WHITEPAPER

Community First: How Data Centers Can Lead Acoustic Responsibility

Acoustics in Outdoor Applications

For years, data centers have been successfully constructed without significant consideration given to outdoor noise control. Additional thermal management has become necessary as server, rack, and power density have increased. With the additional equipment and thermal architecture evolving from computer room air conditioners to chillers and dry coolers, data center noise has become a greater area of focus.

Today's data center designer needs to understand local noise regulations, thermal management, sound sources, and sound attenuation strategies. To be a good neighbor, data center acoustics must be prioritized from day one, with designers understanding local noise regulations and engineering the site and equipment to ensure compliance.

Noise Regulations

Noise regulations, often called noise or sound ordinances, can come in a variety of forms. In the United States of America, there is no uniform noise ordinance for data centers. Instead, individual jurisdictions implement their own requirements.

Ordinances are often created by jurisdictions to ensure resident quality of life while balancing commercial and industrial interests. These ordinances are often designed to set a maximum sound-pressure limit at a specific location. The location is typically at a lot line separating property owners. A lot line defines the legal boundary of property. For example, a data center may be sited beside a residential development and several roads.

As a best practice, data center design teams should use noise ordinances to define local site requirements, including equipment selection, building and equipment siting, and required acoustical barriers or treatments. Begin by identifying and understanding all applicable local noise regulations.

Sound, Sound Power, and Sound Pressure

Sound describes the audible emissions that result from the vibration of molecules within an elastic medium. In the context of a facility or data hall, the “elastic medium” is either air or the building structure itself. (To be audible, vibration conveyed by the structure first must become airborne.) Sound is judged to be noise when it is unwanted, obtrusive, and objectionable — for example, when it interferes with speech, concentration, or sleep.

Sound occurs at a seemingly infinite range of magnitudes, or levels. The loudest sound that the human ear can hear (without damage due to prolonged exposure) is about 10 million times greater than the quietest perceptible sound. Such an extensive range makes an arithmetic scale cumbersome, so a system of logarithmic ratios — called the decibel scale — is applied instead. Other examples of logarithmic measurements include the Richter scale, which measures the energy of earthquakes, and the pH scale, which measures the concentration of hydrogen ions in liquids. This means that a 10 dB change represents a doubling (or halving) of loudness. For example, a sound pressure of 50 dB is perceived as twice as loud as 40 dB.

Sound pressure level, Lp, measures the pressure disturbance caused by sound waves as they pass through the atmosphere. Stated more simply, it describes what human ears hear and what sound meters measure. The extent of the pressure disturbance varies with the strength of the source, that is, with the amplitude of the sound waves. Sound pressure is also affected by the distance between the sound source and the listener (receiver), and by the immediate surroundings. Sound power level, Lw, describes the acoustical energy — measured in watts, W — emitted by the sound source. It is not affected by distance nor by the environment.

“Illuminating” the Difference Between Sound Power and Sound Pressure

Think of sound power as analogous to the wattage rating of a light bulb; both characteristics describe a fixed amount of energy. Sound pressure corresponds to the brightness in a particular part of the room; both attributes can be measured with meters and are affected by the immediate surroundings. In the case of light, the degree of brightness is more than a matter of bulb wattage. Brightness is also influenced by the distance between the bulb and the observer, the color of the room, the reflectiveness of the wall surface, and whether the bulb is covered with a shade. All these factors affect how much light reaches the receiver without changing the light bulb’s wattage.

Similarly, sound pressure depends on more than the sound power (acoustical energy) emitted by the source. Once again, the distance between the source and the receiver is a factor, as is whether the room is carpeted or tiled, and furnished or bare. These and other factors influence how much sound reaches the receiver, even though they do not change the sound source’s wattage (sound power).

No Noise Regulations

The following is a survey of common components found in noise ordinances. In many cases, several may be combined in a single jurisdiction's ordinance. For example, a nuisance component may be combined with A-weighted sound pressure requirements.

Many jurisdictions do not have noise ordinances. While this might seem favorable for data center operators, the opposite is often true. The lack of ordinances can lead to disputes between data center operators and communities.

In jurisdictions where there are no noise regulations, data center design teams may wish to establish a reasonable lot line sound pressure target, then evaluate the anticipated sound pressure to determine if the system noise level is acceptable.

Octave Band Sound Pressure

Human hearing covers a wide range of the sound pressure spectrum. To make this manageable, sound frequencies are typically divided into smaller ranges called octave bands. Each octave band represents a range of sound frequencies in which the highest frequency is twice the value of the lowest frequency. Each octave band is identified by its center frequency. Dividing the range from 45 to 11,200 Hertz yields eight octave bands (Table 1).

Noise ordinances may provide maximum sound pressure limits by octave band. This means that design teams would be required to evaluate the sound emissions from the site using an octave band analysis and compare them against the ordinance to determine whether they meet the requirements.

OCTAVE BAND NUMBER	CENTER FREQUENCY
1	63 Hz
2	125 Hz
3	250 Hz
4	500 Hz
5	1000 Hz
6	2000 Hz
7	4000 Hz
8	8000 Hz

Table 1. Octave band frequencies

Octave band noise ordinances are rare, and instead, weighted sound pressure levels are more common. However, when designing to meet weighted sound pressure, acoustical engineers require octave band data to accurately predict the final weighted sound pressure level.

A-Weighted Sound Pressure

Octave band sound pressure (and power) can be simplified using different weighting schemes. This takes a set of complex data — covering eight octave bands — and simplifies it to a single number. In outdoor applications, the most common single-number descriptors are A-B-C weighting. Each of these systems is based on a linear full-octave-band analysis, which can mask tones and does not provide information about fluctuating or pulsating sound. Weighting curves are used to compensate for the ear's varying sensitivity at different frequencies. A-weighting is used for low-volume (quiet) sound pressure values, while C-weighting is used for high-volume (loud) sound levels. B-weighting is very uncommon.

Most sound meters can automatically calculate and display A-weighted (and sometimes C-weighted) sound values, which makes it easy to objectively verify acoustic performance.

A-weighted sound pressure single-number descriptors are very common for outdoor sound ordinances. Typically, jurisdictions will specify a maximum sound pressure as an A-weighted sound pressure value. This allows easy verification using a handheld sound meter. A-weighted values have come to be trusted as a reasonable quantification of the human perception of sound level. But the ear's insensitivity to low frequencies does not extend to building components, which can be easily vibrated by low-frequency sounds. This can include potential window rattling and other annoyances. Thus, more comprehensive noise codes include C-weighed limits in addition to A-weighted limits.

