can be very similar to the cost of a raised access floor.

Support and Ancillary Spaces

Space must be allocated within a datacom facility for storing components and material, support equipment, and operating and servicing the datacom equipment. Some ancillary spaces may require environmental conditions comparable to those of the datacom equipment, whereas others may have less stringent requirements. Continuous operation of some support spaces is often vital to the facility's proper functioning.

If tape storage is located outside the datacom white space, its requirements are more stringent than for most ITE equipment, especially for temperature and humidity rate-of-change. Temperature range for tape is usually between 59 and 89.6°F, relative humidity between 20 and 80%, and maximum dew point of 71.6°F. Typically, the maximum allowable rate of temperature change is 9°F in a 1 h period. The maximum humidity rate of change is 5% per hour with no condensation. Appendix K of ASHRAE (2021) provides a more detailed discussion, with several examples of both compliant and noncompliant rate-of-change conditions.

Electrical power distribution equipment can typically tolerate more variation and a wider range of temperature and humidity conditions than datacom hardware. Equipment in this category includes incoming service/distribution switchgear, switchboards, automatic transfer switches, panel boards, transformers, and standby generators. Manufacturers' data should be checked to determine the amount of heat release and design conditions for satisfactory operation. Further information and guidance can be found in IEEE *Standards* 446 and 1100. (*Note:* Both standards are currently inactive and will be superseded by the 3005 and 3003 series in the future.)

Air-Cooled Chillers. Air-cooled chillers used for datacenter applications are selected to operate at elevated chilled-water supply and return temperatures to help building owners achieve energy savings during system operation. In central chiller plants, the chilled-water supply temperatures can be as high as 68°F and the chilled-water return temperature can be 86°F. With these water temperature ranges, the supply air temperature to the data hall from precision air-handling units could be in the range of 77 to 79°F.

Chillers used for datacenter applications must be provided with a quick restart feature. Once power is restored after a failure (either utility or generator), chillers must be able to rapidly resume pre-failure operating capacity. The quick restart time varies by manufacturer, and must be selected to meet the project's specific requirements. The air-cooled chillers used for datacenter applications must be selected for ambient temperatures of $n = 20$ years based on ASHRAE weather data or project-specific requirements. Generally, air-cooled chillers are located at the terrace level of a datacenter building or on the roof, mounted on an elevated structural beam grid system. The chiller platform is elevated approximately 11.5 to 13 ft from the terrace slab or roof. The space below is used to route the main chilled-water headers and other services, including pumps. The elevated platform must be provided with handrails and MS plates of sufficient thickness for service personnel to safely move around the chillers. Chiller compressor and condenser fans should be provided with variable-speed features to minimize energy use during part load operation. Chillers and their controls must have the same design redundancy as has been established for the datacom room cooling units.

Uninterruptible Power Supplies (UPSs). These units come in various configurations, including flywheel and diesel rotary UPS (DRUPS), but most use batteries as the energy storage medium. Regardless of type, UPS systems are usually configured to provide redundancy for the central power buses in mission-critical systems, and typically operate continuously at less than full-load capacity.

UPS power monitoring and conditioning (rectifier and inverter) equipment is usually the primary source of heat release. This equipment usually has self-contained cooling fans that draw intake air from floor level or the equipment face, and discharge heated air at the top or back of the equipment. Air-distribution system design should take into account the position of the UPS air intakes and discharges.

Installation of secondary battery plants as temporary back-up power sources should be in accordance with all applicable codes and standards. If ITE is to operate continuously throughout an emergency or accident (business continuity), then the UPS must be air-conditioned with sufficient redundancy and diversity to provide an operable system. Battery durations of more than 30 to 60 min are usually ineffective unless generators are available to maintain cooling, because UPS systems will overheat and shut down within that time period or less. If generators are used, battery durations of only a few minutes are often sufficient: good generators will generally start and stabilize in less than a minute.

Several types of batteries are used with UPS systems. **Flooded lead-acid (wet cells)** are generally considered the longest lasting, but also present the highest initial cost due to facilities. They require special rooms with, among other things, containment for possible acid spills, deluge shower and eye wash stations, as well as hydrogen gas detection and exhaust fans. It is important to locate the battery room close to the UPS room to minimize loop current losses in the large DC conductors. More commonly used today are **valve-regulated lead-acid (VRLA)** batteries, also known as **sealed cells** or **maintenance-free batteries**. These can be colocated with the UPS system, which may be in the datacom room (IEEE *Standard* 1187). However, if the ambient temperature in the datacom room is not appropriate for the batteries, battery life may be adversely affected. VRLA batteries are available in different qualities, but those usually supplied with a UPS system may have a service life of only three to five years before replacement is required, depending on usage conditions.

The newest battery used with UPS systems is the lithium-ion (Li-ion)), commercially introduced to datacom facility service around 2016. Li-ions are still new to UPS usage but are expected to have longer lifespans than VRLAs, although limited usage data is yet actually available. Other advantages are smaller footprint, lower weight, and lower total cost of ownership despite higher initial cost.

Temperature in a battery area is crucial to the life expectancy and operation of the batteries. The optimum space temperature for lead-calcium batteries is 77°F. If higher temperatures are maintained, it will reduce lead-acid battery life; if lower temperatures are maintained, it may reduce the batteries' ability to hold a charge. Recommended ambient temperatures for Li-ion battery types are higher than lead-acid, but should be verified with the battery manufacturers.

Engine-driven generators used for primary or standby power typically have air-cooled radiators that require large volumes of outdoor air when running. It is best to locate them in a room within the building structure, but if that is not possible, they should be in a high-quality enclosure designed specifically for generators. Generators will incorporate heaters to ensure reliable starting in cold weather, and the enclosure should also be maintained at a temperature that enables servicing in cold weather. Designs should ensure that engine exhaust air does not recirculate back to any building ventilation air intakes. Commonly, up to 72 h of fuel oil storage is required, so fuel oil storage tanks and distribution systems need to be integrated into the overall facility design and planning. The governing codes often mandate specific requirements for containment, location of fuel oil storage, fire resistance ratings, etc. However, as has been unfortunately demonstrated in several weather disasters, local codes may impose requirements that can negate the benefits of generators unless all potential conditions are taken into account in the designs. Fuel tanks are heavy when full, but will start to float in a flooded basement as fuel is used, breaking the pipes. Both fill and pressure relief pipes must be located high enough above ground to remain both accessible and impervious to flood waters. Generators themselves must also be carefully located and protected from potential problems.

Not all generators are diesel. Natural gas and propane are also used, and bi-fuel generators are also available that combine the high energy delivery of diesel with the generally lower cost of natural gas. Diesel fuel maintenance is also critical to ensuring reliable backup.

Something rarely considered is fuel delivery. Have at least two sources for emergencies, and make certain suppliers have their own generators so they can pump the fuel that must be delivered.

Other Systems and Considerations

Fire Protection. Datacom fire protection involves a combination of strategies starting with prevention and continuing through detection, suppression, and response to a fire event. The National Fire Protection Association (NFPA) has several standards addressing design, installation, maintenance, and operation of fire protection systems in datacom facilities listed in the Bibliography. Municipalities may have their own codes, and worldwide other fire protection standards will apply as well. Consult local government regulations.

There are several options for providing fire suppression in datacom rooms, but only two classifications: water and gas. Conventional wisdom, invoked by many code authorities, is that gas protects equipment but sprinklers protect people and structures. However, many older (and even some newer) datacom facilities use a code exemption to allow gas only without any water system at all. Often, however, datacom facilities are equipped with both a sprinkler and gaseous suppression systems for a combination of life, structure, asset, and service protection. Water-based systems in datacom rooms should be dual-interlocked pre-action to minimize the chance of an inadvertent discharge or damage from a leak. Fog mist systems have also been approved in some jurisdictions.

Gas-based systems are of two generic types: **inert gas** consisting mainly of nitrogen and argon that reduce the oxygen mixture below levels required to sustain flame; and **clean agent** systems that use chemicals to remove heat and extinguish the fire. There are even systems that combine inert gas with fog mist, or that maintain a hypoxic environment to prevent combustion altogether. Any system must be approved by local code authorities. The primary reason for gas-based systems is to avoid damage to electronic computing equipment. The primary concern is environmental. All are safe for humans, with certain restrictions differing among products, and all have low ozone depletion (ODP) and global warming potential (GWP), albeit to different degrees. Depending on the specifics of the gas-based system, it may be necessary to also provide a ventilation fan to remove the gas after discharge.

Early warning, aspirating smoke detection systems have also become common in datacom facilities with the goal of detecting and suppressing small fires with portable extinguishers before any kind of full system discharge occurs. It is important to coordinate the locations of aspirating sensing tubes with air-conditioner discharge streams, because high air velocities can negate the ability of the aspirating system to draw in air and detect smoke.

Air containment, either hot aisle or cool aisle, has become a common method of improving cooling performance and reducing energy usage. However, when containment systems are retrofit or designed into a new facility, effects on the required detection, suppression, release system, materials of construction, and prevention of fire must be considered. These important considerations are addressed in detail in the NFPA standards. The added obstructions often necessitate modifications to the suppression systems (sprinklers and/or gaseous agent nozzles) to ensure proper suppression release and dispersion. An alternative to suppression system changes can sometimes be **partial containment**, which has been shown to be as much as 80% as effective in improving air control as **full containment**, but does not block existing sprinkler or gas discharge heads. (See the section on Containment.)

Water Concerns. Water damage is always a concern in a datacom facility, which is why conventional sprinklers are to be avoided. However, it is also best to locate the facility, including its central power and cooling systems, above grade if possible. Where this is not practical, significant attention should be given to mitigating water infiltration.

There are other sources of water leakage as well: designs must consider the possibility of leaks from overhead. Datacom rooms and supporting electrical equipment should not be located below bathrooms, pantries, laboratories, or the like. If unavoidable, the space above should have waterproof membrane floors. Liquid piping should also be routed around the datacom room, but if this is not possible, it should be provided with drip pans and leak detectors. Leak detectors should also be provided anywhere water can infiltrate, particularly if it could affect electrical infrastructure. A common problem in buildings that have not been specifically designed for high-availability datacom facilities is the location of primary power switchgear and bus duct terminations in the lowest level of the building, where they are subject to flooding. These conditions cannot likely be changed in existing buildings, but should be avoided in purpose-built facilities.

Acoustics. The rapid increase in datacom equipment density and power draw has brought with it commensurate increases in required cooling. Air cooling requires substantial volumes of air movement, which generates high noise levels that can have deleterious effects on worker communication, concentration, and health. Extremely high sound pressure levels caused by fire suppression systems (as noted previously) or by elevated datacom system fan speeds, may also lead to reduced hard disk drive (HDD) performance, and even to permanent damage, due to acoustically driven vibrations.

Sound level exposure limits in datacom rooms and their associated mechanical/electrical plant facilities are evaluated in terms of exposure levels and duration through a time-weighted average. In the United States, workplace noise emission is governed by the Occupational Safety and Health Administration (OSHA [Annual]) in Code of Federal Regulations (CFR) 1910.95. Other organizations, such as the National Institute for Occupational Safety and Health (NIOSH 1998) and the American Conference of Governmental Industrial Hygienists (ACGIH), promote even stricter limits. Similar regulations with various thresholds exist in other countries. Acoustical declarations pertinent to the system configuration, operation, and environment can be used to forecast exposure in advance of datacom equipment purchases. Additional noise mitigation measures may include engineering controls, administrative controls, and requiring a hearing protection program.

Sound emissions from heat rejection equipment (cooling towers and/or air-cooled chillers) as well as emergency and/or prime power-generating equipment for the datacom facility's mechanical/electrical plant must also be considered. Noise generated outside the building, typically from rooftop chillers, cooling towers, and generators, must also be mitigated so that sound levels in the building are conducive to conducting normal business activities.

Community sound levels, mostly from exterior heat rejection and power-generating equipment, must typically comply with state, regional, or local noise codes, ordinances, guidelines, and/or regulations. Community sound level limits are typically cited at property lines and/or anywhere on the property of a potential complainant.

Sound levels of exterior equipment during normal, emergency, and test operation should allow for relatively easy communication among service personnel, as well as auditory awareness of vehicle and general service activities in the area. A sound level at or below 70 dBA in service areas and equipment yards, with all equipment operating, is an ideal goal.

Vibration. Vibration levels in datacom facilities must be considered as well. Primary sources of vibration that may adversely affect normal operations of datacom equipment are usually roof-mounted support equipment such as air handlers, cooling towers, chillers, and generators. Similar equipment mounted inside the building may also lead to undesirable vibration levels. However, sources of vibration are not limited to external equipment. Vibration may be induced by the datacom equipment itself (e.g., by rack cooling fans). This form of induced vibration may also pose issues for high-density disk drives. The ASHRAE TC 9.9 datacom book series, especially ASHRAE (2008a), provides valuable resources and guidance on vibration-related matters.

Vibration specifications for all datacom equipment should be obtained from the manufacturers. Subsequently, compliance with the specifications must be demonstrated with respect to ambient conditions, including structural dynamic characteristics of the facility's supporting floor. This may require field measurements to establish realistic vibration demands on datacom equipment and its content. Typical datacom equipment can withstand a vibration range of 0.07g without functional degradation. In fact, most datacom equipment is capable of withstanding higher operational shock and vibration amplitudes.

Many geographical locations around the world are considered seismic zones where special design considerations are stated by national and regional codes and guidelines. It is expected that datacom equipment will experience the highest force and deformation demands when subjected to earthquake-induced vibration. Accordingly, special bracing and safety restraints for much of the infrastructure may be required. Use of structural supports and restraints, such as fastening cabinets to the floor, is generally recommended, even at locations where seismic regulations are minimal or do not

exist. These requirements are intended to ensure continuous operation of mission critical facilities such as datacom facilities, during and after both major and minor seismic events.

As described in ASHRAE (2008a), it is important for both the owner and the designer to understand all potential hazards of the region where the data facility is located, including seismic, wind, fire, flood, etc. Clear operational criteria should be established and used for the design of systems. These may include recommendations for structural restraints and bracing beyond what is required by law. Local code requirements must be identified and understood, as well as the requirements set forth in ASCE *Standard 7*, which provides further information and direction.

