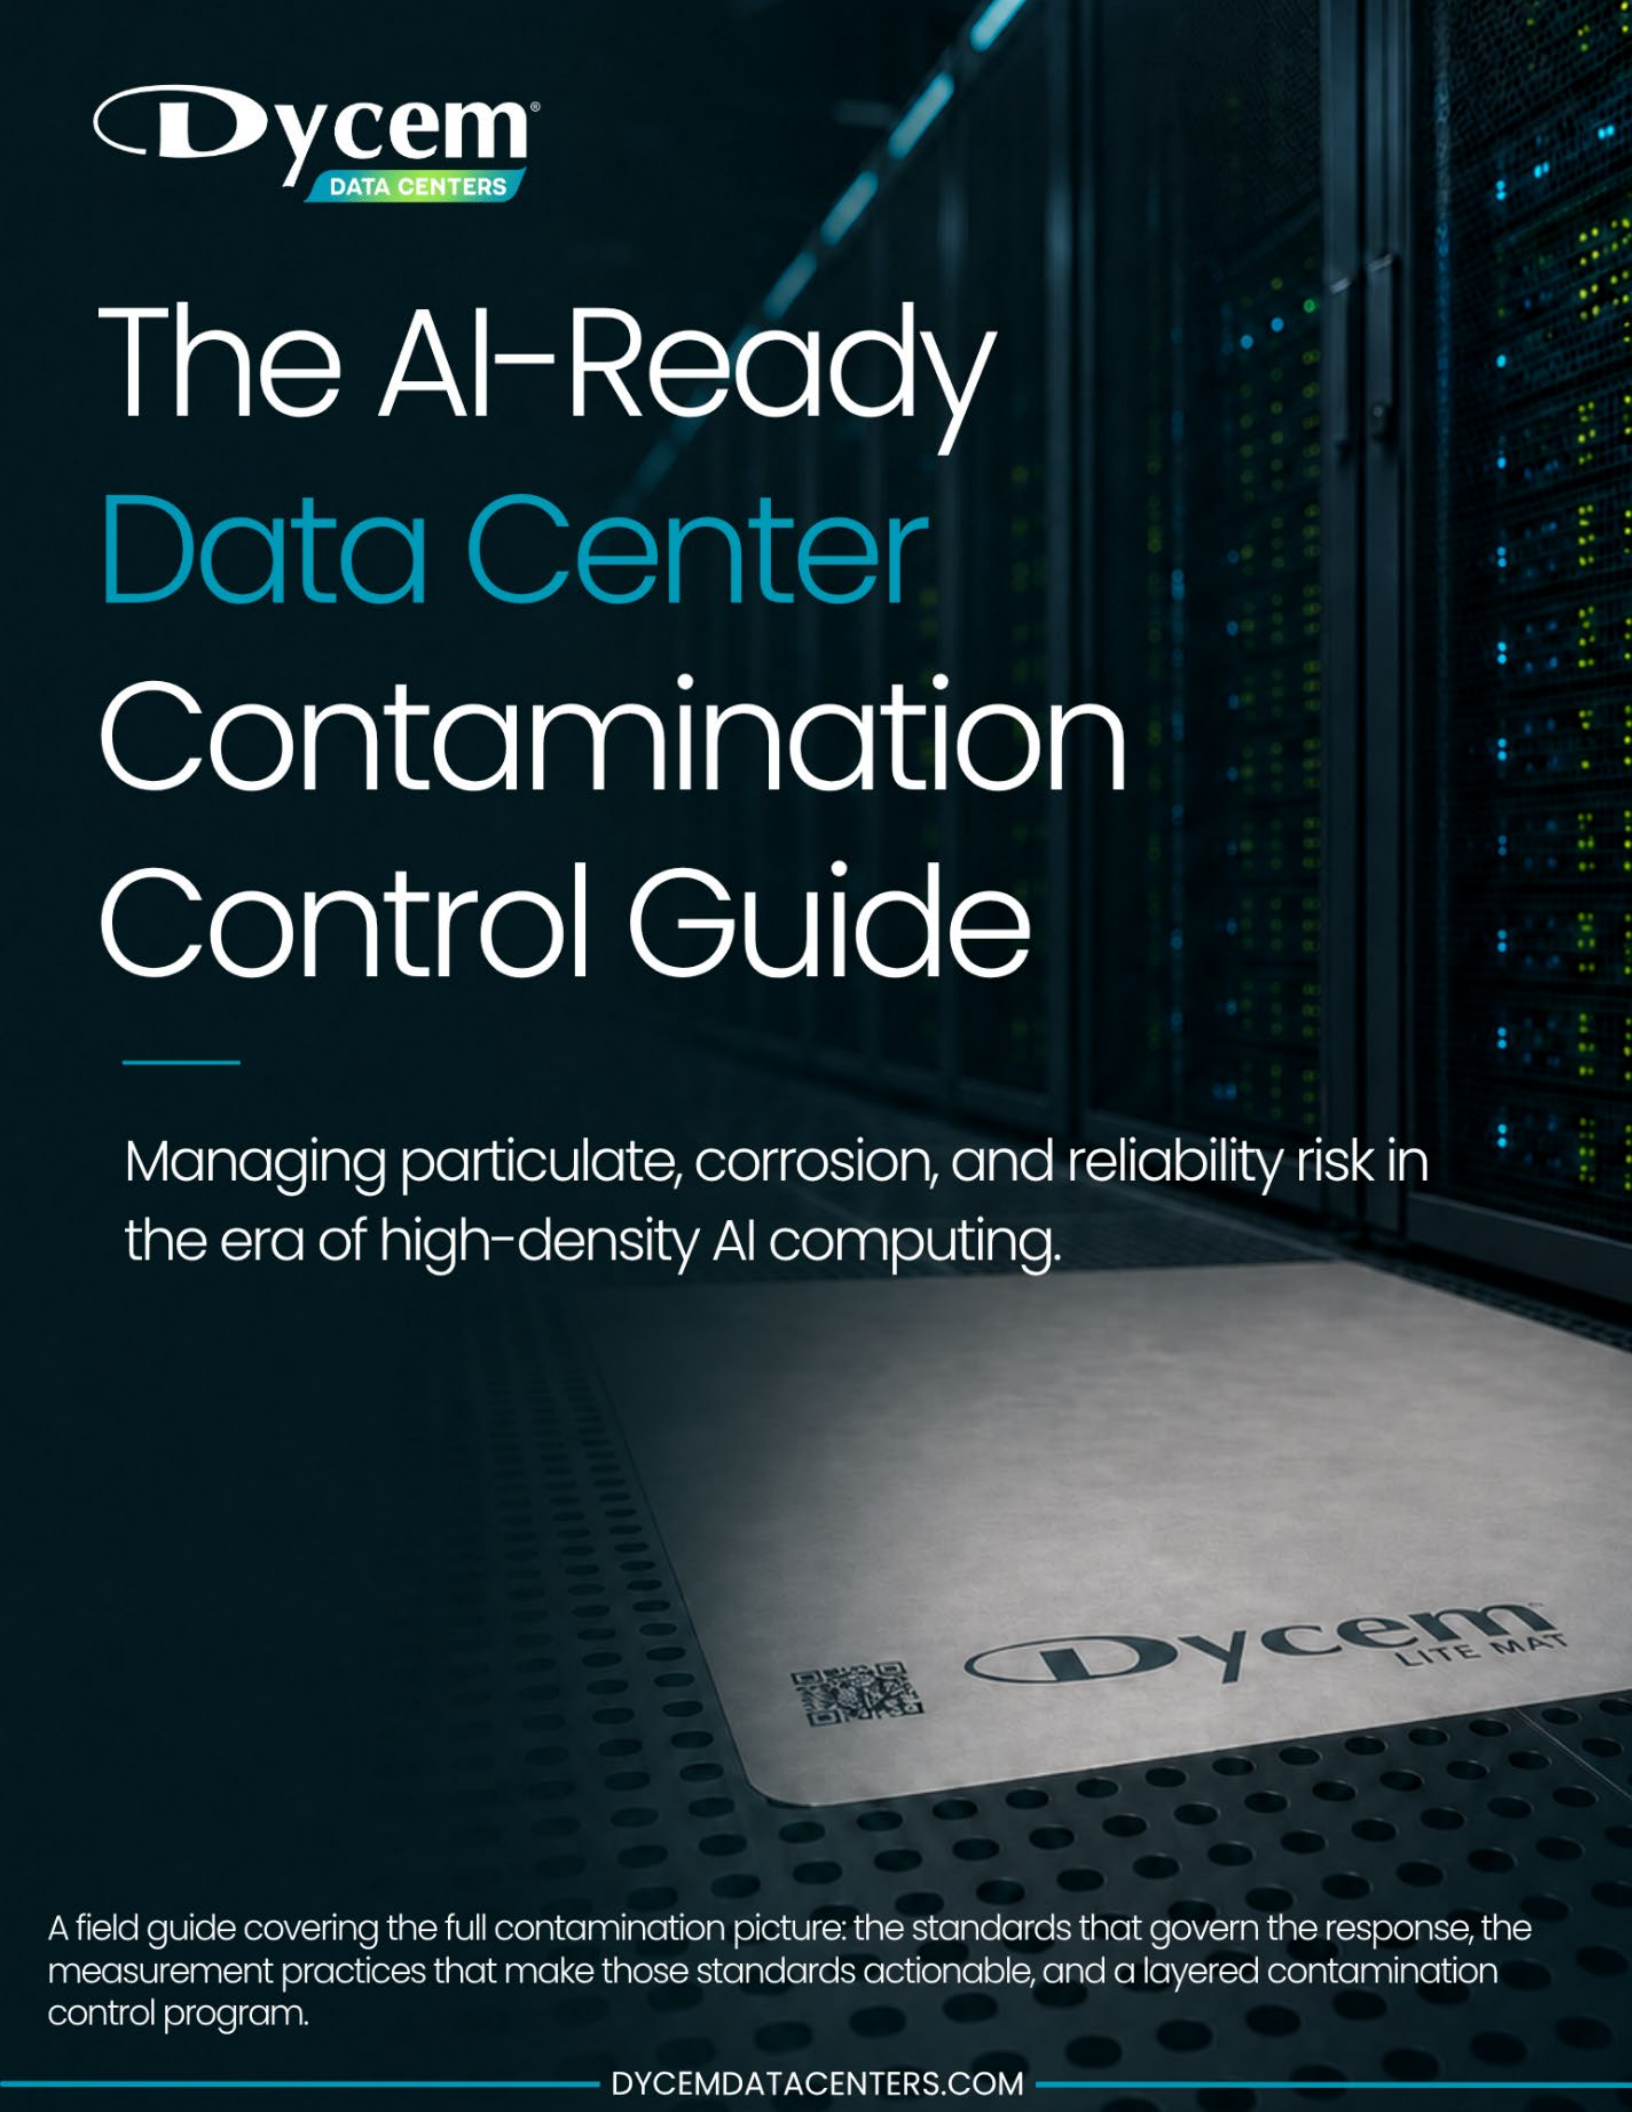
Dycem DATA CENTERS

The AI-Ready Data Center Contamination Control Guide

Managing particulate, corrosion, and reliability risk in the era of high-density AI computing.



A field guide covering the full contamination picture: the standards that govern the response, the measurement practices that make those standards actionable, and a layered contamination control program.

Before you dive in

WHY READ THIS

Speak the standards, not just the sales pitch. Walk into your next design review able to reference ISO 14644-1 Class 8, ASHRAE TC 9.9, and ISA-71.04 G1 by name and know where to go for your contamination control program.

See the cost you're not currently measuring. Understand why contamination is mostly a quiet efficiency and lifespan tax so you can make the case internally with the argument that holds up under scrutiny.

Know where every control fits, and where it stops. Get a clear map of air-side filtration, entry-zone capture, containment, and coolant-loop protection including exactly what floor-level control does and does not do.

Get ahead of the liquid-cooling shift. Most facilities teams haven't yet had to think about coolant-side contamination. Section 12 gets you oriented before you have to manage it under pressure.

Leave with tools, not just concepts. The appendices contain a facility self-assessment scorecard and glossary, so the guide converts directly into program work.

WHAT YOU SHOULD READ

If you are…	Sections to Read
An executive or facility owner building the internal case	Executive summary (1), Why it matters now (2), The economics (3), Business case
A facilities or operations lead running the program	Sections 2–3, 9–11, 13–14, and the roadmap
A design engineer or technical evaluator	The full guide — Parts II, III, and IV are written for you
A procurement or sustainability lead	Sections 3, 15–17