Time response — the amount of time being sampled and averaged — may also be included. The amount of time varies from as short as one minute to as long as 24 hours.

Nighttime Adjustments

Many jurisdictions may vary their requirements based on the time of day. Daytime noise limits may exceed nighttime limits by 5-10 dB. For example, a jurisdiction may have a daytime noise limit of 60 dBA and a corresponding nighttime limit of 50 dBA, reflecting the typically quieter conditions in urban areas at night that can help support natural sleep patterns.

Zoning Adjustments

A single town or city comprises different zoning types—often residential, commercial, and industrial zoning. The jurisdiction may allow more noise in an industrial or commercial setting compared to residential. Accordingly, the jurisdiction may have unique requirements for each zoning type and areas where zoning type changes.

Nighttime Urban Noise Depiction

Ambient Background Sound Pressure Plus an Offset

Some jurisdictions use a stricter noise ordinance that accounts for background sound levels then add an offset. These "over ambient" ordinances allow operators to exceed ambient background sound levels by an offset, which typically ranges from 5 to 15 dB.

Decibel values are expressed on a logarithmic scale, so even small numerical increases can represent large physical changes. For sound intensity, a 6 dB increase corresponds to about four times the sound energy, and a 10 dB increasecorresponds to ten times the sound energy. Perceived loudness does not increase at the same rate; a 10 dB increase is often experienced as roughly twice as loud.

Pre- and Post-Installation Measuring

Jurisdictions may require ambient sound pressure measurements before and after data center installation. The ambient background levels may be established before the data center is operational to understand the baseline levels. Upon completion, additional measurements may be taken for comparison to understand how the data center has affected the background sound and to ensure compliance with the local noise ordinance. Measurements at regular intervals may be made throughout operation to account for any system or equipment changes and to ensure ongoing compliance.

Data center operators should periodically measure sound pressure at the lot lines to verify ongoing compliance, especially as site equipment changes over time.

Comparing Noise Codes in the United States

A systematic study of roughly 500 noise codes was conducted with results published in two papers.[1,2] While there is no average noise code, many trends are discussed in these papers, which can provide useful insights to a community looking to adopt or update their codes. While there are thousands of authorities having jurisdiction across the U.S., these 500 codes cover roughly one-third of the population.

Topics addressed in noise codes are presented in Figure 1. Many codes address noise in more than one way, so the total exceeds 100 percent. Just over half of the codes include a quantified noise level limit. Some codes prescribe setback distances without regard for either subjective or objective sound limits. By far, the most common noise code uses subjective language such as "nuisance" or "objectionable." For codes that include quantified limits, the range varies greatly. 55 dBA is the most common daytime noise limit, while Waco, Texas, provides the loudest limit at 85 dBa (Figure 2).

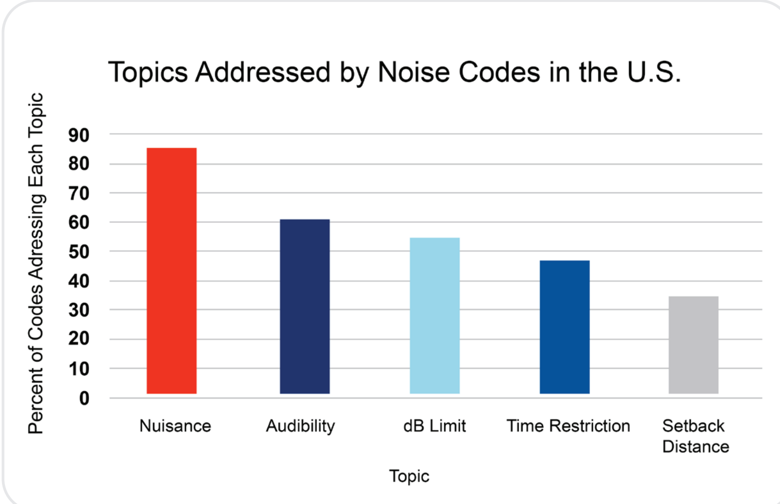

Figure 1 –
Noise codes across the United States include both subjective and objective limits.

Additional insights can be gained by learning from communities that have recently changed or are debating change to their noise codes in response to data center developments. Chandler, Arizona, updated their noise codes in 2022 to address data center noise complaints from within the community.[3] The code does not set a noise limit, rather it requires measurements of existing ambient noise levels, requires mitigation planning to not exceed the existing level, and includes validation by measurement after project completion.

Loudoun County, VA, is one of the most active data center construction markets. The county currently has a 55 dBA limit at residential lot lines. A recent whitepaper is proposing changes which would move that requirement to the data center's lot line and require pre and post-construction noise measurements[4].

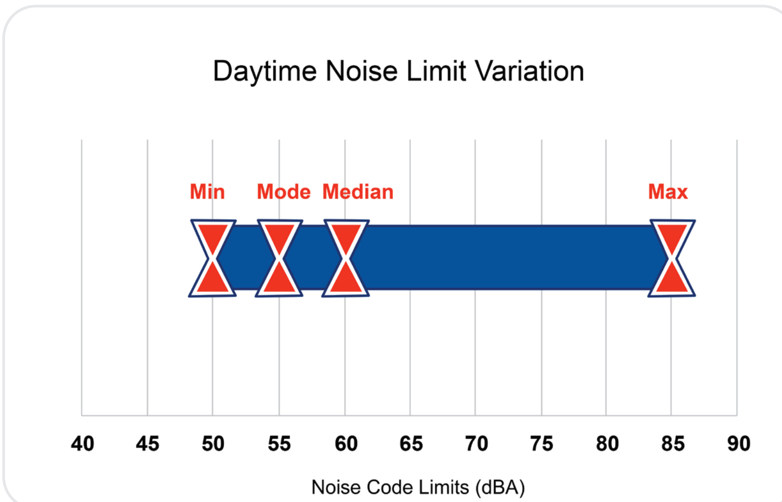

Figure 2 –
For noise codes with objective limits, 55 dBA is the most common daytime limit at residential lot lines.

1. Blomberg, L. 2016 'Preliminary Results of an Analysis of 491 Community Noise Ordinances' NOISE-CON 2016, Providence, Rhode Island, 1139-1142.
2. Blomberg, L. 2018 'Noise ordinance noise level limits, an update of the EPA's 1975 findings' INTER-NOISE 2018, Chicago, Illinois, 3588-3598.
3. https://www.chandleraz.gov/sites/default/files/departments/development-services/PLH22-0053-Ordinance-No-5033-Data-Center.pdf
4. https://www.loudoun.gov/ArchiveCenter/ViewFile/Item/13979

Data Center Site Considerations

Lot Lines

Data center site selection depends on many factors. Power availability and network connectivity are generally the most important, followed by climate and natural disaster risks. After those considerations, regulations and taxes come into play. In most cases, noise codes will be something to plan around but will not be a primary motivator in the site selection. Nearly all noise codes that include a quantified limit will be judged at the lot line of either the source (data center) or receiver (homeowner or commercial building owner). Once the site location is determined, the adjacent lot lines become the focus of concern with regards to meeting the noise code when it exists or mitigating risks to community backlash when codes are not established.