Lighting. In datacom rooms, lighting should usually be centered in the aisles, using narrow fixtures with wide horizontal dispersion patterns similar to lighting of library book stacks. Fixtures should be located above equipment cabinets or cable trays where much of the light energy would be wasted. Fixtures are best suspended 8 to 9 ft above the floor so as to deliver maximum illumination over the heads of technicians and into cabinets. This also enables placing fixtures where there is minimal conflict with overhead air diffusers and airflow. Higher mountings may be necessary to clear other overhead infrastructure, but this disperses more wasteful light energy over the tops of cabinets and other obstructions. An illumination level of 30 footcandles on the vertical surfaces of cabinets is generally sufficient. Although lighting is a small part of datacom room energy consumption, and LED has become the light of choice, proximity sensors should be still used to ensure they are not left on when there is no activity.

Emergency lighting in datacom rooms must consider contained aisles as separate rooms for emergency egress purposes. UL 924-listed battery packs can be added to any PoE lighting fixture to make it into an emergency fixture. This is an advantage over conventional lighting, which requires selected fixtures to be designated and provided as emergency luminaires. Whichever emergency system is decided on, it must meet local requirements for illumination and run time in the event of loss of normal power.

Redundancy, Reliability, and Concurrent Maintainability

It is axiomatic that redundant systems should improve the reliability of a facility, but how much redundancy is justified is always a question. Redundancy decisions must consider business and operational needs and energy efficiency, as well as economic justifications. Unfortunately, redundancy alone does not guarantee increased reliability.

It is not uncommon for large investments to be made in duplicate power and cooling systems that have been configured or installed in ways that defeat or greatly compromise their purposes. Consequently, a careful analysis of all possible failure modes should be an integral part of the design phase of any datacom facility.

The primary goal of any datacom facility should be to align the level of redundancy with the organization's applications criticality, available CAPEX and OPEX budget, and risk tolerance. The combination of these three characteristics is not necessarily the same for datacom facility operators within the same industry. Each operator will weigh these three characteristics in their own unique ways.

The most common level of redundancy incorporated into a datacom facility is to provide for **concurrent maintainability**. This requires a design that allows any element in the power and cooling infrastructure to be shut down and removed from service for maintenance without compromising the computing systems that depend on that infrastructure. This level of redundancy requires two independent facility infrastructure paths supporting the datacom equipment, with no single points of failure, and redundant components or systems to ensure concurrent maintainability.

The term **$N+1$** is often used in the industry to describe a level of redundancy. However, this term alone does not provide sufficient detail to fully understand the level of redundancy provided, or to ensure that the resulting facility is actually concurrently maintainable. $N+1$ can refer to component redundancy or systems redundancy; it can consist of a single path with redundant components, or multiple paths with system redundancy.

Higher levels of redundancy require some degree of duplicate systems, such as two identical and fully load-sharing chiller plants with duplicate piping systems or two full-capacity UPS systems. This is known as $2N$ redundancy. (Note that a smaller installation requiring two units to provide $N+1$ redundancy is automatically $2N$ in nature, but will not function as a true $2N$ system unless the rest of the infrastructure is designed accordingly.) An even more stringent design would have duplicate systems, but with additional redundant components in each. Depending on how the additional redundancy is configured, the systems may be known as $2N+1$, in which an additional unit (e.g., a chiller module) is made available to either of the duplicate systems, or $2(N+1)$ in which both redundant systems each have their own redundant modules. Several methods have been developed to classify levels of redundancy and their resulting reliabilities and uptimes (e.g., ANSI/TIA *Standard 942-B*).

It is standard practice to power datacom equipment from an uninterruptible power supply (UPS). UPSs have two main purposes: to isolate the datacom equipment from power line disturbances; and to maintain ride-through power to the datacom equipment until back-up generators start, or alternatively, long enough to accomplish an orderly shutdown if generation is not available or fails. Most modern ITE is **dual-corded**, meaning it has two full-capacity power supplies and two power cords. To maximize the effectiveness of the dual-corded investment, it has become common to run duplicate branch circuits to each cabinet. This often causes owners to also elect to use duplicate UPS systems, resulting in a $2N$ electrical design that may not be matched by the cooling systems or even by the rest of the electrical infrastructure. It is important to recognize the actual facility vulnerabilities when one element of the infrastructure is more robust than the rest, and to not assume that a facility is more reliable than it actually is.

With today's heat densities, datacom systems cannot be maintained for very long on UPS alone. The usual maximum back-up time is 15 to 20 min, before thermal rise causes a shutdown of the datacom equipment and/or the UPS. High-

performance computers may shut down in minutes or even seconds if cooling is interrupted. It is therefore necessary to have a means of maintaining cooling for the most critical systems until either generators start, or systems can be shut down properly.

It is generally impractical to run large cooling systems on UPSs. If this must be done, the cooling equipment's electrical characteristics make it prudent to use separate UPSs. Further, the substantial power draws and high in-rush currents on compressor start-up and cycling require large and expensive UPS systems. DRUPS or motor generator (MG) sets are often used instead of electronic battery systems to support these kinds of loads.

In many datacom rooms it is not necessary or justifiable to provide full business continuity for an extended period, which would require maintaining full cooling with generators. Even with generators, cooling often must be maintained to the most critical computing systems until both generators and full cooling restart. If a chilled-water system with close-coupled liquid-based cooling has been selected, this can be relatively easy to accomplish. There may be sufficient residual water in header pipes to cool critical systems for several minutes if small, auxiliary pumps are provided on UPS. If not, additional water can be stored in tanks.

Long battery life is of no value if the UPS it supports is without cooling. A UPS generates substantial heat under load, as do the batteries when they take full load after a power failure. Batteries also emit heat as they recharge once the generators start. This heat generation should be considered when choosing the location for the UPS, which is too often relegated to a marginal location. This is sometimes in an electrical or mechanical room that generates additional heat, in a corner of a parking garage, or even in a roof penthouse that is exposed to high sun loads. These kinds of locations can dramatically shorten the actual back-up duration of the UPS, particularly if the batteries are also exposed to continuous heat. The cooling system in the UPS room should have the same level of redundancy as the cooling for the datacom room.

Commissioning

Commissioning (Cx) of critical facilities, including datacom facilities, follows the process described in ASHRAE *Standard* 202-2018 and ASHRAE *Guideline* 0-2019. Additional information is also available in [Chapter 44](#) of this volume.

Commissioning Process Applied to Datacom Facilities. The commissioning process for a datacom facility follows the same fundamental processes as commissioning any building. The difference is the nuances related to critical facilities that must be addressed. As in the generic commissioning process, commissioning starts with the programming phase, with the project goals and objectives being defined and captured in the **owner's project requirements (OPR)** document. During the programming phase, expected reliability and availability must be established for critical ITE loads. This is a very different determination than is normal to typical building OPR requirements, and also involves building performance and facility management considerations. Staffing strategies must also be considered as they will influence the design intent. Sites that do not have continuous staff coverage will require a greater level of automation and remote monitoring. All these factors affect the infrastructure topologies, such as the need for onsite storage and spare parts, and the required training program content.

The **commissioning provider (CxP)** ensures those requirements are addressed and satisfied in the engineering and design team's **basis-of-design (BoD)** document, and that they are further refined and expanded upon in the **construction documents (CDs)** by performing iterative design reviews. The CxP must also ensure that drawings and specifications clearly define the commissioning scope for the CxP, and for the other entities that must participate in the commissioning process. These include the engineering and design team, general contractor and respective subcontractors, select suppliers and/or manufacturers, the owner and facilities management staff, and other entities that may be required. Specification 01 91 13 should be drafted by the CxP and submitted to the owner for approval, and then included in the overall construction documents.

In the construction phase, the CxP ensures that the physical construction of the facility is compliant with the design by reviewing submittals, addenda, and bulletins, and performing onsite progress inspections. The construction-phase verifications culminate in ensuring proper startup and checkout of the installed equipment and systems, including leak/pressure testing of ducting and piping systems; clean and flush; test, adjust, and balance (TAB); NETA testing; and initial energization and checkout by the installing contractors, often with manufacturers' support. Startup and checkout, also referred to as **prefunctional testing**, must include verifying all related monitoring and controls are completed by performing **point-by-point, end-to-end validation** of each monitored or control point. Prefunctional testing and resolution of identified discrepancies is a prerequisite for proceeding to the acceptance testing phase.

Acceptance testing typically proceeds from the simple to the complex, starting with components, then equipment, followed by system-level testing. Acceptance testing involves executing site-specific functional performance tests on the critical infrastructure. This is the most critical part of the datacom facility commissioning process, because any potential point of failure or combination of failure scenarios that is not properly tested could one day be the reason a facility fails. Since datacom facility outages are incredibly costly, thorough testing of the completed facility, and particularly of its redundant systems, is critically important. This can involve testing under low, partial, and design-rated loads to verify that power and cooling systems remain stable during anomalies, and that the environment remains with the OPR and design intent parameters.

Load testing of the infrastructure requires the use of artificial, simulated power and heat loads. Since the vast majority of power sent to the ITE is converted into heat, using resistive electrical load banks allows for simultaneous testing of the electrical power system and the heat rejection (HVAC&R) systems. The CxP should develop a load bank plan in collaboration with the engineer of record and the installing trades to plan and execute load testing. The best

practice is to use true **server simulators**, which are resistive load banks of a size similar to typical IT servers that are temporarily installed in the IT racks. These include fan speed adjustment to simulate actual server airflow, including respective differential temperature rises from server inlets to exhausts. When hot/cold air separation is achieved using rack/row hot/cold aisle containment systems, the testing strategies can become tedious and complex and are beyond the scope of this handbook. Interested readers should consult the ASHRAE Datacom books (www.ashrae.org/technical-resources/bookstore/datacom-series) for more information.

Acceptance testing culminates with **integrated systems testing (IST)** to confirm that the overall site can respond to expected anomalies (e.g., power outages) and restore back to normal operations without impact to the critical loads (typically, computers and other ITE). ISTs are also referred to as **pull-the-plug tests**. Due to the inherent complexity and performance requirements for critical facilities, and the specialized technologies and requirements for redundancies, the level of effort required to perform formal commissioning of datacom facilities can be significantly greater than what is required for office buildings. designed and tested to assure very high levels of reliability. CxPs who commission datacom facilities must be subject matter experts in these specialized technologies and system topologies.

Critical facilities typically have very robust and comprehensive monitoring and control systems, which may include mirrored-redundant central plant controllers so the controls do not become single-points-of-failure. Typical building management systems (BMS) are supplemented with sophisticated electrical power monitoring systems (EPMS), intelligent switchgear with mirrored-redundant programmable logic controllers, and include full integration into critical equipment, resulting in literally hundreds of possible monitoring points for a single system. Monitoring and control systems are typically the last to be completed, yet still need to be tested and validated before proceeding with testing the critical infrastructure that they monitor and control. They must be tested and validated as they would for any other system, but it must also be verified that their interface with the critical datacom infrastructure is correct. This requires site-specific prefunctional test scripts with embedded software verification, field panel programming, head-end graphics, monitoring points, and control sequences of operations.

Thorough CxP services include transitioning to the operations phase, which includes ensuring the project is thoroughly documented for posterity in organized **system operations and maintenance manuals (SOMMs)**. The CxP should also be responsible for ensuring the facility O&M staff receive thorough site-specific training, which needs to be substantially greater than for office buildings. Of particular importance is understanding how to maintain continuous operations while concurrently performing equipment maintenance, and without causing inadvertent shutdowns.

3.2 AIR COOLING

Air-Cooling System Configurations

Datacom equipment rooms can be conditioned with a wide variety of equipment types and configurations. These include packaged computer room air-conditioning units and central-station air-handling systems. Air-handling and refrigeration equipment may be located either inside or outside the datacom equipment rooms, and should incorporate economization or free cooling to the greatest extent possible. Some climates and ITE now enable use of free cooling year round. This section concentrates on cooling systems using some type of refrigeration.

The following systems and configurations are the most commonly used solutions to providing sufficient cooling to air-cooled datacom equipment.

Computer Room Air-Conditioning (CRAC) and Computer Room Air-Handling (CRAH) Units. Despite the development of a variety of newer cooling technologies, CRAC and CRAH units remain the most common datacom cooling solutions. They can be packaged in different ways, but are specifically designed for datacom equipment room applications. They should be built and tested in accordance with the requirements of ANSI/ASHRAE *Standard* 127 and ANSI/AHRI *Standard* 1360.

CRAHs are special-purpose chilled-water air handlers designed for datacom applications. CRACs are compressorized cooling systems and are available in several configurations, including direct expansion (DX) air-cooled, DX water-cooled, and versions that include a water- or refrigerant-cooled economizer coil. Dual-fluid systems are available as well, consisting of a chilled-water circuit plus a DX air-cooled or water-cooled circuit. They are used in applications where full redundancy is required with limited space for CRACs, where the buildings central chilled water system is not available 7 × 24 × 365, or where municipal water may be interrupted and reverting to air cooling is necessary to maintain operations. Both CRAH and CRAC units are available in either downflow or upflow designs. Downflow units are used primarily for underfloor air delivery and have top air returns. They can be installed for either free air return or, for more efficient operation, with a ducted return. Downflow units are available in at least three different versions: (1) with standard height and fixed, built-in fans; (2) with standard height and fans that can be lowered into the raised floor on site; and (3) with extended height, where the separate fan section is located in the raised floor and the heat exchanger section is placed on the fan section above the raised floor. Upflow units discharge air overhead, often into ducts, and can have either front, rear or bottom air returns.

CRAC and CRAH units are usually located in the datacom equipment room, but may also be in mechanical galleries adjacent to the datacom room, or installed remotely and ducted to the conditioned space. Ducted designs require consideration of the relatively low-static-pressure designs of CRAC and CRAH units when plug-type fans are used. Long ducted designs may require conventional forward-curved fans that can work against higher static pressures. Although

these fans do not use EC motors, variable-frequency drives (VFDs) can still be used to control fan speed and improve energy efficiency.