Contents

Before you dive in 2

WHY READ THIS 2

WHAT YOU SHOULD READ 2

PART I 7

1 Executive summary 7

2. Why contamination control matters more in the AI era 9

2.1 The density shock 9

2.2 The cooling transition 10

2.3 Chronic cost versus acute risk 10

3. The economics of dirt: Where the cost shows up 11

3.1 The cooling-energy tax 11

3.2 The performance tax 11

3.3 The hardware-life tax 11

3.4 The warranty and compliance exposure 11

3.5 The acute tail: outage risk 11

PART II 13

The Science 13

4 The Types of Data Center Contamination 13

4.1 Particulate 13

4.2 Gaseous and corrosive contamination 14

4.3 Zinc whiskers and metallic filaments 14

4.4 Electrostatic discharge (ESD) 14

4.5 Moisture and deliquescent humidity 15

5 Mechanisms of harm 15

6 How contamination moves: the cycle and the reservoirs 17

6.1 The entry vectors 17

6.2 Resuspension: Why settled is not gone 17

6.3 The hidden reservoirs 17

6.4 The self-sustaining loop 18

PART III 19

The Standards 19
7. The standards and frameworks landscape 19
7.1 ISO 14644-1 — air cleanliness classification 19
7.2 ASHRAE TC 9.9 — thermal and contamination guidance 20
7.3 ANSI/ISA-71.04 — corrosion severity 20
7.4 ANSI/ESD S20.20 — electrostatic discharge control 20
7.5 EN 50600 and the DCA Anti-Contamination Guide — the operational frameworks 20
7.6 GHG Protocol and ISO 14044 — where sustainability and contamination control meet 21
8. Measurement, monitoring, and documentation 22
8.1 The measurement toolkit 22
8.2 Why the records matter 22
PART IV 24
The Contamination Control Program 24
9. Defense in depth: the four-layer model 24
10. The zoned program 27
ZONE 1 · ENTRANCE & RECEPTION 27
ZONE 2 · TRANSITION / BUFFER 27
ZONE 3 · WHITE SPACE 28
WHOLE FACILITY · ONGOING 28
11. Floor-level control: mechanism, fit, and scope 28
11.1 The foot and wheel traffic vector 28
11.2 Capture versus retention 29
11.3 What floor-level contamination control does not do 29
12. Liquid cooling: the new contamination frontier 30
12.1 Why the fluid side is unforgiving 30
12.2 The four fluid-side risk categories 30
12.3 Monitoring the loop 31
12.4 Operational discipline 31
13. Construction, commissioning, and retrofit 32
13.1 New builds: cleanliness as a commissioning deliverable 32

13.2 Liquid-loop commissioning 32
13.3 Retrofit in live facilities: the hard case 32
14. Cleaning operations and contractor qualification 34
14.1 Scope of a proper cleaning program 34
14.2 Qualifying a data-center cleaning contractor 34
15 Airflow behavior: the hidden variable in contamination control 35
15.1 Why airflow is a contamination risk, not just a cooling concern 35
15.2 Making airflow visible: CFD 36
15.3 Where CFD falls short 37
16. When something goes wrong: contamination incident response 38
PART V 39
The Business Case 39
17. Building the internal case 39
17.1 Lead with the cost nobody is counting 39
17.2 The five costs, in the order they persuade 39
17.3 Anchor the case to standards 40
18. Sustainability accounting done honestly 40
19 Implementation roadmap: the first 90 days and the first year 41
19.1 Phase 0 — Design and planning (before construction, and up to ~90 days before turnover) 41
19.2 DAYS 0–30 · BASELINE 42
19.3 DAYS 30–90 · DESIGN AND CLOSE THE URGENT GAPS 42
19.4 MONTHS 3–12 · OPERATE, VERIFY, INSTITUTIONALIZE 42
FROM THE PUBLISHER 44
20. Where Dycem Fits 44
20.1 The Contamination Pathway Dycem Addresses 44
20.2 Published Performance Claims 44
20.3 Product Selection 45
Product Overview 45
APPENDICES 46
The Toolkit 46

A. Facility self-assessment scorecard 47
View more data center definitions here. 49
C References and sources 50
D Disclaimers 53

PART I

1 Executive summary

For most of the data center era, contamination control was treated as a line item associated with janitorial services. It was seen as a matter of visual cleanliness rather than reliability or uptime protection and that framing no longer holds. The move to high-density AI compute has changed the underlying physics of the white space, and with it, the stakes of keeping that room clean.

Racks that once drew tens of kilowatts now draw hundreds. Cooling has moved from air alone to liquid delivered directly to the chip. Component geometries have shrunk while conductor spacings have tightened. Each of these shifts makes the same particle, the same trace of corrosive gas, and the same metallic whisker more consequential than it was a hardware generation ago. Contamination is no longer primarily an aesthetic concern; it is a reliability and energy-efficiency concern that sits directly on the critical path of the most expensive assets in the building.