While rural areas often have less attractive infrastructure, recent power limitations are requiring many projects to bring their own power generation. This additional constraint makes rural selections more attractive, given the open space available. Lot line distances at receiving locations will often be measured in miles rather than feet. These large distances are very helpful in reducing the noise levels at the receiving lot lines. However, such large distances in rural communities introduce two complicating factors. First, the existing background noise level in a rural community can be extremely low. Even over vast distances, it can be very difficult to match such low levels. Additionally, atmospheric effects can bend sound waves towards the ground, which would otherwise be lost to the sky under certain weather conditions. One example is provided in the sidebar to illustrate such complications. Another disadvantage of very large distances is the reduced performance of sound walls, which will be discussed in more detail in the mitigations section.

Rural Data Center

What is a Lot Line?

A lot line defines the legal boundary of property. Many jurisdictions define a lot line as an infinite plane that extends perpendicularly from the ground to the sky. When noise codes define sound level limits at the lot line, the sound level applies to the entire surface area of that imaginary plane. This becomes very important in urban areas where tall buildings place people at many different elevations relative to the noise source.

The Complexities of Atmosphere on Sound Propagation

Noral Stewart was recently asked to reflect upon his 50-year career as an acoustic consultant in noise control.[5] In one notable project example, he described a noise complaint that surfaced after decades of unchanged operation at an industrial plant. Further investigation revealed that the plant had recently taken one building offline. This building didn't generate any noise, but it did generate heat and happened to be located between the noisy buildings and the affected homeowner. The heat from this building had caused much of the sound to bend upwards into the sky. Once the building was taken offline, this thermal sound barrier was lost, and the homeowner noticed a significant increase in sound at their property. While this is an unusual example, it exemplifies just one of the complications in noise control that only a qualified expert can identify and address.

5. https://noisenewsinternational.net/special-interview-noral-stewart-talks-acoustics-and-noise-control/

Data Center Site Considerations

With the rise of AI, suburban areas have been experiencing a large expansion of data centers. The lot lines between the source and receiver will be far more compressed than in rural areas, yet the distances are still typically large enough to contribute significantly to sound level reduction. Most atmospheric effects will no longer be significant, but the geography of land, especially modest elevation changes, can become very important. There will often be housing located on multiple adjacencies to be considered. Most acoustic engineers will create 3D models of the land and buildings, both on the site and in affected areas around the site, to make sound predictions across all the lot lines from potentially hundreds of individual sound sources.

Urban areas have long attracted data centers because placing them closer to end users reduces latency for both sending and receiving data. The small distance between lots provides little sound level reduction.

Additionally, the vertical nature of urban buildings adds significant complications. A recent paper[6] summarized the problem as follows:

"... data centers have limited area to put their cooling equipment and often need to maximize their rooftop layout with units; this inherently puts a limit to how many floors of data halls a data center can build.

If the design strategy of an urban data center is to be the tallest building, that strategy depends on no other neighboring property being built in the future or putting some of the responsibility of future construction on the developers of the new buildings, which can be problematic."

In all likelihood, neighboring high-rises will tower over the urban data center's rooftop, which eliminates the usefulness of sound walls due to the line of sight over the walls. There are practical limits to how tall the sound wall can be due to structural limitations.

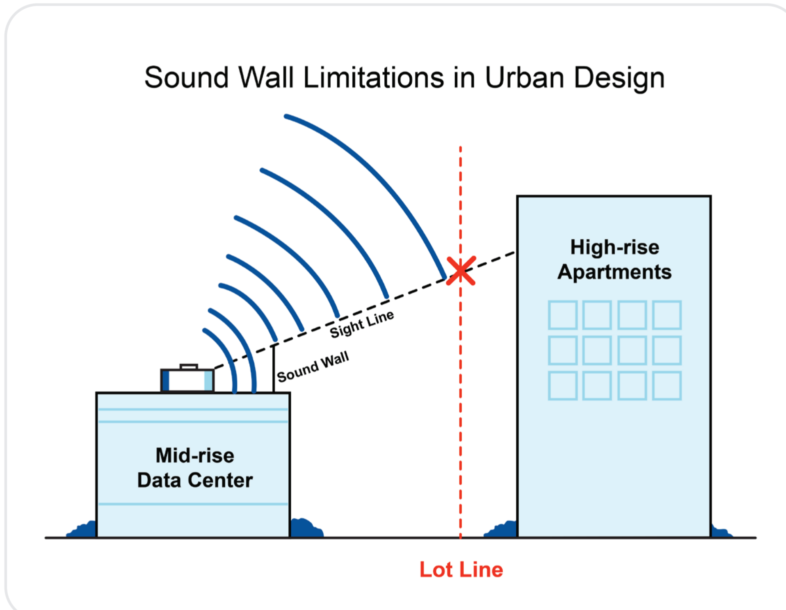

Figure 3 – Rooftop sound walls are limited in height by structural constraints, and nearby high rise buildings will see no benefit from them because their line of sight clears the barrier.

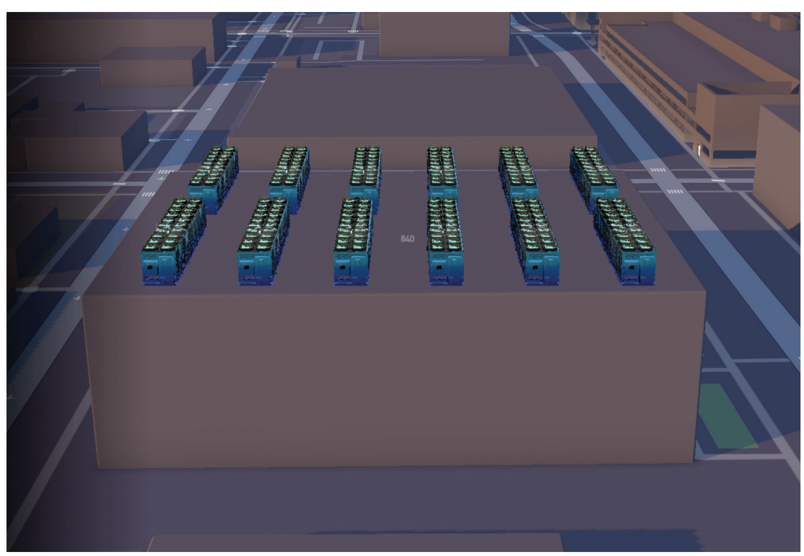

Visualization of Urban Data Center with Rooftop Equipment

6. Hamilton, S. 2024 'Challenges in Urban Data Center Design' NOISE-CON 2024, New Orleans, Louisiana, 716-723.

Campus Design

A data center is often made up of multiple facilities arranged across a larger campus. Earlier, the importance of distance in relation to sound levels between the source and receiver was discussed. There are opportunities to optimize the campus design for sound considerations after the site has been selected and the lot line noise requirements have been clearly defined. Rural, suburban and urban sites will have their own unique challenges.