If CRACs or CRAHs are located in datacom rooms without a raised floor, or in rooms with a raised floor that is not used for cooling, overhead air delivery is required, which means using upflow cooling units. The return air grilles on these units are typically at the front, so the units should be located at the ends of the hot aisles and the supply air should be ducted into the cold aisles using overhead supply and containment (see the following sections on Hot Aisle/Cold Aisle and Containment).

With either placement, temperature and humidity sensors must be located to properly control air delivery in order to keep inlet air conditions to the datacom equipment within specified tolerances.

Centralized Air-Handling Systems. Traditionally, many telecommunications central office facilities and datacom facilities use overhead or side-wall air delivery on a non-raised floor with central-station air handlers. Larger, centralized air handlers, typically either roof mounted or adjacent to the datacom space, have been gaining popularity as air-side economizer-based solutions have become more common. These air handlers may include DX cooling coils, chilled-water coils, a combination of both, adiabatic cooling sections, and/or indirect economizer solutions (such as air/air heat exchangers). Larger air handlers may use fan arrays consisting of multiple direct-drive plug fans. Side-wall air conditioners deliver supply air from CRAH/CRAC units through side throw grilles installed on the mechanical room walls. Generally, this arrangement is used in large datacom facilities with no raised floor and with hot-aisle containment systems.

Central air-handling systems are also used to serve non-critical areas in a datacom facility, such as office spaces and meeting rooms.

Control of Variable-Speed Fans. Virtually all air-conditioner fans are now speed controlled to minimize energy use when cooling demand is low. Conventional motors can be equipped with variable-frequency drives (VFDs), but they are expensive and add their own inefficiencies. More common today are direct-driven fans using electronically commutated (EC) motors that are highly energy efficient and inherently variable speed using a low voltage control.

There are several ways to sense the cooling environment and send control signals to fans. The most common are underfloor pressure, cold-aisle containment pressure, differential pressure, supply air temperature, and return air temperature. The choice of approach can depend on the air-conditioner design as well as on the room layout and air delivery approach.

Air Distribution

Traditionally, telecommunication spaces have had no raised floor, with overhead ducted air delivery, whereas datacom facilities have used raised-flooring systems as supply air plenums. Although that is changing in many datacom facilities, underfloor air is still often used. See the preceding section on Datacom Facilities for more information on raised-floor selection and installation.

Underfloor Air Delivery. The interstitial space under the raised floor creates a large-volume air plenum that, if properly configured, can deliver relatively uniform air pressures across the entire room area. One of the major benefits of this configuration is that the raised floor represents a common supply air path from multiple cooling units to multiple server racks. However, because the floor plenum is also often used for piping, power and cable, there are many potential obstacles to airflow that upset air distribution, cause significant pressure changes, and can be challenging to mitigate.

Underfloor air is delivered to cold aisles and balanced using a range of airflow tiles. These are available in 25% open, 56 to 63% open, and fan boosted, both dampered and undampered. Tiles with variable dampers enable adjustment of airflow to match the requirements in each part of the floor. Tiles are also available with directional vanes that allow air to be directed toward the cabinets as needed. However, the addition of a damper to any airflow tile results in reduced airflow and cooling capacity, even when the damper is 100% open. Even distribution of air through the airflow panels is primarily a function of the evenness of the underfloor static pressure. However, pressures are also altered by the very existence of the airflow tiles. Therefore, though it may be tempting to install high-airflow tiles everywhere in an attempt to provide sufficient air delivery to all cabinets, the quantities and locations of different tile types must be balanced in conjunction with the available air volume and static pressures. As with any fluid, air will take the path of least resistance, so too many high-air flow tiles in one area can result in air starvation for equipment in other areas.

The static pressure under a raised access floor is not high (generally not over 0.2 in. of water, so it can be difficult for air to be delivered from a floor tile all the way to the top of a tall cabinet. Per thermodynamics, warm air does not rise, as is commonly stated; rather, cool air (being more dense) falls, displacing warmer air and forcing it upward. Cooling with air from below, therefore, is contrary to the laws of physics and, although it has worked very well for decades, it requires more careful control as cabinet heat densities rise. Tiles with integrated air-booster fans or tiles with variable dampers, either manually or motor controlled, can be used in combination with raised-floor pressure-controlled airflow systems to ensure sufficient cold-air supply at any time and at any place, as well as to push air upward. However, just as with too many high-airflow tiles, tiles with integrated air-booster fans may solve a spot cooling problem but impair cooling to other cabinets because fans will take the air they need.

Computational fluid dynamics (CFD) modeling is generally recommended to confirm air flow patterns and to adjust cooling designs for maximize cooling effectiveness. Where redundant cooling systems are used, it is particularly

important to examine raised-floor designs under cooling failure-mode scenarios. Further information on CFD modeling is provided in the section on Computational Fluid Dynamic (CFD) Analysis.

In short, air balance with underfloor systems can be challenging, and specifications should be carefully examined when selecting from the wide range of air flow tiles now available.

Leakage between the floor tile joints will also reduce the expected airflow. Likewise, air will leak and be wasted through unsealed gaps at the raised-floor-to-wall or raised-floor-to-column junctions, and through any floor cutouts for power, cable, or liquid lines from the raised-floor plenum that are not correctly sealed. A properly installed and maintained floor should leak no more than 2% of the delivered air.

Non-Raised Floor Air Delivery Solutions. Datacom rooms that do not incorporate a raised floor, or that use a raised floor for purposes other than air delivery, have several options available. Common air delivery solutions include overhead, above-row, row-based, door-based, and side-wall methods.

These air delivery methods can be used in datacom rooms with or without raised floors. They can be used in new data center designs, or to increase capacity or add redundancy within existing data centers. If used to supplement new or legacy data centers with perimeter CRAC/CRAH cooling units, they can increase cooling capacity in specific areas to address hot spots created by high-density datacom equipment. They can also be very useful in small computer rooms with only a few rows of cabinets. Above-row and row-based air delivery systems are available in refrigerant, condenser water, or chilled water-based solutions.

Overhead Air Delivery. Delivering air overhead can be accomplished in two ways. The more common approach has been via ducts from upflow CRACs or CRAHs. The CRAC or CRAH cooling units are typically located around the perimeter of the datacom room or in an equipment gallery adjacent to the datacom room. The ducts must be large enough to convey the required air volume to the ITE in each aisle, and at velocities and pressures that enable air flow to be easily adjusted and balanced in each aisle. It is important that one duct runout serve only one cold aisle to enable proper adjustment of airflow to the loads in each aisle. This can make planning ducts for future cold aisles and IT cabinet rows challenging. The large ducts can also add to a crowded overhead infrastructure when power, cabling, and lighting must also be coordinated in ways that ensure maintenance access. Irrespective of planning difficulties, overhead air delivery may still provide more effective airflow management than conveying air underfloor because the delivered air volume can be controlled by the measured temperatures in each aisle. This means the system can dynamically increase or decrease airflow to each cold aisle in response to actual conditions.

Above-Row Air Delivery. Above-row air delivery incorporates close-coupled cooling methods by placing cooling units above the datacom cabinets. The cooling units deliver cool air in close proximity to the datacom equipment, and pull hot return air back into themselves from the hot aisle, resulting in very short air delivery and return paths, limited opportunity for air bypass or recirculation, and reduced fan energy. The cooling units are similar in depth and width to the datacom equipment cabinets, so an above-row cooler can be placed above each datacom cabinet or above the datacom row, based on the cooling capacity required. The above-row cooling units can be attached to and supported by the datacom cabinets, or can be suspended from the structure above, which enables datacom cabinets to be moved or replaced without affecting the above-row cooling units. Above-row coolers can use either chilled water or refrigerant for heat transfer.

Above-row air delivery creates a highly efficient, close-coupled cooling design. Above-row units are sometimes combined with under-floor air delivery systems to provide supplemental cooling for high density ITE cabinets, but they can be challenging to integrate with air containment barriers. They can also complicate the integration of power distribution and network pathways above the datacom cabinets. Since they deliver only sensible cooling, some form of humidity control may also be necessary.

Door-Based Cooling. Door-based cooling units incorporate cooling coils into the doors of datacom equipment cabinets. Cooling doors can be mounted on the front, as precoolers, or on the rear to cool exhaust air before it is discharged into the room. Rear-door heat exchangers (RDhX) are the most common, and may be passive, relying on ITC fans to move air through low-resistance doors, or may have integrated fans to enable higher cooling capacities. Door cooling units are affixed to individual cabinets, replacing standard doors, and add several inches of depth. It is critical that units on raised floor systems are coordinated with floor openings, and incorporate air seals if they supplement under-floor air delivery, so that liquid hoses can move without kinking or damage when doors are opened and closed.

Since door cooling units are non-redundant and provide only sensible cooling, they must be integrated with more common cooling systems to maintain cooling when doors are opened or taken offline for maintenance. It is also critical to maintain the temperatures of liquid coolants to the door-based coolers above the dew point, which also means maintaining compatible humidity control within the room.

Although door-based cooling is often supplemental to an overall cooling system, it is possible to use it as a total cooling solution if heat loads and cooling capacities can be matched. In these cases there is no actual hot or cold aisle because intake and exhaust are maintained at the same temperatures, so not only is air containment not required, but it is also necessary to ensure good air circulation throughout the room. Humidity control must then be addressed another way.

The close coupling of cool supply air to datacom equipment air intakes, and warm exhaust air to cooling unit returns, can result in an excellent cooling solution, both in cooling and energy efficiency.

Row-Based Cooling. Row-based cooling is another close-coupled system that can deliver highly efficient cooling and energy efficiency. Units consist of air delivery and cooling equipment installed within cooling cabinets that are similar in

depth and height to the datacom equipment cabinets. Cooling cabinet widths vary with cooling unit capacities and colling type (chilled water, condenser water, or refrigerant), but most are either 12 or 24 in. wide. Row-based cooling is also referred to as **in-row cooling (IRC)** by some manufacturers.

Cooling unit cabinets are intermixed within the rows of datacom equipment cabinets. The numbers and locations of cooling units required will vary for each customer application, and must be verified based on manufacturer recommendations, redundancy requirements, engineering computations, and CFD analysis. It is sometimes satisfactory to place row-based coolers only at the ends of rows, but more often it is necessary to intersperse them between groups of datacom equipment cabinets.

Row-based cooling units supply cold air from their entire frontal surfaces at relative low speed to flood the cold aisles with supply air. Some units incorporate directional vanes to control airflow direction, and others blow air sideways across ITE cabinet faces. Regardless of type, row-based cooling units are most efficient when incorporated into an air delivery containment system, either cold-aisle or hot-aisle.

When row-based cooling is combined with underfloor air delivery from existing perimeter CRAC/CRAH cooling units, cold-aisle containment generally provides an advantage over hot-aisle containment.

Side-Wall Air Delivery. Side-wall air delivery can be used in both raised-floor and non-raised-floor environments. This design supplies air to the datacom room above the floor on which the datacom equipment sits. If a raised floor plenum exists, it is not used for this air delivery system. The goal is to deliver cooling air evenly but at relatively low velocity, so that it floods the room. Air is delivered perpendicular to the IT aisles through the large supply area provided by the long side walls. The external units can be either centralized or modular cooling systems, the latter being positioned adjacent to the datacom room external wall (e.g., indirect evaporative cooling [IEC] systems). Side-wall air delivery systems more generally use large **fan wall units (FWUs)**, which consist of high-capacity cooling coils and plug fan arrays. These units are generally located inside a mechanical gallery

In a side-wall air delivery design, the hot aisle is usually contained so that the warm ITE exhaust air can be routed out of the datacom room via a ceiling plenum and returned to the cooling units. The cooling air from the discrete external cooling units is often delivered into a large plenum, and then through the side wall grilles, in order to create a more even distribution of cooling air. The cooling can be double-sided if opposite walls of the datacom room are both external.

FWUs incorporate control valves to modulate chilled-water flow through the cooling coils to ensure constant supply air temperature. Proper space planning during the concept design stage is critical when using large side-wall delivery systems. Sufficient space must be ensured within the mechanical gallery and corridors to achieve smooth airflow and provide for movement and service of mechanical equipment. The mass loading details of this equipment must be provided to the structural engineers during concept design. Care must be also be taken when designing a side-wall supply system to ensure sufficient width in both side and cold aisles to avoid developing high air velocities as the air turns into the cold aisles. High air velocities create separation and low pressure in front of the IT equipment, particularly at the ends of the aisles, which is especially important to avoid with high-density IT equipment.

Effects of Air Mixing. Air mixing occurs in two ways: (1) when hot air discharged from computing equipment recirculates back to the air intakes, thereby increasing the inlet air temperature at the computing equipment; and (2) when cool supply air bypasses the computing equipment and mixes with hot discharge air, thereby lowering return air temperature. If air paths through or between the datacom equipment racks exist, then some of the cool supply air will bypass the datacom equipment, and some of the discharge air will recirculate back to the front equipment intakes. To avoid air bypass and recirculation, blanking or filler panels must be used in all open and unused equipment rack spaces. The more complete the separation, the more effective and energy efficient the cooling system will be. Aisle containment is a highly effective method of further reducing air mixing problems.

Cool bypass air reduces the return air temperature, decreasing the system ΔT , which in turn decreases the effective cooling capacity of the air-conditioner coils, requiring an increase in airflow to meet the load. Hot-air recirculation increases ITE inlet temperature. If the supply air temperature has been set toward the upper limit of the ASHRAE recommended envelope, this can result in air delivered to the ITE that is warmer than the design temperature, potentially increasing ITE fan speeds and energy use and shortening ITE service life.

Hot Aisle/Cold Aisle. The first step in avoiding air mixing is to arrange cabinets in a hot aisle/cold aisle configuration. This means that racks and cabinets are installed facing back-to-back and front-to-front, negating the effect of legacy cabinet arrangements where hot-air discharge from one row of cabinets directly entered the intakes of cabinets in the next row. Hot aisle/cold aisle configurations assume that all datacom hardware has been designed with industry-standard front-to-rear airflow. Equipment that uses a nonconventional airflow pattern must be dealt with using special racks, cabinets, and air deflectors, as discussed previously, to achieve proper cooling within a datacom room designed predominantly for front-to-back cooling.