Three ideas organize everything that follows.

- **Contamination is a chronic cost before it is an acute risk.** Dramatic failure like a corroded connector, a whisker short, or an outage is what gets attention and a well-run facility spends heavily to keep those events uncommon. But that spending doesn't touch the quieter problem: contamination steadily raises cooling energy, eats into thermal headroom, and shortens hardware life, none of which ever shows up on an incident report. That daily drain runs regardless of your uptime tier and it's where most of the recoverable money sits.
- **No single control is sufficient, because contamination has no single source.** It enters through the air handlers, walks in on shoes and wheels, is generated inside the room by legacy materials and maintenance activity, and circulates inside the coolant loop. A credible program is layered: each layer covers a vector and the others structurally cannot
- **Credibility is the organizing principle.** Every target is mapped to a recognized standard, every mechanism to published literature, and every boundary is stated plainly including the boundaries of floor-level control, the layer this guide's publisher makes its living in.

The guide is organized in five parts:

- Part I establishes what changed in contamination control and what it costs.
- Part II covers the science behind it: what contamination is, how it damages hardware, and how it moves.
- Part III maps the standards landscape and measurement practices that make it actionable.

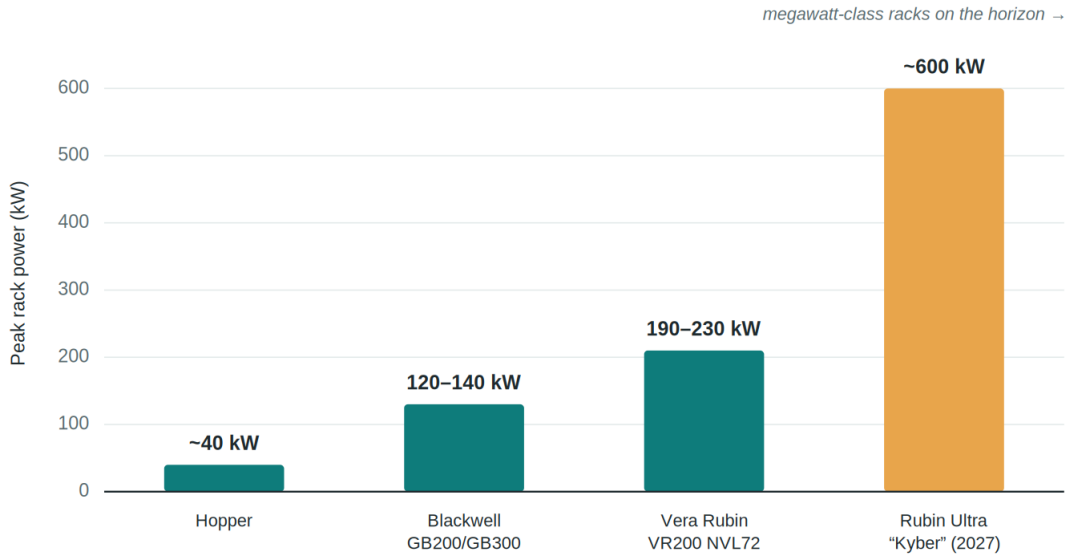
- Part IV lays out the contamination control program itself. Details the layered model, the zoned implementation, floor-level capture, liquid cooling, construction and retrofit, cleaning operations, and incident response.
- Part V builds the business case and closes with an implementation roadmap.
- The appendices convert the guide into working tools: a self-assessment scorecard, a glossary, and a sourced reference list.

2. Why contamination control matters more in the AI era

Four structural changes have moved contamination control from the margins to the core of data center reliability. None is speculative; all are visible in current hardware roadmaps and published environmental guidance. Taken together, they explain why environmental practices that were adequate for a 10-kW rack are not adequate for a 200 kW one and why the gap widens with every hardware generation.

2.1 The density shock

Rack power density is climbing faster than most facilities were designed to accommodate. A Hopper-generation rack drew on the order of **40 kW**. Blackwell-class NVL72 systems (GB200/GB300) draw roughly **120–140 kW**. Vera Rubin VR200 NVL72 racks arriving in the second half of 2026 are estimated at approximately **190–230 kW**, and the 2027 Rubin Ultra "Kyber" design is specified in the region of **600 kW**, with megawatt-class racks on the horizon. While exact figures vary by configuration and source, the major point is the trajectory upwards with density.