Recall that in rural areas, there are typically large distances between the data center and the effected neighbors. In these cases, the potential to use campus layout to improve sound is not likely to exist. Placing a chiller yard at the furthest corner of the data center campus relative to a particular neighbor won't proportionally impact the sound travel distance by a significant amount if that neighboring lot is a mile away.

Suburban areas offer the best opportunity to leverage campus design. Placing equipment emitting higher noise levels farther away from the nearest residential lot line can be very helpful. For equipment mounted on grade (like a chiller yard), arranging the data hall between the sound source and residential lot line can offer the benefits of a very large sound wall. However, sites tend to be quite complicated with multiple residential lot lines in multiple directions. Additionally, the sound sources are not a single object but a huge array of varying equipment types. It is therefore essential to engage a qualified professional to map sound propagation from all sources to all receivers simultaneously, helping to balance these factors and determine the optimal campus layout. A recently published paper does an excellent job of describing this iterative process. In this example, a developer was designing three adjacent data centers as separate projects, but the acoustic consultant found it was essential to adjust the campus layout of all three as a single campus to optimize the layout and minimize the cost of noise mitigation.[7]

In most urban data centers, there isn't really a campus. They tend to be a single facility with nearly all mechanical equipment placed on the rooftop. In some cases, there might be some benefit to considering which adjacent lot is the most sensitive to noise and arranging the rooftop layout accordingly. However, the quantity of equipment typically fills the entire rooftop, so reshuffling equipment will not provide much benefit. It is worth noting again that in an urban area, the vertical dimension is important to consider. If noise-sensitive neighbors tower over the data center rooftop, the specific layout of equipment will have little impact on neighbors at higher elevations. Efforts to select quieter equipment with enclosures will yield more useful results.

Facility Design

The decisions made regarding site location and campus design sometimes have more freedom from constraints than the design of the facilities. Many data centers have historically relied on repeatable, standardized designs to accelerate construction and reduce overall build time. Increasingly, however, developers are adopting more customized, site-specific designs to address unique operational needs, local requirements, and project conditions. There are a few options available to reduce noise at neighboring lot lines through conscientious facility design because a structure can serve the same function as a sound wall. Most commonly, this applies to on-grade noise sources being tucked up next to the facility on the less-sound-sensitive side. For rooftop equipment, some data centers will have two or more rooftop elevations. This enables the higher portion of the building to again serve as a sound barrier for equipment mounted on the lower elevation roof. Additionally, parapets can be added to the edge of the roof to block sound by serving as a small sound barrier.

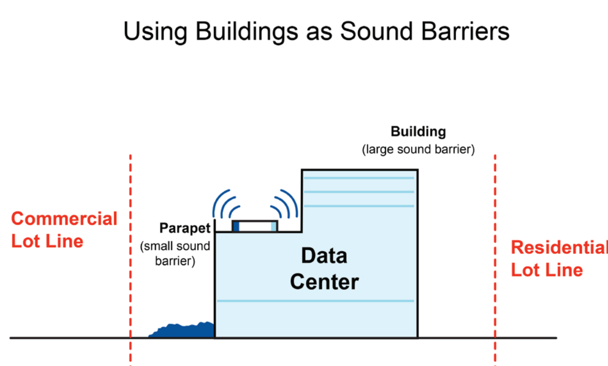

Figure 4 – The data center can sometimes be used as a sound barrier if the loudest equipment can be placed on the least noise-sensitive side.

7. Conaway, M., Baldassano, J. 2024 'It's only one more data center! When scope creep becomes sound emissions creep' NOISE-CON 2024, New Orleans, Louisiana, 206-215.

Sound Sources

There are a variety of different systems and types of equipment that can produce sound within a data center site.

Outdoor Sounds Sources Include:

- Thermal Management systems components, such as chillers and heat rejection equipment
- Electrical generators and transmission systems
- Ancillary systems

Indoor Sounds Sources Include:

- IT equipment
- Thermal Management systems such as air handlers, chillers, coolant distribution units, and pumps and support equipment

It is important to understand all the various sound sources if acoustical modeling and predictions are needed. Thermal management equipment operation can vary from site-to-site and by geography. In some cases, chillers might operate 24 hours per day, or sporadically, while dry coolers are modulated and cycled, creating a dynamic sound environment. Understanding this will be necessary when demonstrating compliance with local noise ordinances prior to construction.

Chiller Sound

Chillers are typically divided into two categories — air-cooled or water-cooled. The former includes a full refrigeration system with heat rejection. This packaged unit allows the chiller to produce chilled water for equipment cooling and rejects heat to the surrounding ambient air. Water-cooled chillers require separate heat rejection equipment, such as cooling towers or dry coolers. Industry standards are used to ensure uniform sound testing and data presentation, which leads to better comparisons between equipment models and manufacturers.

Air-Cooled Chillers

Air-cooled chillers sound sources include the refrigerant compressor(s), refrigeration system, and condenser. There are a variety of compressor types available in the marketplace, including scroll, helical rotary (also called screw), and centrifugal. These compressors each have unique sound characteristics. The condenser rejects heat by circulating ambient air through a coil with hot refrigerant flowing through it. Manufacturers typically design and optimize the air-cooled chiller to provide the greatest amount of heat transfer in a small footprint while producing the least amount of noise possible.

North American air-cooled chillers are typically tested in accordance with AHRI Standard 370. This standard allows manufacturers to test their equipment either within a controlled laboratory environment or outdoors in a free field environment. Due to the size of the chillers, it's common to perform outdoor testing. The result of the test is octave band sound power ranging from 63 to 8000 Hertz, which is useful for sound prediction. The standard requires manufacturers to perform the sound testing at a set of standardized AHRI thermal rating conditions, which allows easier product comparisons.