Containment. Containment further segregates the supply and return airflow paths by preventing mixing at the top of the equipment racks and at the ends of equipment rows. There are two classifications of containment: hot-aisle containment (HAC) and cold-aisle containment (CAC). Either of these can be full or partial. Full containment requires solid panels above cabinets and sealed to cabinets and ceiling, with door panels enclosing row ends. Partial containment usually means row end doors only, or plastic strips (which still have air leakage). There is also rack-based containment, which is commonly associated with active or passive chimneys. These main types of containment are illustrated in [Figure 11](#).

Computational Fluid Dynamic (CFD) Analysis

Overview. One of the main challenges to maintaining the high availability required for datacom rooms is delivering cooling effectively and efficiently to all the equipment, wherever it is in the room, and by whatever means the cooling is provided. Complexities created by widely variable heat densities, plus the disruptions to airflow patterns created by obstacles in the air paths, make it difficult to envision air movement in the space without modeling tools that provide graphic representations.




Figure 11. Examples of Main Types of Containment

Underfloor cooling can be particularly difficult to envision and optimize because it is driven by many design factors, such as underfloor plenum geometry, piping and cabling obstacles, and perforated tile types and their airflow characteristics. CRAC/CRAH performance parameters (particularly fan type and orientation) and their locations relative to tile rows and cabinet heat loads are also major factors. However, CFD modeling is also very useful for above-floor cooling designs that use overhead or row-based cooling, particularly if more than one type of cooling system is used in the same space.

Reliable CFD simulations require building an accurate 3D computer-generated model of the datacom room and all its components. A model is only as good as its input data, regardless of the program's sophistication, so of most practical importance is the way in which the user defines the space, the cooling, and IT equipment in it. A CFD model needs to represent the physical room geometry, and anything that might add or stimulate airflow or heat transfer, such as fans and vents. It must also include items that impede airflow, such as underfloor pipes and cables, and interactions with the surrounding environment, such as columns and oddly shaped walls and cabinets (the boundary conditions). The granularity requirements for the model may be different at different points in a project, and most definitely after the room is occupied, since the level of actual data will become more refined as the project progresses.

Depending on the sophistication of the modeling tool, CFD models can be run for a total air-cooled room and its equipment or for individual components (e.g., individual racks and cabinets, or even specific pieces of hardware in them). The numerical data obtained from each level is used to process further levels.

Generally it is not necessary to model individual items to evaluate a room design. The physics-based representations for the majority of items present in datacom rooms are likely to be already included in a data center specific toolset. If a general-purpose CFD tool is used, the user may need to create models for a wide variety of components, some of which could have operational parameters that are outside the range of normal building environmental modeling. In any case, the user should verify with the CFD tool provider that all items are appropriately modeled, ASHRAE RP-1675's final report (Hu et al. 2022) provides some guidance for CFD modeling of data centers.

Application. Several simulations are commonly completed for datacom rooms. These may be based either on assumed datacom equipment layouts and projected heat densities, or on actual datacom equipment installations, and often include

- Testing different cooling system types and strategies
- Comparing different arrangements and positions of cooling, power, and computing hardware
- Optimizing cooling paths, including raised-floor height, ceiling height, return air plenum size, duct sizes, and containment
- Testing cooling effectiveness with part-load configurations and examining failure modes, particularly in redundant designs intended to maintain adequate cooling during maintenance and equipment failure conditions

- Determining where the highest-heat-density datacom equipment is best located from a cooling perspective

Although CFD is a powerful tool, it is also easy for it to be misapplied and misinterpreted. Datacom facilities are complex, and infrastructure and equipment must be simplified for models to be practical. It is critical, therefore, that the modeler understands the key elements of the datacom facility and the fundamentals of CFD modeling for the application of CFD to be successful. Proper use of CFD can help the designer identify and avoid most major cooling problems before construction occurs, but expecting the model to be a 100% accurate representation of the finished installation is not realistic.

In the conceptual design phase of most enterprise facilities, the modeler will probably not know detailed information about the datacom equipment types or configurations. Similarly, the precise locations and routings of cables and other physical infrastructure may not be known, and even the cooling system manufacturer or model may not yet have been selected. In such instances there is little point in putting excessive detail into the model, but at the same time the modeler must interpret the results accordingly: that is, they must understand that the predictions are limited to high-level system design decisions and must recognize that performance will likely be a best-case solution because best practice has been assumed. Nevertheless, a knowledgeably developed CFD model at an early stage can help avoid design decisions that would become problematic in the future, such as inadequate space for the required numbers and sizes of air conditioners, or insufficient raised-floor height.

Where real facilities are being modeled, the models must be more representative of the actual installation. This normally means basing the model on a physical survey of the facility, infrastructure, and datacom equipment configurations. Even so, the real infrastructure and equipment cannot be represented in ultimate detail. For example, a bundle of cables will be represented by an approximate obstruction or resistance to airflow rather than explicitly modeling each and every cable.

To ensure that judgments are made appropriately and that the model is accurate, it is best to compare simulation results with actual measurements of airflow and temperature once the facility is completed and operational. This is generally regarded as a **calibrated model**, because actual conditions can be measured and compared. Even this will have discrepancies from reality, which will be difficult, if not impossible, to resolve, but at least they will be known. Then, and only then, should the model be used for sensitivity studies to upgrade the facility, troubleshoot problems, or make deployment decisions.

Although CFD's primary focus for datacom facilities is determining the effectiveness and efficiency of cooling delivery to the computing equipment in the computer room, it can also be used to analyze such things as airflow around air-cooled chillers, condensers, generators, and other critical equipment. However, a CFD analysis on equipment located outdoors cannot use the commercially available CFD tools developed specifically for computer room modeling. Developing a CFD model and configuring the simulation of outdoor equipment is significantly more complex and requires a higher level of expertise from the CFD modeler.

CFD is also a recommended component of The Green Grid's (www.thegreengrid.org) most recent performance assessment tool: the Performance Indicator (PI). The PI is an extension of the PUE metric (see the section on Energy Efficiency) and examines a composite of energy efficiency (PUE), thermal conformance, and thermal resistance. Each of these parameters can be optimized in different ways. The PI illustrates them in a spider diagram format as an aid to achieving an efficient and cost-effective balance among the variables. Linking the PI parameters to a calibrated CFD simulation gives a clearer picture of how cooling is performing in the room, and allows scientific analysis of which physical and operational changes will deliver the most effective improvements.

For further details, see the section on CFD and Flow Network Analysis under the Liquid Cooling heading.

3.3 LIQUID COOLING

Liquid-Cooling System Configurations

The datacom industry is beginning to need to implement alternative cooling technologies required to deal with the heat rejection of the higher-power CPUs, GPUs, accelerators, and other high-power ASIC devices. Simply put, traditional air-cooled technologies are limited in their ability to remove the increasing amount of heat density within silicon while managing the lower chip temperatures required for next-generation high-performance computing applications. Liquid cooling will become essential and may be integrated with a facility-level cooling system in various ways, including the following.

Modular Room-Based Systems. The most common liquid cooling requires that facility chilled water be delivered to a heat exchanger (often called a cooling distribution unit [CDU]) located in or adjacent to the datacom room. The CDU has piping that connects to the datacom equipment; this is called the technology cooling system (TCS). The TCS connections may be to a centralized heat exchanger at the datacom equipment rack or may connect with the datacom equipment itself (e.g., multiple connections per rack). An example of this configuration is shown in [Figure 12](#). The secondary loop or TCS is responsible for distributing the coolant in the loop to pick up the dissipated heat in IT and deliver it to the CDU. The secondary loop consists of supply/return row manifolds, supply/return rack manifolds, server-level cooling loops, hoses and tubes, valves, quick disconnects, sensors, and controllers.

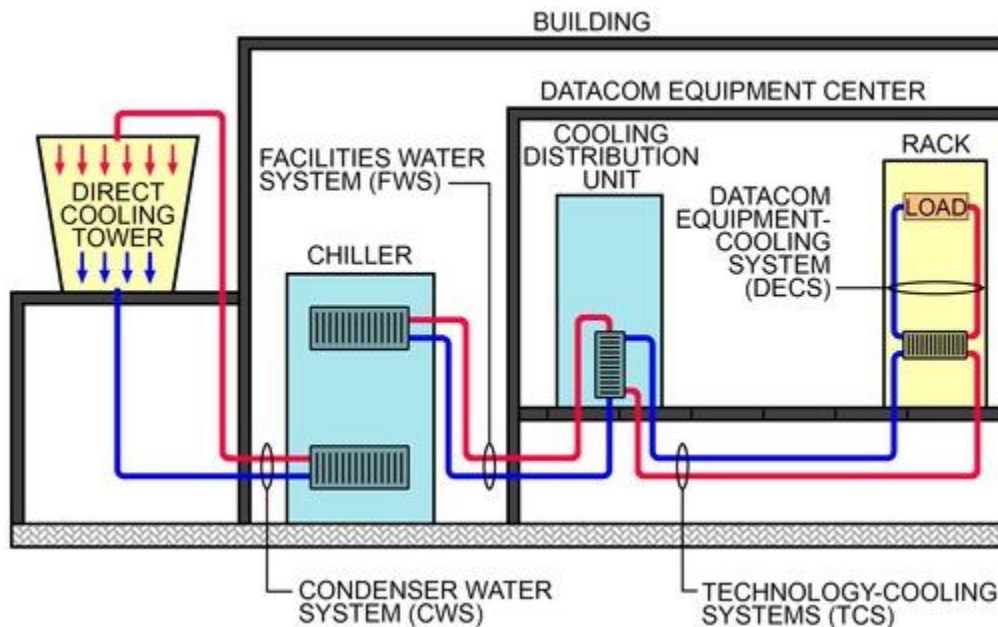


Figure 12. Typical Liquid Cooling Systems/Loops Within Datacom Facility (ASHRAE 2015a)

Quick disconnects in the secondary loop or TCS are required for easy access to each server or rack. Using quick disconnects allows hot-pluggable servers, which means that each server or rack can be disconnected and then reconnected without coolant leaks while the other servers are still operating. Classifications of quick disconnects are necessary as a guideline for selecting the appropriate quick disconnects for liquid-cooling systems. Quick disconnects can be classified based on applications, materials, design parameters.

The fluid in the technology cooling system may be chilled water, deionized water, reverse ionization (RO) water, refrigerant, or other liquids, such as different concentrations of propylene glycol or ethylene glycol. The coolant distribution unit typically also contains pumps, valves, temperature/pressure/flow sensors monitoring and control, and operating software. Refrigerant-based systems have many of the same components as well as compressors, condensers, and/or pumps and related control components. CDUs can only have compressors on the primary side, because compressors can only operate with gases. CDUs that operate with refrigerant on the secondary side need to be operated using pumps, since refrigerant leaves cold plates as a mixture. One of the most important functions of the CDU, or whatever alternative distribution and control mechanism may be utilized, is to maintain coolant temperatures above the dew point. It can be easy to create condensation with liquid-cooled systems if not properly controlled.

Pure (100%) liquid cooling is only available in a full-immersion approach, as discussed later. Except for immersion, the datacom room needs to support a hybrid of air and liquid cooling. The more server components cooled by liquid, the less demand for working fans. To maintain liquid-cooling performance in the event of a failure in the loop, large mutual headers in the secondary piping design may be suitable as reservoirs to keep the coolant temperature within the acceptable range until the failed equipment is restored. This may require small, supplemental pumps on UPS for the most critical equipment. Note that redundancy of the proposed design is of vital importance in the liquid-cooling method. A chilled-water reservoir can also be used as a backup when the primary cooling system fails.

Direct Component Liquid Cooling. This type of system delivers the cooling medium directly to each individual datacom equipment chassis, and then often straight through to the components. This approach is typically used in high-performance computing (HPC) or supercomputing platforms. However, due to the rapidly increasing power densities of components, direct liquid solutions are expected to become more common in a wider range of installations. They require completely dedicated piping distribution, specialized heat exchangers, and related components between the liquid cooling equipment and the facility climate control systems. In this type of cooling system (usually using cold plates on chip, as shown in [Figure 13](#)), the heat is rejected directly to the TCS/FWS water loops, which thereby removes this portion of the heat from the room.

Immersion Cooling. In this type of system ([Figure 14](#)), the datacom chassis are fully or partially immersed in a nonconductive liquid bath. Immersive solutions can be chassis level, where servers mount in standard racks and scale vertically, or tank level, where servers are suspended vertically and scale horizontally. The cooling medium (typically single- or two-phase dielectric) completely surrounds the devices, and circulates through the datacom enclosures or individual chassis subsystems. The pumped fluid transfers the heat to a dedicated coolant-to-water heat exchanger, which is connected to the facility chilled-water or condenser-water loop. The thermal mass of the liquid vat often enables these systems to “ride-through” a cooling failure with little or no supplemental circulation. As with direct component liquid cooling, immersion offers the benefit of rejecting heat directly to the TCS. FWS water loops and thereby from the room completely. In this system, however, with all components immersed in the cooling fluid, nearly 100% of the heat would be rejected, which could effectively remove the need for auxiliary air-cooling infrastructure ([Figure 15](#)).

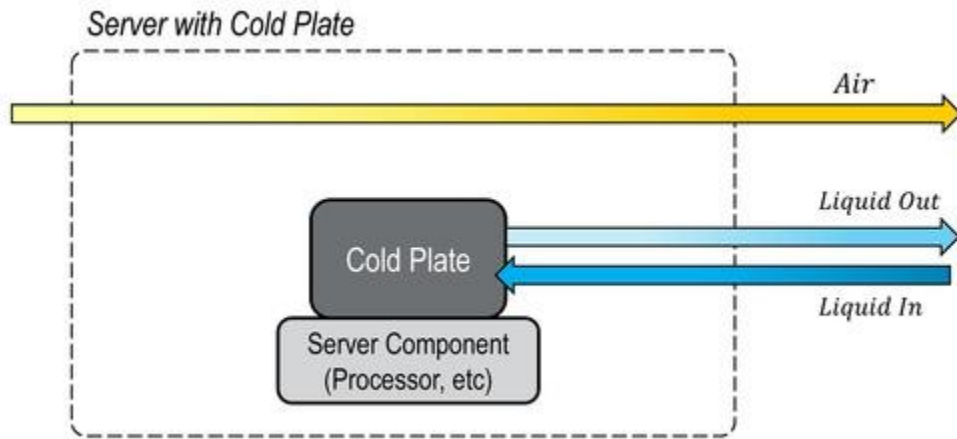


Figure 13. Direct Liquid Cooling with Cold Plate on Computer Chip Created for ASHRAE Handbook by Mark Fisher.