Density can have a major effect on contamination for 3 reasons. First, it concentrates value: a single rack can now represent the capital cost of an entire legacy row, so any environmental factor that shortens hardware life is amplified in dollar terms. Second, it concentrates airflow: high-density halls move dramatically more air per square foot, which means more suspended particulate passes through and over the equipment per hour, and settled particulate is more

readily re-entrained. Third, and most important, density leaves no thermal margin to absorb contamination. High-density racks already run close to their temperature ceiling, so a dust film or loaded filter that a legacy rack would have shrugged off now pushes chips into throttling. The thermal margin that used to absorb it is gone.

2.2 The cooling transition

At these densities, air cooling reaches a practical ceiling somewhere around 40–50 kW per rack. Beyond it, liquid cooling is no longer optional. Vera Rubin and its successors are designed as fully liquid-cooled, direct-to-chip systems, and the industry's rack-scale reference designs assume liquid as the primary heat-removal path.

It is tempting to conclude that liquid cooling makes air-side contamination irrelevant. The opposite is closer to the truth: the transition splits one contamination problem into two. The air side still governs the room which includes the people in it, the storage and networking equipment not yet under liquid, the power distribution, and every optical connector in the hall. Meanwhile the liquid side introduces an entirely new contamination regime: particulate in cold-plate microchannels, galvanic corrosion at mixed-metal joints, and biological growth in warm-water loops. The near-term operational reality is that facilities teams now own both problems simultaneously, usually with monitoring practices designed for only one of them.

2.3 Chronic cost versus acute risk

The most useful way to think about contamination is to separate two very different cost profiles. **Acute risk** is the discrete failure event like an unplanned outage. It is visible, attributable, and easy to argue about it after the fact. **Chronic cost** is the quieter drift. It could be marginally higher fan and pump energy or filters loading faster than expected. It rarely produces a single dramatic incident, which is part of why it is under-managed.

In a high-density, energy-constrained facility, the chronic cost is where most of the money lives. Framing contamination control purely as outage protection undersells it. The more defensible case is that it protects both uptime and efficiency over the full-service life of the asset.

3. The economics of dirt: Where the cost shows up

It is worth being precise about the cost categories a contamination program protects against. Vague appeals to “risk” invite skepticism; a named cost structure invites measurement. Five taxes on the facility cover this.

3.1 The cooling-energy tax

Insulating films on heat-transfer surfaces raises the temperature differential the cooling system must overcome. Loaded air filters raise fan static pressure and fan energy. Fouled coolant channels raise pump energy and localized temperatures. Each effect is individually small and collectively persistent, meaning it shows up in PUE and in the cooling share of facility power, every hour of every day, invisibly folded into a utility bill nobody itemizes. In a facility where cooling is already one of the largest energy lines, small, persistent increases compound into meaningful money.

3.2 The performance tax

Modern accelerators run against two ceilings: a power limit and a temperature limit. When cooling is healthy the chip is power-limited: it draws its full rated wattage and delivers full throughput. An insulating film makes the chip run hotter for the same work, and past a point it clocks itself down to stay under the temperature ceiling. The chip isn't doing more. It's just running slower to stay cool, so you're paying full power for less computing. In synchronous AI training the penalty compounds, because the run moves in lockstep and waits for the slowest node at every step. In simple terms, a contamination problem confined to a few racks can set the pace for a job spanning hundreds. It's rarely traced back to contamination, because throttle events live in the GPU telemetry stack while particle counts and inlet temperatures live in the facilities stack. The two are seldom joined, so these events may never be investigated fully.

3.3 The hardware-life tax

Corrosion, ESD events, thermal stress, and fan wear all shorten equipment life. The effect is statistical rather than dramatic such as optics that degrade early. Because replacement is budgeted as a routine capital cycle, the acceleration is rarely traced back to environment. At AI-class hardware prices, even small changes in effective service life move real capital.

3.4 The warranty and compliance exposure

Major OEMs increasingly state environmental requirements to be ISO 14644-1 Class 8 particulate conditions in hardware installation and support documentation. Where dust- or corrosion-related damage occurs in an environment that demonstrably failed those conditions, warranty and support positions can be contested. The takeaway is simple: keep records. A facility that measures and logs its environment can prove it was compliant when a claim is challenged. One that doesn't is stuck arguing from memory.