For data centers, air-cooled chillers are often configured to produce chilled water at temperatures warmer than those used for typical comfort cooling applications. As a result, the compressor, refrigeration system, and condenser operate differently, which results in different sound power levels. Manufacturer selection software or catalogs may allow chiller operating conditions to be specified during product selection. Specifying the expected operating conditions during design conversations helps ensure that the sound data accurately represents real-world operation. Partnering with manufacturers that use advanced laboratory testing and predictive acoustical modeling is essential to verify that any proposed mitigation approach will perform as required under the unique data center thermal and environmental conditions.

Sound Sources *(cont.)*

Water-Cooled Chillers

Water-cooled chillers are often selected for their energy efficiency over air-cooled models. Typically, water-cooled chillers and their pumps are installed in a mechanical equipment room. These rooms are commonly constructed of masonry, either cement block or poured concrete. Neither of these materials absorbs or transmits a significant portion of the sound energy, so much of it reflects into the room. This reflected sound adds to the sound emanating from the sources, greatly increasing the overall sound level in the mechanical equipment room. Mechanical equipment room surfaces may be treated with absorptive material to reduce the reflected sound. A series of cooling towers, or, more commonly in modern data center design, dry coolers, is used to reject heat to the ambient environment.

Historically, manufacturers have published water-cooled chiller sound pressure obtained in accordance with AHRI Standard 575. AHRI Standard 1280 is a newer option to obtain water-cooled chiller sound power.

Heat Rejection Equipment: Cooling Towers and Dry Coolers

Heat rejection equipment, typically in the form of cooling towers or dry coolers, is installed outside of the data center facility to reject heat to the ambient environment. This makes the heat rejection system a significant outdoor sound source. There are several different sound standards used for heat rejection equipment, including CTI Code ATC-128, CEN EN13487, and others. It is critical to ensure the sound data type — power or pressure — is the same when making comparisons. Ensure the measurement distance — often 30 feet or 10 meters — is consistent between comparisons.

Cooling Tower

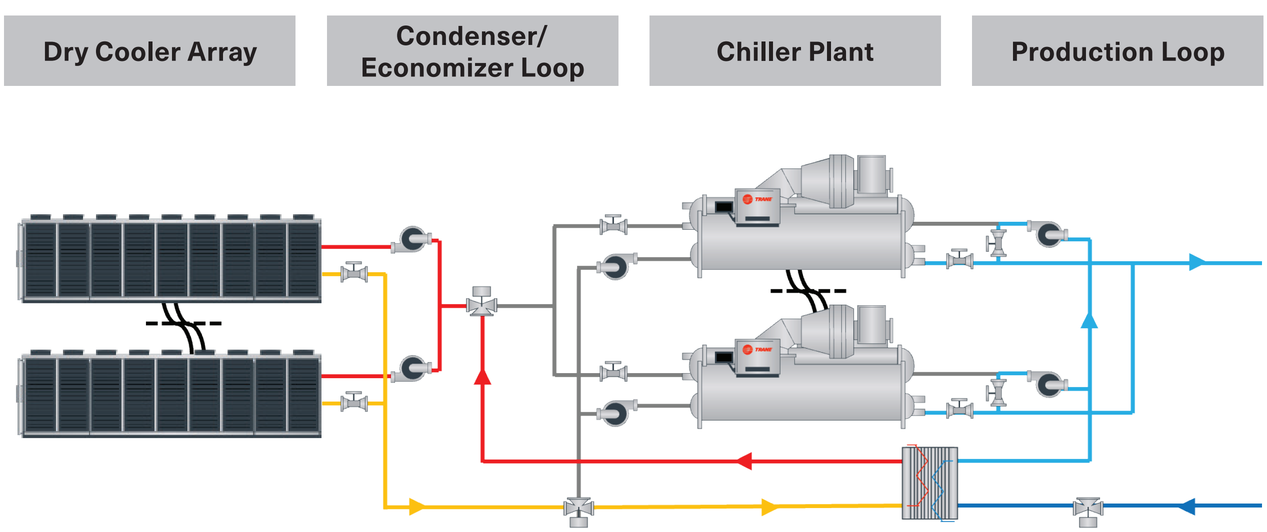

Reference Diagram | Trane 1GW AI Factory Reference Design

Sound Attenuation and Mitigation Strategies

Distance Attenuation

Attenuation due to distance is at the foundation of every outdoor noise problem. Larger distances are helpful in reducing noise, but what is less obvious is how the distance factor becomes less effective over large distances. To illustrate this, consider an idealized example of a point source over a reflecting plane. Figure 5 illustrates how sound propagates hemispherically away from the source. The sound pressure level will decrease as the sound energy is spread over larger hemispheres. If no obstacles or other sources of attenuation exist, the sound pressure level will decay at a rate of 6 dB with each doubling of distance. Figure 6 graphs this relationship and adds some real-world distances for context.

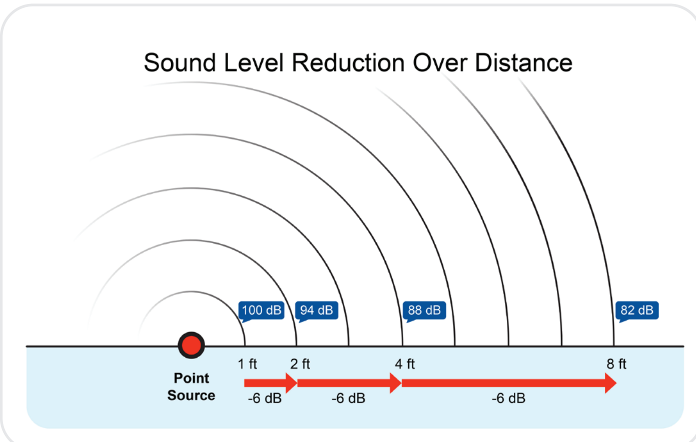

Figure 5 – An idealized point source over a reflecting plane will experience a reduction in sound pressure level of 6 dB per doubling of distance.