Piping and Distribution Systems

Facility water distribution systems that serve datacom equipment should be designed to the same standards of quality, reliability, and flexibility as other datacom room support systems. Piping quality standards typically require steel or copper piping with welded, bolted flanged, or brazed joints. Although grooved joint coupling failures are rare, they can be sudden and catastrophic when they occur, so measures to prevent such failures are recommended. Where flexible pipe connections are required at equipment, steel braided couplings are preferred over rubber materials. Reliability and flexibility are achieved through a piping configuration that can be expanded or modified as needed to accommodate changes in datacom equipment or major cooling components without requiring extensive system shutdowns. Every section of the main piping, and every major component and valve, should be configured so that it can be isolated and replaced without degenerating the system design below its design level (N , $N+1$, $2N$). Maintenance should include a valve exercise schedule where valves are closed and then reopened annually to loosen debris from the valve seats and break free any rust in rotating parts. Filters and strainers should then be cleaned.




Figure 14. Liquid Immersion Cooling System Servers in Liquid Bath Created for ASHRAE Handbook by Mark Fisher.

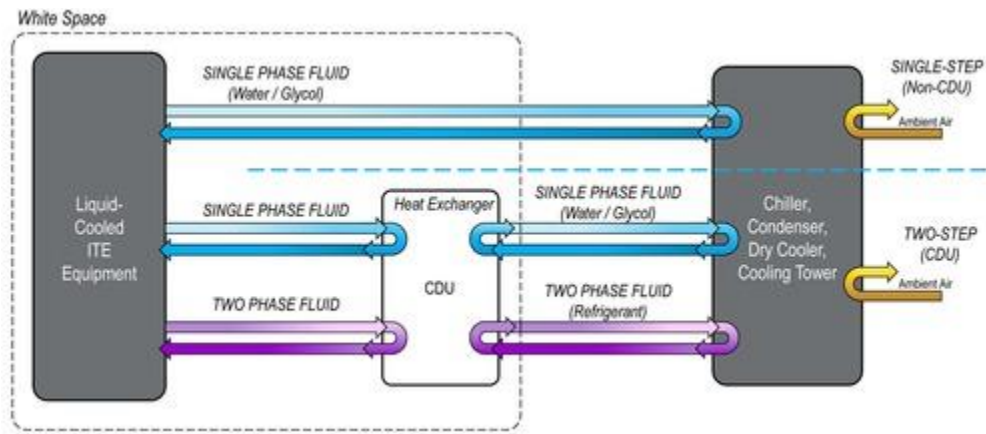


Figure 15. Liquid Immersion Cooling System Approaches Created for ASHRAE *Handbook* by Mark Fisher.

Figure 16 illustrates a double-ended, looped chilled-water distribution system with sectional valves and multiple valved branch connections. The branches could serve air handlers or liquid-cooled datacom equipment, and the valves allow modifications or repairs without a complete system shutdown. Additional piping concepts are detailed in ASHRAE (2015a).

Liquid Coolants

Various kinds of fluids used in TCS cooling loops, including glycol-based liquids, water with additives, refrigerants, and dielectric fluids. The selection of a coolant is a critical task, as one must investigate wetted material compatibility and equipment serviceability, as well as liquid maintenance and operational needs.

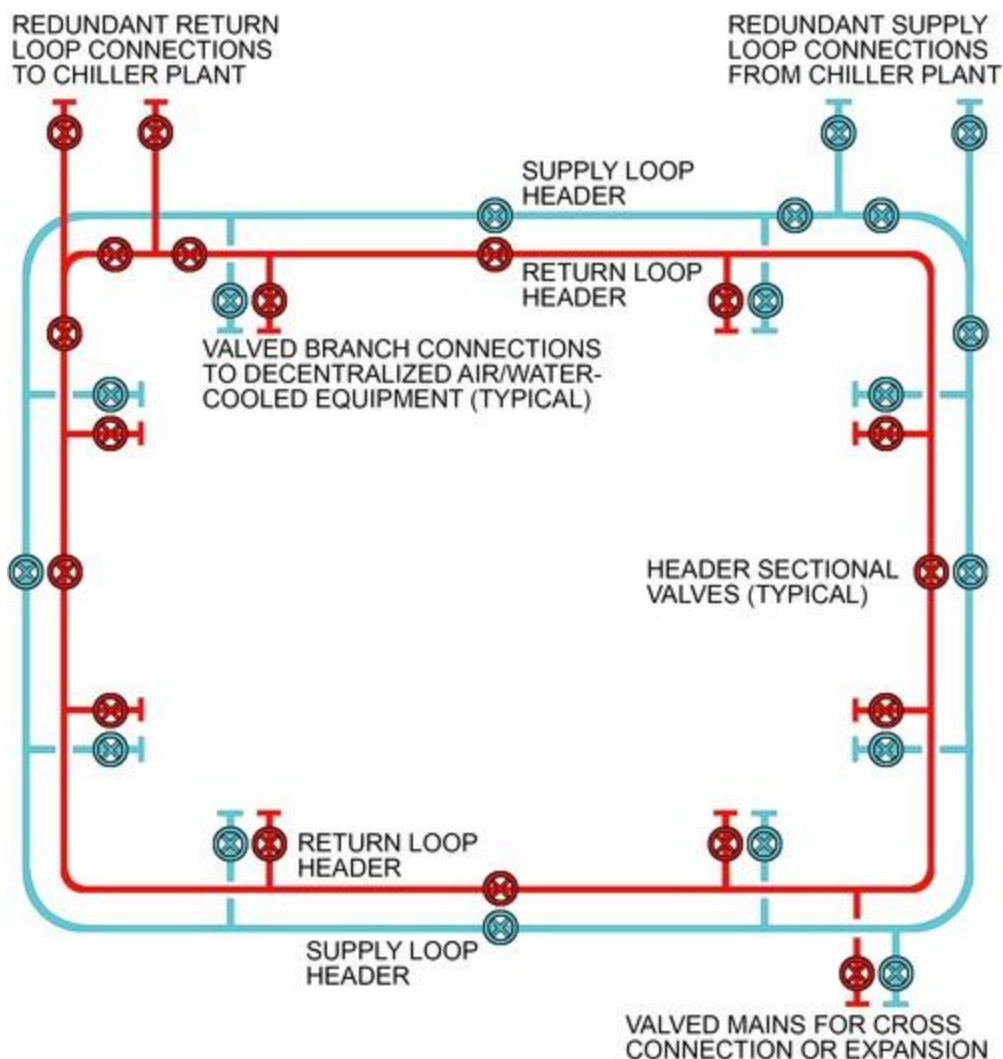


Figure 16. Example of Chilled-Water Distribution Piping System

Water with additives is one of the most common fluids. Water is used as a coolant because of its high thermal mass; however, additives lower its heat transfer capacity so their impact on the heat transfer properties needs to be accounted for. Another concern is that the concentration of additives in water can degrade with time, leading to corrosion and bacterial growth.

Glycol-based fluids are often used to minimize bacterial growth and corrosion, as well as to lower the freezing temperature in cold climates, although they can significantly reduce water heat transfer properties. Propylene glycol and ethylene glycol are the most frequently used additives. At glycol levels above 25% there will be no bacterial growth in the liquid. Glycol-based fluids also have additives and inhibitors, which must be periodically checked for concentration.

Due to their high global warming potential (GWP), most commonly used refrigerants, such as R-134a, are being phased out. As a result, it is critical to find a “green” refrigerant that can be used in the system without negatively influencing its performance. Physical properties, thermal performance, stability, flammability, GWP, toxicity, and availability are all aspects that go into selecting a climate-friendly refrigerant. Preliminary investigation is showing R-515b to be the most appropriate alternative for R-134a in two-phase rack-level cooling systems. Other potentially suitable replacements are R-513a, R-1234yf, and R-471A.

Computational Fluid Dynamic (CFD) and Flow Network Analysis

To perform a complete investigation of a hybrid (air and liquid) solution, both computational fluid dynamics (CFD) and flow network modeling (FNM) should be run in parallel to create a clear picture of the proposed liquid cooling approach. Datacom facility cooling systems can be considered as networks of flow paths through components such as screens, filters, fans, pumps, ducts, bends, valves, heat exchangers, and cold plates. FNM is a generalized methodology of calculating system-wide distributions of flow rates and temperatures in a network representation of a cooling system. Deployment of liquid-cooled server systems that have not been analyzed with FNM might experience issues with pressure, flow rates, or cooling limits. This method focuses on the interactions among the various components incorporated in the cooling network to determine system-wide performance. The thermohydraulic performance of the cooling network is predicted using the imposition of conservation of mass, momentum, and energy in the flow network.

The simulation process starts with a single liquid-cooled electronic model using a CFD tool to measure the thermohydraulic characteristics of the cooling system. Then the performance data are used as the input for a CFD/FNM model of a server cooling loop consisting of several liquid-cooled electronics. All the components (tubes, QDs, valves, tees, etc.) are included in the cooling loop FNM model. Pressure drop and temperature change of the cooling loop are used as the input for the next step, which is rack modeling. In the rack flow network model, elevation effect and uniformity of flow distribution are checked. In the final step, the flow network of the full data center can be modeled, and the secondary flow loop can be analyzed in detail by showing flow distribution in the loop, system pressure, secondary pressure drops, and temperature rise in the secondary loop. Rack and datacom room level CFD analysis explores the air stream circulation, impedance of the system, required air flow rate, supply and return air temperature to racks and cooling equipment, and ASRAE allowable temperature of the ITE.

3.4 WATER USAGE AND ENERGY EFFICIENCY

Conservation and sustainability practices should be at the forefront of every design project today, but in many cases they are now required by code, environmental reality, or both.

Water usage in datacom facilities has gained much attention in recent years. Although water usage does not generally contribute as much to the total cost of ownership (TCO) as does energy efficiency, water has become a very precious commodity in many areas, and there are several environmental regulations that restrict water usage by datacom facilities.

Water Usage Effectiveness (WUE™)

WUE is a site-based metric, developed and popularized by The Green Grid (Patterson et al. 2011), to assess the water used on site for operation of the datacom facility. It presents a comprehensive evaluation of water usage in a datacom facility, where it is affected by a range of factors such as location, IT load, quality of available water source, type and efficiency of cooling equipment, and humidification loads. The formula for calculating WUE is the annual water usage (in litres) consumed by the entire datacom facility, divided by the IT equipment energy usage (in kWh). The physical unit for WUE is liters per kilowatt-hour (L/kWh).

$$\text{WUE} = \frac{\text{Annual water usage, litre}}{\text{Annual power consumed by IT equipment, kWh}}$$

Energy Efficiency

Energy efficiency is at the forefront of modern building design. Datacom rooms are large energy users. They are difficult to consistently operate at peak efficiency because of their dynamic natures, and also because reliability must

usually take precedence over the risks of absolute maximum-efficiency operation. Nevertheless, there are many opportunities to design and operate datacom rooms at high energy efficiencies.

Because cooling typically accounts for the highest energy use (after the IT equipment itself), it is often a primary focus for energy-saving measures such as economization. Datacom facilities were excluded from energy code requirements until the 2010 edition of ASHRAE *Standard* 90.1, which eliminated the exclusion and required datacom facilities to have some means of economization.

In response to industry concerns about potential economizer reliability issues, the challenges of installing economizers on existing high-rise buildings, and the prescriptive nature of *Standard* 90.1, the new ASHRAE *Standard* 90.4, Energy Standard for Data Centers, was developed. *Standard* 90.4 is written to specifically address datacom facility efficiency in a nonprescriptive manner, and to recognize the balance between energy efficiency and reliability that is critical to datacom facility design. This standard is considered a "sister standard" to *Standard* 90.1. It was officially incorporated into *Standard* 90.1 in the 2016 edition as a recognized method of confirming datacom facility energy efficiency in the design stage by using whatever best practices techniques best suit the aggregate needs of the project (e.g., space, location, climate, cooling approach, budget).

Standard 90.4 uses new metrics for both mechanical and electrical efficiencies that were developed specifically to simplify conformance calculations in the design phase of a project, as well as to make it easy to demonstrate compliance to the AHJ. *Standard* 90.4 applies to data centers (called "computer rooms" in *Standard* 90.1) with IT design loads above 10 kW, power densities above 20 W/s, and mechanical and electrical systems dedicated to the datacom facility. Since *Standard* 90.4 is recognized as an alternative compliance path in *Standard* 90.1, *Standard* 90.4 can be used whenever the code enforcement jurisdiction recognizes the applicable 90.1 standard.

Power Usage Effectiveness (PUE™)

PUE is an efficiency metric developed and popularized by The Green Grid. Since the concept was introduced (see, e.g., Rawson et al. [2007]), the metric has been revised to make it more understandable and the methods and reporting of measurement numbers more reliable, culminating in the 2014 release of a joint TGG/ASHRAE TC 9.9 publication (ASHRAE 2014b). In 2012, TGG transferred ownership of the PUE metric to the ISO/IEC JTC 1/SC 39 Work Group 1 standard committee to further develop and expand on the work previously done by TGG. JTC 1/SC 39 WG1 expanded on the PUE metric and published ISO/IEC *Standard* 30134-2, Key Performance Indicators, Part 2: Power Usage Effectiveness (PUE). This standard introduced new PUE derivatives, such as design PUE (dPUE) and partial PUE (pPUE).

PUE measures how effectively an operating datacom facility delivers energy to the datacom equipment inside. The formula for calculating PUE is simply the energy consumed by the entire datacom facility (measured at the meter for the facility or room) divided by the energy consumed by the facility's datacom equipment. The PUE is an average of 12 months of measurement.

$$\text{PUE} = \frac{\text{Total facility energy}}{\text{Datacom equipment energy}}$$

It is important to understand that the PUE metric was developed as a means for individual operations to monitor and track their own energy efficiencies. It was never intended as a means of comparing the efficiencies of different datacom facilities, because too many conditions, including climate zone and level of redundancy, can affect the number. It is also important to understand that an enterprise can take steps to reduce its total energy consumption, yet achieve a worse PUE. Extensive consolidation and virtualization, for example, and the purchase of ENERGY STAR® rated servers, could significantly lower the datacom equipment energy number in the denominator of the PUE equation. However, unless a massive renovation of the power and cooling systems was also done, which would not likely be justifiable in most facilities, the energy consumption of those systems would not be reduced in the same proportion as the datacom loads. Although that would result in a larger PUE quotient, it should still be recognized as a very positive step, because total energy use has still been reduced.

It should also be recognized that the PUE metric, despite the existence of the dPUE metric, is impractical to use as a means of quantifying projected energy efficiency in the design stage of a datacom facility. A significant amount of effort would need to be invested in calculating the large number of electrical path efficiencies and losses, and precision energy modeling would be necessary to develop a realistic number. Note also that the calculated PUE number would still not result in a number likely to be realized when the facility is put into operation, potentially misleading owners into expecting something that is unachievable. It is for these reasons that different metrics were developed for use in the design stage, as set forth in the ASHRAE *Standard* 90.4.

Partial PUE (pPUE) is a complementary metric to the existing definition of PUE. pPUE allows datacom facility managers to easily focus on a certain portion of their facility within a defined boundary condition for the analysis, and to neglect the other components that contribute to total energy. pPUE is used mostly to enable managers to exclude portions of their facility that are complicated to measure from the analysis, or that are intended for use only in certain situations. In short, pPUE is mainly suitable for management purposes, enabling a more complete understanding facility design, whereas the full PUE is the proper metric for evaluating and tracking the energy efficiency of the complete, operating facility.

Partial-Load Operation

A datacom facility is dynamic in terms of electrical and mechanical loading. However, the design of a datacom facility cooling system, whether single plant or modular, must be based on the maximum anticipated datacom equipment load of the space, even though this maximum load is rarely, if ever, achieved.

Even if the maximum design load is someday realized, the day-one load at move-in will most always be much lower than the ultimate design load in order to provide for long-term growth. Additionally, over the course of its lifetime (which may be 10 to 20 years or more), the datacom facility load constantly fluctuates. The load also changes in both density and location within the datacom space as systems are installed in one location and decommissioned in another.

These below-peak, fluctuating loads mean that the cooling plant operates in part-load conditions nearly of the time. It is therefore critical to ensure that the cooling plant selection can modulate capacity over the full, anticipated range, and that it also has good part-load efficiency.

Economizers/Free Cooling

Economizers use ambient outside air temperature to remove heat without using mechanical refrigeration. The process is commonly referred to as **free cooling**. Although it is not totally without cost, in that the mechanical movement of air or liquid is still necessary, the energy to remove the heat is far less than the energy required for refrigeration. There are three fundamental approaches for air-cooled ITE: (1) direct air, (2) indirect air (single-step), and (3) indirect fluid (two-step) economizers, shown in [Figure 17](#).

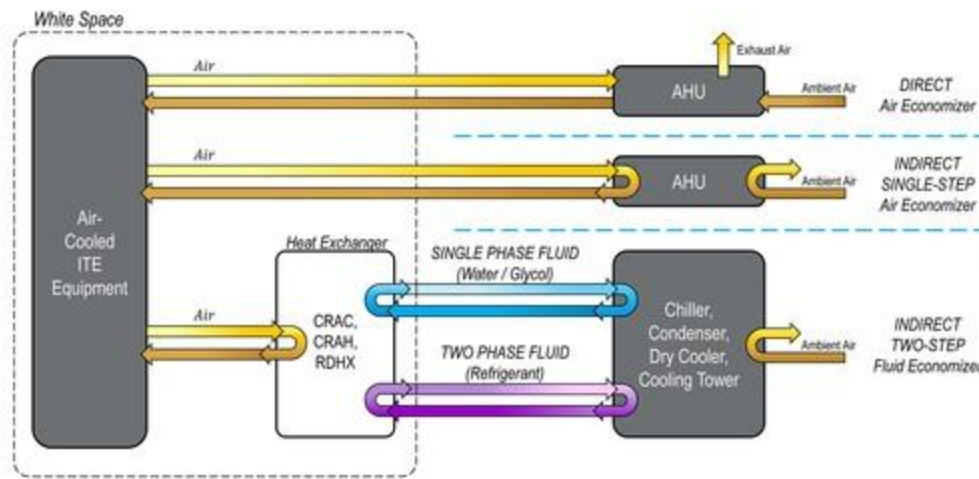


Figure 17. Illustration of Typical Air Economizer approaches Created for ASHRAE Handbook by Mark Fisher.

Direct air economizers are systems that draw air into the data hall directly from the outdoors, and exhaust a comparable amount of hot air back out.

Indirect (single-step) air economizers exchange hot data hall air with cooler outdoor ambient air using an air-to-air heat exchanger.

Indirect (two-step) fluid economizers have two heat exchangers and use an intermediate fluid to move heat between the air inside the data hall and the ultimate heat rejection to the atmosphere. These systems may use single-phase fluids, such as water or glycol, or two-phase fluids, such as refrigerant. Additional steps may be added by including intermediate heat exchangers to transfer energy between different fluids. Although they require more energy than the other two systems to move fluids, indirect, two-step fluid systems still avoid the need for mechanical refrigeration, making them highly efficient economizers.

Economizer systems for datacom facilities can also be differentiated by the type of ultimate heat rejection to the atmosphere: dry, wet, or hybrid. **Dry heat rejection** provides free cooling whenever the ambient dry-bulb conditions are suitable. **Wet heat rejection**, also called evaporative or adiabatic cooling, consumes water to provide free cooling, leveraging the ambient wet-bulb conditions. **Hybrid heat rejection** is able to operate either wet or dry, automatically taking advantage of the benefits of one mode or the other, depending on ambient conditions. Wet heat rejection is typically the most energy efficient, but dry mode saves water and enhance freeze protection, as well as protects against the contingency of drought conditions.

Datacom facilities can often make greater use of economization than office buildings because of the higher temperatures that are usually used. The percentage of the year for which partial (integrated) or full economizer operation is available is a function of the ITE Class (see [Table 1](#)), the cooling system design, and the climatic conditions. Classes H1, A1, and A2 ITE (the most common equipment types) should enable significant use of integrated, or even full, economizer operation for much of the year in most climates. For example, the use of Class A3 ITE, which has an allowable inlet air temperature of up to 104°F, should make it possible to use chiller-less operation with water-side free cooling year-around in many climates. Similarly, for Class A4 ITE, with allowable air inlet

temperatures up to 113°F, year-round chiller-less operation with either air- or water-side free cooling should be possible in most climates.

Economization for datacom facilities is similar to other economization designs, except that reliability, particularly during changeovers between mechanical refrigeration and economization, must be paramount. The unique environmental and reliability requirements of datacom facilities have spawned several data-center-specific cooling products with built-in, innovative economizer cycles. Alternatively, equipment must be selected to enable seamless and highly reliable transitioning. Since the primary energy users in a datacom facility cooling system are the refrigerant compressors, it is common to leverage favorable ambient conditions by integrating economizers into cooling systems so as to minimize annual compressor use. This usually means selecting chillers that can reduce capacity and mix condenser water flow until full economizer operation is achieved.

Detailed Explanations of Free Cooling Approaches. **Direct air economizers**, as shown in [Figure 18](#), introduce ambient air directly into the space so that it flows through the datacom equipment to remove the heat, thereby providing free cooling. This system offers the highest theoretical energy efficiency of any free cooling system. However, since outdoor air is often not as clean or at a humidity level that matches the IT equipment requirements, close attention must be given to the filtration system as well as to supplemental cooling. Supplemental systems may be evaporative (wet, hybrid) or redundant DX (hybrid, dry). Regardless of type, the goal is to ensure adequate cooling when outdoor conditions deviate too much from required free cooling parameters. Note that one reason for the allowable range in the ASHRAE (2021) *Thermal Guidelines* is to enable the continuation of free cooling operation through those periods of the day when outside air conditions exceed the recommended range for relatively short periods of time.

Indirect Single-Step Air Economizers incorporate air-to-air heat exchangers such as plate or heat transfer wheels into the system or AHU. Free cooling can start whenever the ambient air (dry bulb for dry economizers, wet bulb for hybrid economizers) is lower than the datacom facility return air dry bulb temperature. [Figures 19](#) and [20](#) show examples of air-to-air heat exchangers performing this function. In both cases, the heat exchanger provides environmental protection as the datacom room is not directly exposed to outdoor air temperatures, humidity, and/or contaminants.




Figure 18. Typical Direct Air Economizer (Free Cooling) Created for ASHRAE *Handbook* by Mark Fisher.

Indirect Two-Step Fluid Economizers with single phase fluid, use a fluid such as water or glycol to transfer heat from the data hall to an outside heat rejection device. For CRACs ([Figure 21](#)), a room-air-to-water heat exchanger/free cooling coil is added and connected to an outdoor dry, wet, or hybrid heat rejection device such as a dry cooler (dry) or cooling tower (wet, hybrid). The free-cooling coil is located upstream of the mechanical refrigeration coil. The water loop usually serves the free-cooling coil first, and then picks up the heat rejection from the indoor compressors prior to ambient heat rejection. Free cooling in integrated economizer mode can start a few degrees below datacom room return air temperature.




Figure 19. Typical Indirect Single-Step Air-Side Economizer with Air-to-Air Plate Heat Exchanger Created for ASHRAE *Handbook* by Mark Fisher.




Figure 20. Typical Indirect Single-Step Air-Side Economizer with Air-to-Air Heat Transfer Wheel Created for ASHRAE *Handbook* by Mark Fisher.

For CRAHs (Figure 22), a means to use ambient air to remove heat from the recirculating chilled-water system is added. An example is free-cooling coils integrated with the chillers (dry). Free cooling in integrated economizer mode can start a few degrees below the return water temperature to the chiller, precooling the return water and reducing or eliminating the mechanical refrigeration load.




Figure 21. Indirect Two-Step Single-Phase Fluid Economizer with Dual-Coil CRAC Unit Created for ASHRAE Handbook by Mark Fisher.

Indirect Three-step Fluid Economizers with single-phase fluid include an additional fluid-to-fluid heat exchanger between separate fluid circuits. To provide economizer cooling between a condenser water loop and a chilled-water loop, as shown in [Figure 23](#), a liquid-to-liquid heat exchanger is included to precool warm heat rejection water from the datacom room prior to final cooling by a water chiller. Various piping arrangement alternatives are available to allow for bypassing either the heat exchanger or the water chiller depending on the relative temperatures of the two liquid loops.




Figure 22. Two-Step Indirect Single-Phase Fluid Economizer with CRAH and Air-Cooled Chiller Created for ASHRAE Handbook by Mark Fisher.

Indirect Two-Step Fluid Economizers with two-phase fluid use a fluid such as refrigerant to transfer heat via phase change from the data hall to an outside heat rejection device such as an air-cooled condenser or cooling tower. For CRACs ([Figure 24](#)), the air-cooled condenser (dry) can be used for economizer cooling by adding a set of valves, a liquid receiver, and a refrigerant pump. As ambient conditions allow, the compressor(s) are shut off and the pump activates to move the refrigerant between the indoor evaporator coil and the outdoor condenser coil. A typical system has multiple circuits, which allows for partial economizer cooling.




Figure 23. Indirect Three-Step Single-Phase Fluid Economizer with Water-Cooled Chiller Created for ASHRAE Handbook by Mark Fisher.

For CRAHs ([Figure 25](#)), the refrigerant circulates between the evaporator and condenser coils by thermosiphon, though a refrigerant pump may also be used if needed. As ambient conditions allow, the compressor(s) stage off and cooling shifts from the active condenser to the passive (dry) free cooling condenser. A typical system has multiple circuits, which allows for partial economizer cooling.

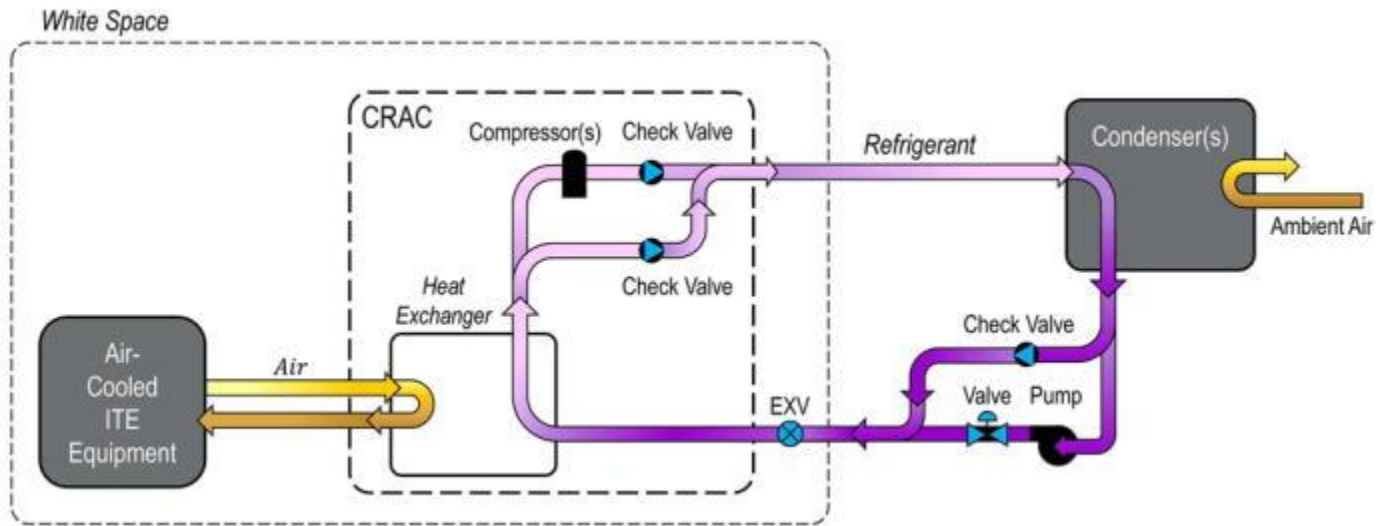


Figure 24. Indirect Two-Step Two-Phase Fluid Economizer with CRAC and Refrigerant Pumps Created for ASHRAE Handbook by Mark Fisher.