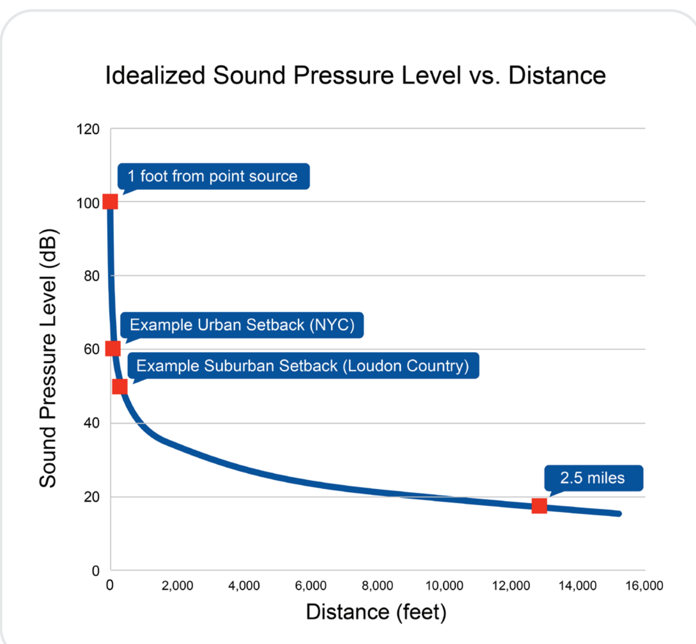

Figure 6 - Sample calculations of sound pressure level at various distances: (Urban) New York City requires a minimum structure-to-structure distance of 90 feet for buildings taller than 125 feet. (Suburban) Loudon County has proposed a 300-foot minimum setback for new data centers. Several community noise codes require 2.5-mile setbacks for data centers relative to schools and hospitals. Note, this is purely hypothetical, based on a single sound source under the most idealized sound propagation assumptions.

Environmental Officer Setting Sound Level Meter for Monitoring

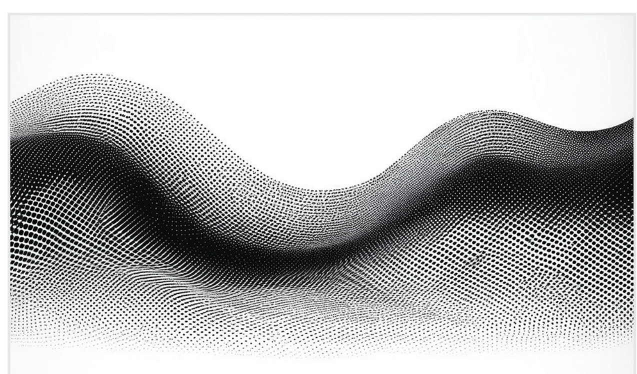

Sound Visualization

Sound Attenuation and Mitigation Strategies *(cont.)*

Barrier Walls

A barrier reduces noise by blocking the line-of-sight between a sound source and receiver. However, the wave nature of sound propagation makes this benefit complicated. The size of a sound wave varies by frequency. Over the range of human hearing, these wavelengths are as long as 50 feet and as short as 0.5 inches. For a barrier to block noise, the size of the barrier needs to be a significant fraction of the wavelength. This becomes difficult to achieve when low frequencies are involved because they require very large sound barriers. In effect, barriers create shadow zones where significant noise reduction is achieved. The size of the shadow zone depends on the wavelength of sound, the height of the barrier, the distance from the source to the barrier, and the distance from the receiver to the barrier. These calculations are standardized in the field of acoustic engineering and best managed in 3D software, which can calculate barrier effectiveness from an arbitrary number of sound sources across a landscape of receivers. Sound barriers can be highly effective over short and medium distances, but the details are quite complicated and should be evaluated by an expert alongside all the other mitigation options. In other words, a sound barrier is not always the right solution.

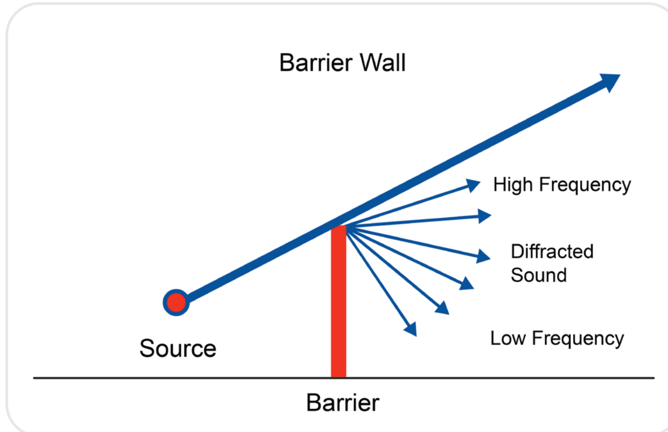

Figure 7 – Sound barriers can be very effective at blocking noise, but their performance is reduced at low frequencies and the receiver must be in the shadow zone of the barrier. For additional information, see the Trane Engineers Newsletter Live entitled "Acoustics in Outdoor Applications."[8]

Sound Enclosures

For a sound barrier to be useful, it must cast a sound shadow over the receiving area. This can become impractical over long distances because the height of the required barrier becomes unrealistic. At this point, a sound enclosure becomes a more viable solution. The concept of enclosing a noise source is obvious, but the details are complicated. All noise sources around a data center require substantial airflow to function. This means that an enclosure must have substantial open areas through which air can pass. Fundamentally, this airflow requirement renders sound isolation impossible. Sound is sometimes compared to water in that it will find any leaky path available (noise control textbooks call these flanking paths). There are two ways around this airflow limitation. One is to build a small building around the equipment with air inlets and outlets designed into the building, which forces the sound to reflect around corners. The cost of such a solution can rival the equipment cost, and the airflow is usually significantly impacted, which could derate the performance of the equipment. The other option is to enclose the unit with large acoustic louvers. Relative to the first approach, this can be more cost-effective due to lighter construction. The maximum effectiveness is reduced, but when engineered properly, the size of the open area in the louvers will not derate the equipment housed inside.

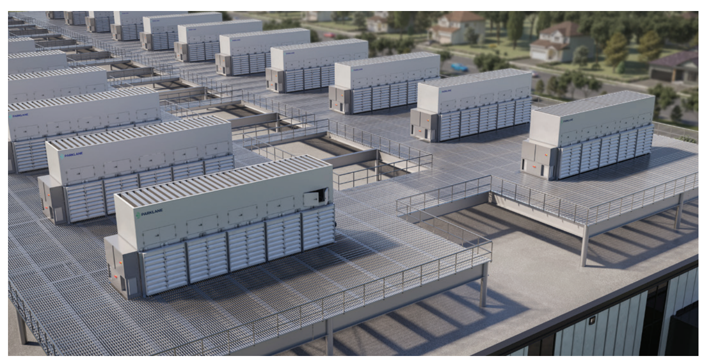

Figure 8 - An example of at-source discharge silencer(s) with hinged louver intake silencer(s) solutions via Parklane

Because of the complexity, the enclosure and equipment manufacturers should work together on a noise-reduction target set by an acoustic engineer, who can weigh enclosure cost and performance against other mitigation options.

8.Trane. 2016 "Acoustics in Outdoor Applications" Engineers Newsletter Live. La Crosse, Wisconsin, APP-CMC060-EN.