X-Factor Reliability Analysis

The x-factor analysis is a unique control strategy developed by the IT sub-committee of TC9.9 to increase the number of hours that can be cooled in 100% economizer mode without negatively impacting ITE reliability on an annualized basis. It uses a floating temperature set point to increase economizer hours without sacrificing server reliability.

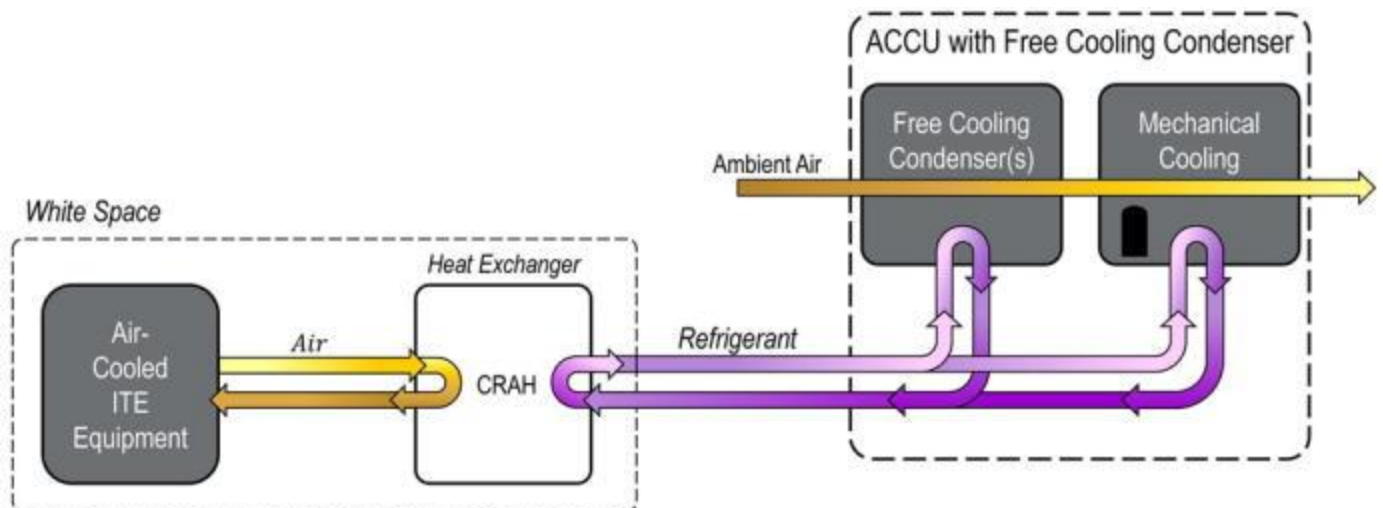


Figure 25. Indirect Two-Step Two-Phase Fluid Economizer with CRAH and Active/Passive Thermosiphon Refrigerant Loop Created for ASHRAE Handbook by Mark Fisher.

ITE equipment is more likely to fail at higher than at lower temperatures, and statistical data on failure rates as a function of inlet temperature has been made available from ITE manufacturers. The x-factor concept is to trade off excursions into the high end of the allowable operating range during the summer with operation at the low end of the allowable operating range in the winter. An annualized x-factor relative reliability value can then be calculated. This concept was first introduced in the 3rd edition of *Thermal Guidelines*. More recently, time-weighted example calculations and reliability tables can be found in Appendices H and I of the 5th edition of *Thermal Guidelines* (ASHRAE 2021).

Liquid Cooling of ITE as Means to Increase Economizer Use

The gradual migration from air cooling to liquid cooling, particularly for high-density ITE, allows for reduced approach temperatures of the heat rejection system and increased economizer operation for the liquid cooling classes (see [Table 2](#)). One advantage of liquid cooling, especially direct-to-chip liquid cooling, is that it can reduce the approach temperature between the chip operating temperature and the ambient heat rejection temperature. This is possible primarily because fewer heat transfer steps can occur between the chip and the ambient air. It is also because liquid-to-component and liquid-to-liquid heat exchangers have lower approach temperatures than heat transfer using air at lower

density. As such, if the choice is available, liquid cooling will likely allow for a greater number of economizer hours than air-cooling, and the potential for cooling without the use of mechanical refrigeration.

4. RESOURCES

ASHRAE Datacom Series

This series comprises 14 books produced by TC 9.9 and is a valuable resource. To keep pace with the rapidly evolving datacom industry, some books have been revised several times, with new editions containing updated information. New titles are also planned for the future.

These books are equally useful for experts and people new to this industry. The following includes brief descriptions of each book.

Thermal Guidelines for Data Processing Environments, 5th ed. (ASHRAE 2021). The trend toward increased equipment power density in datacom facilities presents significant challenges to thermal design and operation. Undesirable side effects include decreased equipment availability, wasted floor space, and inefficient cooling-system operation.

Avoiding a mismatch between datacom equipment environmental requirements and those of adjacent equipment, or between datacom equipment requirements and facility operating conditions, requires a standard practice solution to datacom equipment interchangeability that preserves industry innovation.

ASHRAE (2021) provides a framework to align the goals of equipment hardware manufacturers, facility designers, operators, and managers. This book covers four primary areas: equipment operating environment specifications, facility temperature and humidity measurement, equipment placement and airflow patterns, and equipment manufacturers' heat load and airflow requirements reporting. The fifth edition features clarified wording throughout, changes due to research on the effects of high relative humidity and gaseous pollutants on the corrosion of IT equipment, and a new environmental class for high-density server equipment.

IT Equipment Power Trends, 3rd ed. (ASHRAE 2018). Datacom equipment technology is advancing at a rapid pace, resulting in relatively short product cycles and an increased frequency of datacom equipment upgrades. Because datacom facilities and their associated HVAC infrastructure are typically built to have longer life cycles, any modern datacom facility needs the ability to seamlessly accommodate the multiple datacom equipment deployments it will experience during its lifetime.

Based on the latest information from leading datacom equipment manufacturers, ASHRAE (2018) provides datacom equipment power trend charts through 2025 to allow datacom facility designers to more accurately predict future equipment loads. It also provides ways of applying the trend information to datacom facility designs today.

Included is a review of various air- and liquid-cooling system capabilities and considerations for handling future loads, plus an invaluable appendix containing terms and definitions used by datacom equipment manufacturers, the facilities operation industry, and the cooling design and construction industry.

Design Considerations for Datacom Equipment Centers (ASHRAE 2009a). The design of computer rooms and telecommunications facilities differs in fundamental ways from the design of facilities used primarily for human occupancy. As the power density of datacom equipment continues to increase, this difference has grown more extreme.

This book covers basic design considerations for data and communications equipment centers. The Datacom Facility Basics section includes chapters on datacom design criteria (temperature, temperature rate of change, relative humidity, dew point, and filtration), HVAC load, computer room cooling (including both air and liquid cooling), and air distribution.

The section on Other Considerations includes chapters on ancillary spaces (battery plants, emergency generator rooms, burn-in rooms and test labs, and spare parts rooms), contamination, acoustical noise emissions, structural and seismic design and testing, fire detection and suppression, commissioning, availability and redundancy, and energy efficiency. This book does not cover electrical or electronic systems design and distribution.

Liquid Cooling Guidelines for Datacom Equipment Centers, 2nd ed. (ASHRAE 2013). Datacom equipment today is predominantly air cooled. However, with rack heat loads steadily climbing, the ability of many datacom facilities to deliver either adequate airflow rates or sufficient chilled air is now being stretched to the limit. These trends in the heat load generated by datacom equipment can have detrimental side effects, such as decreased equipment availability, wasted floor space, and inefficient cooling system operation. This situation is creating a need for implementing liquid cooling solutions.

The overall goals of liquid implementations include aspects such as transferring as much waste heat to the facility liquid-cooling loop as possible, reducing the overall volume of airflow needed by the racks, and reducing processor temperatures to improve computer performance.

This book includes definitions for liquid and air cooling as they apply to the datacom equipment, describing the various liquid loops that can exist in a building that houses a datacom space. The book also bridges the liquid-cooling systems by providing guidelines on interface requirements between the chilled-water system and the technology-cooling system. It also outlines the requirements of liquid-cooled systems that attach to an electronics rack and are implemented to help datacom room thermal management.

Structural and Vibration Guidelines for Datacom Equipment Centers (ASHRAE 2008a). The typical life span of datacom equipment is often three to five years, and sometimes far less. On the other hand, the anticipated life span

of the mechanical and electrical infrastructure is 15 to 20 years, and the building's structure can last 20 to 50 years. Consequently, the building's infrastructure and structure may eventually house and support many vintages of datacom equipment.

This book is divided into four main sections. Part 1 gives an overview of the best practices in the design of datacom facilities, including recommendations for new and renovated building structures, building infrastructure, and datacom equipment. Part 2 covers design of new and existing structures. In Part 3, structural considerations of the building's infrastructure, raised-access floor systems, and vibration sources and their control are discussed in detail. Part 4 covers shock and vibration testing, seismic anchorage systems, and analysis of datacom equipment.

Best Practices for Datacom Facility Energy Efficiency, 2nd ed. (ASHRAE 2009b). Sustainable design, global warming, dwindling fuel reserves, energy use, and operating cost are becoming increasingly more important. These issues are even more important in datacom facilities because of their large, concentrated use of energy (which can be 100 times the usage of an office building); 24/7 operations have about three times the annual operating hours of other commercial properties.

The intent of this publication is to provide detailed information to help minimize the life-cycle cost to the client and maximize energy efficiency in a datacom facility.

This book covers many aspects of datacom facility energy efficiency, including environmental criteria, mechanical equipment and systems, economizer cycles, airflow distribution, HVAC controls and energy management, electrical distribution equipment, datacom equipment efficiency, liquid cooling, total cost of ownership, and emerging technologies. There are also appendices on topics such as facility commissioning, and operations and maintenance.

High Density Data Centers—Case Studies and Other Considerations (ASHRAE 2008b). Data centers and telecommunications rooms (datacom facilities) that house datacom equipment are becoming increasingly more difficult to adequately cool because datacom equipment manufacturers continually increase datacom performance at the cost of increased heat dissipation. The objective of this book is to provide a series of case studies of high-density datacom facilities, and a range of ventilation schemes that demonstrate how loads can be cooled using a number of approaches.

Particulate and Gaseous Contamination in Datacom Environments, 2nd ed. (ASHRAE 2014a). Particulate and gaseous contamination monitoring, prevention, and control in datacom environments have gained greater importance because of an increase in datacom equipment reliability concerns arising from many factors: mission-critical societal dependence on computers; continued miniaturization of electronic circuit features; elimination of lead from printed circuit board solder metallurgies; proliferation of datacom equipment into locations with high levels of sulfur-bearing contamination; increased use of free-air cooling to conserve energy; and expansion of the allowable temperature-humidity datacom equipment envelope.

This book describes in detail the procedures necessary to ensure that airborne contaminants will not be a significant factor in determining datacom equipment reliability. It also includes the description of a landmark ASHRAE gaseous contamination datacom facility survey that found that silver corrosion rate is a much better predictor of corrosion-related hardware failures, compared to the prior practice of relying on copper corrosion rate for failure predictions.

Real-Time Energy Consumption Measurements in Data Centers (ASHRAE 2010). Datacom facilities are dense and complex environments that house a wide variety of energy-consuming equipment. With both datacom equipment and associated facility equipment, there are thousands of energy consumption monitoring points accumulated. If a datacom facility operator cannot monitor a device, that device cannot be controlled. In addition, for a datacom facility to reach its optimal energy efficiency, all equipment on the datacom and facilities side must be monitored and controlled as an ensemble.

Datacom equipment and facilities organizations in a company typically have different reporting structures, which results in a communication gap. This book is designed to help bridge that gap by providing an overview of how to instrument and monitor key power and cooling subsystems in ways meaningful to both groups. It also includes numerous examples of how to use energy consumption data in calculating power usage effectiveness (PUE).

Green Tips for Data Centers (ASHRAE 2011). The datacom industry is focused on reducing energy. This focus is driven by increasing energy costs and capital costs to add more datacom facility capacity. Combined with the rapid growth in the industry, and the increase in the power used by the datacom equipment, it is important that every datacom facility operator understands the options for reducing energy.

This book gives datacom facility owners and operators a clear understanding of energy-saving opportunities. It covers the building's mechanical and electrical systems as well as the most promising opportunities in technology. In addition, the book's organization follows a logical approach that can be used for conducting a preliminary energy assessment.

PUE™: A Comprehensive Examination of the Metric (ASHRAE 2014b). Power usage effectiveness (PUE), the industry-preferred metric for measuring the actual infrastructure energy efficiency for datacom facilities, is an end-user tool that helps boost energy efficiency in datacom facility operations. This book provides a high level understanding of the concepts surrounding PUE, plus in-depth application knowledge and resources for those implementing, reporting, and analyzing datacom facility metrics.

It gives actionable information useful to a broad audience, ranging from novice to expert in the datacom equipment industry. This includes executives, facility planners, facility operators, datacom equipment manufacturers, HVAC&R manufacturers, consulting engineers, energy audit professionals, and end users.

PUE was developed by The Green Grid Association, a nonprofit, open industry consortium of end users, policy makers, technology providers, facility architects, and utility companies working to improve the resource efficiency of information technology and datacom facilities worldwide. Since its original publication in 2007, PUE has been globally

adopted by the industry, and The Green Grid has continued to refine the metric measurement methodology with collaborative industry feedback. For further details, see the section on Power Usage Effectiveness in this chapter.

Server Efficiency—Metrics for Computer Servers and Storage (ASHRAE 2015b). This book consolidates information on current server and storage subsystem energy benchmarks for use in selecting the appropriate IT hardware solutions. Each chapter describes a metric and its target market, includes examples of data generated from the subject benchmark or tool, and provides guidance on interpreting the data. This book supplies the information needed to select the best measure of performance and power for a variety of server applications.

IT Equipment Design Impact on Data Center Solutions (ASHRAE 2016a). With everything from smart phones to thermostats generating data, back-end IT systems are experiencing massive hardware demands. Datacom facilities must have a footprint that is flexible, scalable, and adaptable. This book provides guidance in making the critical datacom facility infrastructure equipment selections and design configurations best suited for the evolving datacom facility.

Advancing DCIM with IT Equipment Integration (ASHRAE 2019a). The modern datacom facility is inundated with hundreds to thousands of connected devices and sensors. Data center infrastructure management (DCIM) tools provide visibility across both the management and operations layers to more holistically optimize the cost and performance of the datacom facility. This book depicts how a well-implemented and maintained DCIM system helps safely maximize the efficient use of power, cooling and space resources through a comprehensive connective framework.

ANSI/ASHRAE STANDARD 90.4-2019, ENERGY STANDARD FOR DATA CENTERS (ASHRAE 2019B)

This standard provides a performance-based (non-prescriptive) alternative to ANSI/ASHRAE *Standard* 90.1 for demonstrating compliance with minimum datacom facility efficiency in the design stage. It balances the need for energy efficiency with the concurrent need for reliability in high-performance datacom facilities.

ANSI/ASHRAE STANDARD 127-2020, METHOD OF TESTING FOR RATING AIR CONDITIONING UNITS SERVING DATA CENTERS (DC) AND OTHER INFORMATION TECHNOLOGY EQUIPMENT (ITE) SPACES

This standard establishes a uniform set of test requirements for rating air conditioning unit classes that are used to condition datacom (DC) and other information technology equipment (ITE) spaces. Such units must be able to be tested using an air enthalpy method and facilitate heat transfer across at least one heat exchanger. The standard provides dual units of measurement (SI and IP).

ANSI/AHRI STANDARD 1360-2017, PERFORMANCE RATING OF COMPUTER AND DATA PROCESSING ROOM AIR CONDITIONERS

This standard establishes a uniform set of rating conditions for each of the various types of computer and data processing room (CDPR) air conditioners. It incorporates the set-up and testing requirements from ASHRAE *Standard* 127 and 37.

DATA CENTER HANDBOOK, 2ND ED. (JOHN WILEY & SONS, 2021)

A comprehensive summary of datacom facilities design, construction and management in 37 chapters written and edited by 58 recognized experts in their fields.

ANSI/TIA STANDARD TIA-942-B-2017, TELECOMMUNICATIONS INFRASTRUCTURE STANDARD FOR DATA CENTERS

The Telecommunications Industry Association's *Standard* TIA-942 specifies minimum requirements for telecommunications infrastructure of data centers and computer rooms, including edge data centers, single-tenant enterprise data centers and multitenant Internet hosting data centers.

The TIA-942 standard provides requirements for a wide range of subjects related to the design of data centers, including telecommunications cabling, structured cabling to support intelligent building infrastructure, telecommunications pathways (e.g., conduits and cable trays), cabinets, computer rooms, telecommunications entrance rooms, coordination with access providers (carriers), site selection, and data center rating.

ANSI/TIA-942-B Addendum 1, published January 2022, describes infrastructure requirements and design guidelines of edge data centers (EDC). These data centers are smaller than core data centers and support edge computing in order to reduce data processing latency. They are operated without staffing and need significant detailed remote monitoring.

ANSI/BICSI *STANDARD* 002-2019, DATA CENTER DESIGN AND IMPLEMENTATION BEST PRACTICES

This standard provides requirements, guidelines, and best practices applicable to all types of data centers, including enterprise, colocation, hyperscale and edge, being planned, constructed or in operation. Requirements and recommendations cover site location, construction, power supply, air conditioning and ventilation, cabling and network, security, buildings systems, commissioning, and maintenance. This standard also covers assessing risk, reliability, growth, and other factors as part of the design process to balance function, performance, and business needs.

ANSI/ASHRAE *STANDARD* 202-2018, COMMISSIONING PROCESS FOR BUILDINGS AND SYSTEMS (ANSI APPROVED; IES CO-SPONSORED)

This standard can be considered the minimum acceptable commissioning scope and list of activities necessary to adequately commission a facility, including those that house data centers and other critical facilities. *Standard* 202 describes the commissioning process, the roles of the principal agents and stakeholders, and provides a framework for developing design documents, specifications, procedures, documentation, and reports that streamline the ultimate commissioning process. It also describes the general requirements for a training program to ensure continued successful system and assembly performance. As an enforceable standard, it is written in code language and sets the minimum requirements necessary to meet the intent.

ASHRAE GUIDELINE 0-2019, THE COMMISSIONING PROCESS

This guideline presents best practices for applying whole-building commissioning to facilities, including critical facilities such as data centers. The principles and processes included apply to all phases of new construction and renovation projects, and provide a uniform, integrated, and consistent approach to commissioning. ASHRAE *Guideline* 0 includes the total building commissioning process (TBCxP) as defined by National Institute of Building Sciences (NIBS). As a guideline, it is written as recommended best practices and may, in some instances, exceed the minimum requirements set forth in *Standard* 202.

ANSI/BICSI *STANDARD* 009-2019, DATA CENTER OPERATION AND MAINTENANCE BEST PRACTICES

This standard provides policies and procedures related to the effective operation, management, and maintenance of data centers. Specific areas include standard operations, maintenance and emergency operations, moves, adds, changes and cut-overs, security, and governance.

THE GREEN GRID, *WHITE PAPER* 79. DATA CENTER AUTOMATION WITH A DCIM SYSTEM.

This paper provides a new DCIM functional model that links automation with the physical environment, and explains how to implement it according to the business needs.

THE GREEN GRID, *WHITE PAPER* 68. THE PERFORMANCE INDICATOR: ASSESSING AND VISUALIZING DATA CENTER COOLING PERFORMANCE.

This paper presents an indicator that considers three key cooling performance metrics: PUE ratio, IT thermal conformance, and IT thermal resilience. This model enables operators to shift focus from just energy efficiency to the complete installation and operation of ITE to meet business needs.

DIN EN 50600; INFORMATION TECHNOLOGY—DATA CENTRE FACILITIES AND INFRASTRUCTURES

DIN EN 50600 represents the first European standard that uses a holistic approach to creating comprehensive specifications for the new construction and operation of a data center. It defines requirements for the construction, power supply, air conditioning and ventilation, cabling and security systems, and defines criteria for the operation of data centers.

ISO/IEC 22237 SERIES: INFORMATION TECHNOLOGY—DATA CENTRE FACILITIES AND INFRASTRUCTURES

The ISO/IEC 22237 series of standards consists of seven modules that cover general concepts, building construction, power distribution, environmental control, telecommunications cabling, security systems, management and operational information. The series defines common aspects of data centers, including terminology, parameters, and reference models addressing both the sizes and complexities of their intended purposes. The facilities and infrastructure aspects required to support data centers are also described. The series specifies a classification system, based upon the key criteria of availability, security, and energy efficiency over the planned lifetime of the data center, for the provision of effective facilities and infrastructure. The series also details the issues to be addressed in a business risk and operating cost analysis, enabling application of the classification of the data center.

EUROPEAN COMMISSION—THE EUROPEAN CODE OF CONDUCT FOR ENERGY EFFICIENCY IN DATA CENTRES

This document was created in response to the increasing energy consumption in data centers, and the need to reduce the related environmental, economic and energy supply security impacts. The aim is to inform and stimulate data center operators and owners to reduce energy consumption in a cost-effective manner without hampering the mission critical function of data centers. The *Code of Conduct* aims to achieve this by improving understanding of energy demand within the data center, raising awareness, and recommending energy efficient best practices and targets.

REFERENCES

- ASHRAE members can access *ASHRAE Journal* articles and ASHRAE research project final reports at technologyportal.ashrae.org. Articles and reports are also available for purchase by nonmembers in the online ASHRAE Bookstore at www.ashrae.org/bookstore.
- Acton, M., P. Bertoldi, and J. Booth. 2021. *Best practice guideline for the EU Code of Conduct on Data Centre Energy Efficiency*. European Commission, Ispra. JRC123653.
- AHRI. 2017. Standard for performance rating of computer and data processing room air conditioners. ANSI/AHRI *Standard* 1360 (I-P)-2017. Air-Conditioning, Heating & Refrigeration Institute, Arlington, VA.
- ASCE. 2016. Minimum design loads and associated criteria for buildings and other structures. ASCE/SEI *Standard* 7-16. American Society of Civil Engineers, Reston, VA.
- ASHRAE. 2016. Energy standard for buildings except low-rise residential buildings. ANSI/ASHRAE/IES *Standard* 90.1-2016.
- ASHRAE. 2019. Energy standard for data centers. ANSI/ASHRAE *Standard* 90.4-2019.
- ASHRAE. 2012. Method of testing for rating computer and data processing room unitary air conditioners. ANSI/ASHRAE *Standard* 127-2012.
- ASHRAE. 2021. *Thermal guidelines for data center environments*, 5th ed.
- BICSI. 2014. Data center design and implementation best practices. ANSI/BICSI *Standard* 002-2014. Building Industry Consulting Service International, Tampa, FL.
- CISCA. 2007. *Recommended Test Procedures for Access Floors*. Ceilings & Interior Systems Construction Association, Oak Brook, IL.
- Consumer Protection Agency. 2020. Evaluation of Cell-to-Cell Propagation in Lithium-Ion Batteries Containing 18650 Sized Cells. Consumer Protection Agency, Bethesda, MD. tinyurl.com/LithIonCPA.
- EIA. 2005. Cabinets, racks, panels and associated equipment. EIA/ECA *Standard* 310-E-2005. Electronics Industries Alliance through Electronic Components Industry Association, Alpharetta, GA.
- Hu, B., D. Patel, and L. Phan. 2022. Guidance for CFD modeling of data centers. ASHRAE Research Project RP-1675, *Final Report*.

- IEEE. 1995. Recommended practice for emergency and standby power systems for industrial and commercial application. *Standard* 446-1995. Institute of Electrical and Electronics Engineers, Piscataway, NJ.
- IEEE. 2008. Recommended practice for installation design and installation of vented lead-acid batteries for stationary applications. *Standard* 484-2002 (R2008). Institute of Electrical and Electronics Engineers, Piscataway, NJ.
- IEEE. 2005. Recommended practice for powering and grounding electronic equipment. *Standard* 1100-2005. Institute of Electrical and Electronics Engineers, Piscataway, NJ.
- IEEE. 2013. Recommended practice for installation design and installation of valve-regulated lead-acid batteries for stationary applications. *Standard* 1187-2013. Institute of Electrical and Electronics Engineers, Piscataway, NJ.
- NFPA. 2019. Installation of sprinkler systems. *Standard* 13. National Fire Protection Agency, Quincy, MA.
- NFPA. 2014. Inspection, testing, and maintenance of water-based fire protection systems. *Standard* 25. National Fire Protection Agency, Quincy, MA.
- NFPA. 2017. National Electric Code®. *Standard* 70. National Fire Protection Agency, Quincy, MA.
- NFPA. 2018. Handbook for electrical safety in the workplace. *Standard* 70E. National Fire Protection Agency, Quincy, MA.
- NFPA. 2019. National fire alarm and signaling code handbook. *Standard* 72. National Fire Protection Agency, Quincy, MA.
- NFPA. 2017. Fire protection of information technology equipment. *Standard* 75. National Fire Protection Agency, Quincy, MA.
- NFPA. 2016. Fire protection of telecommunication facilities. *Standard* 76. National Fire Protection Agency, Quincy, MA.
- NFPA. 2019. Water mist fire protection systems. *Standard* 750. National Fire Protection Agency, Quincy, MA.
- NFPA. 2018. Clean agent fire extinguishing systems. *Standard* 2001. National Fire Protection Agency, Quincy, MA.
- NIOSH. 1998. Occupational noise exposure—Revised criteria; Criteria for a recommended standard. National Institute for Occupational Safety and Health, Cincinnati, Ohio. www.cdc.gov/niosh/docs/98-126/pdfs/98-126.pdf.
- OSHA. Annual. Occupational noise exposure. 29 CFR 1910.95. *Code of Federal Regulations*, Occupational Safety and Health Administration, Washington, D.C. www.ecfr.gov.
- Patterson, M., D. Azevedo, C. Belandy, and J. Pouchet. 2011. The Green Grid water usage effectiveness (WUE). Green Grid Data Center Sustainability Metric *White Paper* 35.
- Pommerenke, D., and D. Swenson. 2014. The effect of humidity on static electricity induced reliability issues of ICT equipment in data center. ASHRAE Research Project RP-1499, *Final Report*.
- Rawson, A., J. Pflueger, and T. Cader. 2007. The Green Grid data center power efficiency metrics: PUE and DCiE. *White Paper* WP#06. C. Belady, ed. The Green Grid, Beaverton, OR.
- TIA. 2017. Telecommunications infrastructure standard for data centers. ANSI/TIA *Standard* 942-B-2017. Telecommunications Industry Association, Arlington, VA.
- UL. 2016. Emergency lighting and power equipment. ANSI/UL *Standard* 924. Underwriters Laboratories, Northbrook, IL.
- Zhang, J., R. Zhang, R. Schmidt, J. Gilbert, and B. Guo. 2019. Impact of gaseous contamination and high humidity on the reliable operation of information technology equipment in data centers. ASHRAE Research Project RP-1755, *Final Report*.

BIBLIOGRAPHY

- ASHRAE. 2011. Gaseous and particulate contamination guidelines for data centers. ASHRAE TC9.9, *White Paper*. tinyurl.com/mrxk9dj5.
- ASHRAE. 2014. Data center networking equipment—Issues and best practices. ASHRAE TC 9.9, 2014 *White Paper*.
- ASHRAE. 2016. BACnet™: A data communication protocol for building automation and control networks. ANSI/ASHRAE *Standard* 135-2016.