Sound Attenuation and Mitigation Strategies *(cont.)*

Chilled Water System Attenuation

System designers have multiple options available to attenuate a chilled water system. Air-cooled chillers have several sound sources, including compressors, refrigeration system components, and condenser fans. Manufacturers design and optimize the components for reduced sound generation. Condenser fans can operate at a reduced speed to produce less noise, but this generally means the chiller will function with reduced capacity or an elevated leaving chilled water temperature.

Examples of Other Air-Cooled Chiller Attenuation Options:

- **Refrigeration system treatments** include compressor, suction and discharge piping, and system accessory wraps that are designed to reduce noise from the chiller refrigeration system.
- **Discharge silencers** can be chiller-mounted or ground-supported sound baffles, which are purpose-built for the chiller model. This category of attenuation is typically installed atop the chiller and designed to address airflow and condenser fan noise.
- **Intake louvers/silencers** are designed to address the air intake noise generated by the airflow and condenser fans. Models may be attached to the chiller along the perimeter. Systems with these options can increase chiller footprint by several feet in all directions.
- **Sound barrier walls** can be used to reduce sound transmission to the lot line. See the corresponding section for further discussion.
- **Sound enclosures** are custom designed and custom built structures engineered to fully encapsulate chillers, delivering maximum sound attenuation.

Water-cooled chilled water systems require a cooling tower or more commonly in modern data centers, dry cooler, for heat rejection. These devices include large fans to circulate ambient air through the heat transfer surface to reject heat to the ambient environment.

Fans can be engineered to operate more quietly by optimizing their design and running them at reduced speeds to minimize noise output. Systems should be designed and evaluated with any night setback or condenser fan limitations to ensure sufficient thermal performance.

Another way to reduce sound is to integrate evaporative free cooling — often called “adiabatic assist”— into the chiller or dry cooler. This approach decreases the heat rejection load on the condenser or dry cooler, allowing the heat rejection fans to operate at lower speeds, which in turn reduces overall sound levels. Similarly, an air cooled condenser can incorporate evaporative cooling.

By spraying water into the condenser’s incoming airstream, the dry bulb temperature is lowered, enhancing heat rejection and allowing the condenser fan to operate at reduced speeds — ultimately decreasing overall noise production.

When designing a cooling system for a data center, there will be many competing requirements. Manufacturers bidding on the project must demonstrate that their equipment meets the specified performance criteria, including, in many cases, limits on the sound power each unit may emit at the design point.

The design point itself (addressed later in this section) is typically expressed as the specific outdoor air conditions in which the equipment must deliver enough cooling capacity to keep the indoor facility within optimal operating temperatures.

6N+1 Redundancy

105°F: *50-year extreme design day*

Redundant Chiller Not in Use

	FANS	dBA
Chiller 1	1000	99
Chiller 2	1000	99
Chiller 3	1000	99
Chiller 4	1000	99
Chiller 5	1000	99
Chiller 6	1000	99
Chiller 7	0	0
TOTAL FAN SOUND		**107**

105°F: *50-year extreme design day*

Utilizing Redundant Chiller

	FANS	dBA
Chiller 1	870	96
Chiller 2	870	96
Chiller 3	870	96
Chiller 4	870	96
Chiller 5	870	96
Chiller 6	870	96
Chiller 7	870	96
TOTAL FAN SOUND		**104**

96°F: *Annual extreme design day*

Utilizing Redundant Chiller

	FANS	dBA
Chiller 1	726	92
Chiller 2	726	92
Chiller 3	726	92
Chiller 4	726	92
Chiller 5	726	92
Chiller 6	726	92
Chiller 7	726	92
TOTAL FAN SOUND		**100**

Figure 9 – Hypothetical example of leveraging additional equipment to deliver lower sound.

Consider a design strategy where the cooling and compute loads are designed together in blocks. One common approach is to design for 6N+1 redundancy. This means that seven chillers will be dedicated to the block of compute even though the cooling load can be served by six chillers. The extra chiller is included in the block to reduce risk by always having one chiller on standby. Additionally, it is common for these six chillers to be selected to provide this cooling load on the hottest day in several decades. This extra chiller and designing for a worst-case scenario are necessary to ensure the data center can remain operational under unlikely and extreme conditions. However, this redundant chiller and extreme operating capability offer two opportunities to potentially reduce the sound levels.

An analysis was performed on a hypothetical 6N+1 block of air-cooled chillers operating in the Upper Midwest. Designing for the hottest day expected in 50 years, six chillers could accommodate the cooling load at a total sound power level of 107 dBA. On that one day in 50 years, it would be possible to run all seven chillers to serve the same load at a reduced level of approximately 104 dBA. The key is to run the condenser fans at 6/7th speed. An additional 4 dBA of sound reduction was predicted with all seven chillers serving the full load on the hottest day in an average year.

There are several important considerations for this analysis. The size of the benefit may not be immediately obvious but recall that doubling the distance to an outdoor sound source is roughly equivalent to a 6 dBA reduction. By using the chillers already in the 6N+1 design, a reduction of sound was possible, which would equate to doubling the lot line distance.

It is common in the building industry to select equipment which can serve the required load on the hottest day of a year, and sound mitigation plans are made around this design day. Data centers demand much higher resilience and often select equipment capable of full load on the hottest day in 50 years, but there is no logical reason why the sound mitigation plan should be based on this ultra-rare design day. When manufacturers investigate how their equipment will operate at lower ambient temperatures while still providing the full cooling load, this is called ambient relief. The analysis above shows that the difference between the hottest day in 50 years and hottest day in an average year provides significant ambient relief which can allow the condenser fans to spin slower, generating less sound.

It is important to note that the compressor in this analysis had sufficient headroom at these outdoor air conditions to accommodate the reduced condenser fan speeds, while providing the full cooling load without producing excessive sound. This may not always be true, especially in more demanding locations (when the design day temperature is even hotter).

When manufacturers provide sound data at all operating conditions and use control strategies that reduce fan speeds even on the hottest typical day of the year, system sound levels can be reliably derated. This approach can significantly reduce the need for costly mitigation measures — such as barriers, walls, enclosures, or campus redesigns — while still maintaining acceptable sound levels for neighboring properties on every day in an average year.

Not all data center projects will allow for separate design conditions to be implemented for sound. Fortunately, there are other tools available. Adiabatic assist adds on demand water mist to provide evaporative pre-cooling of the ambient air before it reaches the condenser coils.

From a sound perspective, the reduced air temperature can translate to lower fan speeds. In addition to being useful at the design point, it also gives operators the ability to employ peak shaving during the hottest part of each day, so it has usefulness well beyond enabling lower sound levels on the rare design-day condition.

There can also be significant sound benefits to using water-cooled chillers, which decouple the heat rejection solution from the refrigeration solution. Water-cooled chillers are easier to enclose or locate indoors, allowing the heat rejection equipment to take priority in the control strategy for minimizing sound without compromising chiller performance.

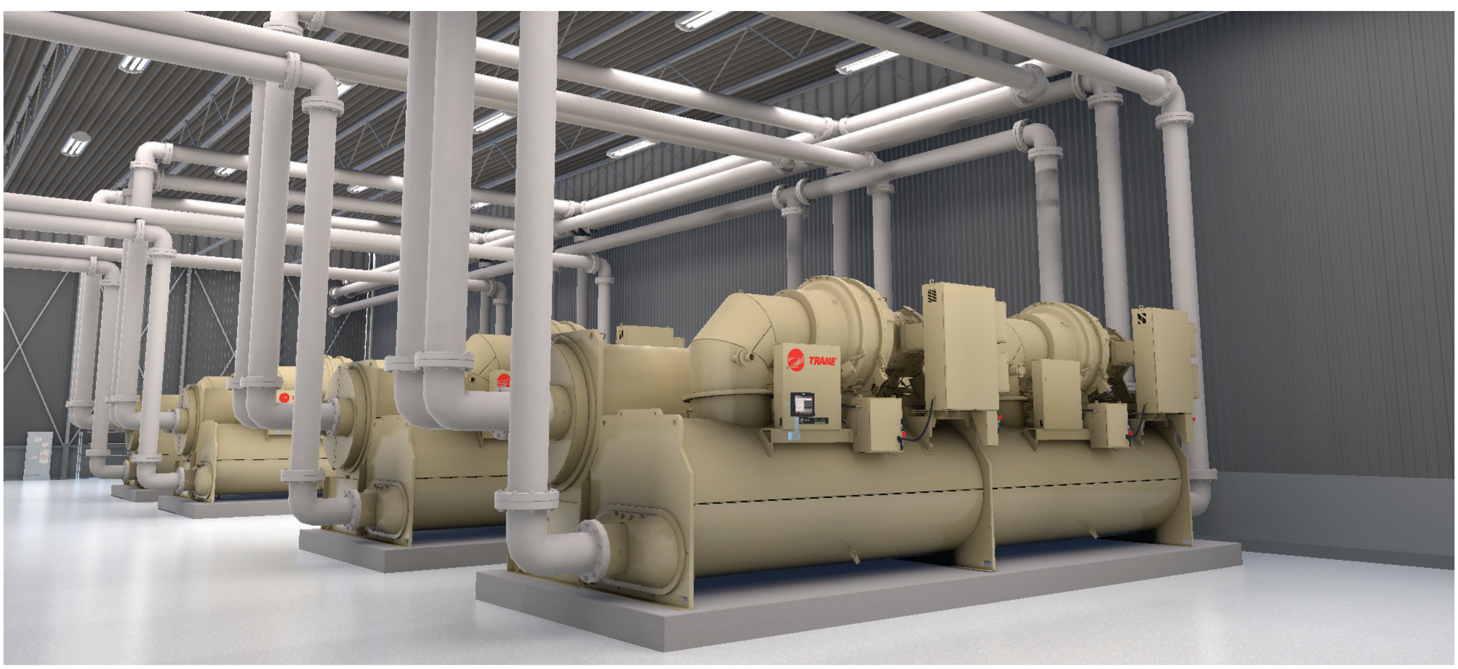

CenTraVac® Water-Cooled Data Center Chillers Model CDHH

Conclusion

As hyperscale and colocation data centers continue to expand in scale, density, and urban geographic reach to enable AI, their acoustic impact has emerged as a critical dimension of responsible development. No longer a peripheral consideration, sound must be treated as a core design parameter — one that influences site selection, campus layout, equipment strategy, and long-term operational planning.

Across diverse jurisdictions, noise regulations and community expectations vary widely, requiring data center designers to navigate an increasingly complex landscape of ordinance types, measurement methods, and compliance challenges. Whether regulations are highly prescriptive or entirely absent, thoughtful acoustic design is essential to maintaining positive relationships with surrounding communities.

Effective acoustic performance cannot rely on a single tactic. Rather, it requires a holistic and collaborative approach — one that integrates acoustical engineering, equipment selection, mechanical design, and site planning from the earliest stages of development. Understanding sound sources, modeling their propagation, and evaluating mitigation strategies such as barriers, enclosures, operational controls, and advanced thermal management technologies allows designers to address noise concerns proactively and efficiently. These solutions must balance acoustic goals with thermal performance, energy use, reliability, and cost, recognizing that data centers face competing priorities that must be reconciled through informed engineering decisions.

By embracing acoustic responsibility and partnering early with stakeholders and collaborative thermal management system innovators like Trane, data center developers can build facilities that coexist harmoniously with their surroundings. This means working closely with manufacturers, acoustic consultants, and communities to ensure that sound is managed effectively at every stage — from concept to commissioning and through ongoing operation.

Ultimately, prioritizing acoustics not only supports regulatory compliance but also builds trust, enhances community well-being, and strengthens the long-term sustainability of data center development. In a world increasingly dependent on digital infrastructure, designing quiet, considerate, and community-aligned data centers is both an engineering and social imperative.

About the Authors

Derrick Knight, Lead Acoustic Engineer

Derrick joined Trane as an acoustic engineer in 2014 after more than a decade as an acoustic consultant and lab test engineer, building on degrees in Physics and Architectural Acoustics. At Trane, he has led sound rating and lab improvement projects, contributed to quieter product designs, and serves as a subject matter expert in acoustics across the organization.

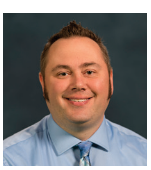

Eric Sturm, Lead Applications Engineer

Eric joined Trane in 2006 after earning his mechanical engineering degree from the University of Wisconsin–Platteville and later served as a marketing engineer and product manager in Customer Direct Services, supporting and training customers worldwide on TRACE™ 700. In his current role as an applications engineer, he specializes in acoustics, airside systems, indoor agriculture, and indoor air quality.

Disclaimer: *This is for informational purposes only. Trane believes the facts and suggestions presented here to be accurate. However, final design and application decisions are your responsibility and will affect actual financial and energy efficiency results. Trane disclaims any responsibility for actions taken on the material presented.*

Trane – by Trane Technologies (NYSE: TT), a global climate innovator – creates comfortable, energy efficient indoor environments through a broad portfolio of heating, ventilating and air conditioning systems and controls, services, parts and supply. For more information, please visit *trane.com or tranetechnologies.com.*