3.5 The acute tail: outage risk

Discrete failure could be a failure like a clogged cold plate cooking a GPU. These events are rarer than the chronic categories but carry the familiar catastrophic price tags Downtime for AI training and inference facilities is routinely estimated in the hundreds of thousands of dollars per hour. The rare outage gets the program funded. The everyday chronic costs are what get funded properly.

PART II

The Science

What contamination is, the physical mechanisms by which it damages hardware, and how it moves through a facility.

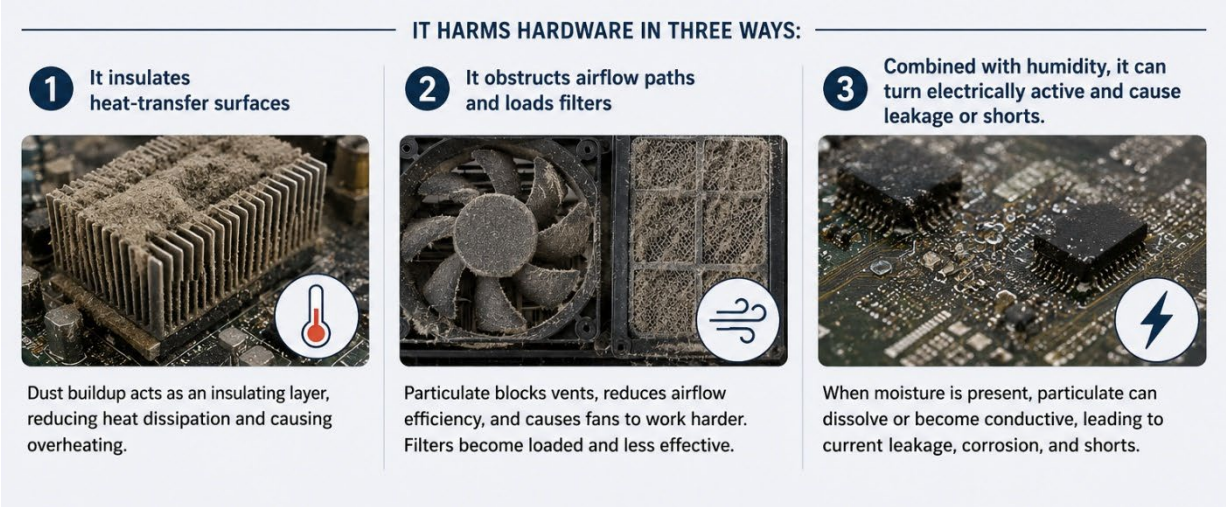
4 The Types of Data Center Contamination

Contamination arrives through multiple sources and acts through physical mechanisms, and the controls that address one category are frequently useless against another. Controlling contamination well starts with naming the categories precisely.

4.1 Particulate

Particulate is the solid material that floats in the air and settles on surfaces — mineral dust, textile and paper fibers, skin flakes, salt (especially in coastal and de-iced environments), and construction residue like gypsum and concrete fines. It harms hardware in three ways:

- It insulates heat-transfer surfaces
- It obstructs airflow paths and loads filters
- combined with humidity, it can turn electrically active and cause leakage or shorts.



The recognized cleanliness target for data halls is ISO 14644-1 Class 8, the same benchmark cited in major IT-equipment installation documentation.

Particulate also behaves in two very different ways depending on how fine/small it is. Larger particles like visible dust, fibers, and grit settle quickly, pile up on floors and in voids, and move mainly by being tracked and re-suspended by traffic and airflow. Fine particles stay airborne

far longer, slip through lower-grade filtration, and settle on internal surfaces including optics and heatsink fins. The practical consequence: a program focused only on visible cleanliness manages the coarse dust and misses the fine, while a program focused only on-air filtration where it manages the fine dust while the coarse reservoir quietly rebuilds underfoot. Both fractions matter, and each is controlled at a different point.

4.2 Gaseous and corrosive contamination

Reactive gases cause corrosion on the circuit boards and connectors. Gaseous contamination is frequently the more damaging category in industrial, urban, and coastal sites and it is invisible to the eye, which means a hall can look immaculate while its air is actively corroding connectors.

Sources are both external and internal. External: industrial emissions, vehicle exhaust, marine air, agricultural activity, and diesel exhaust during extended generator runs. Internal: off-gassing from construction materials, sealants, and packaging; batteries; and cleaning chemistry itself, which is one reason cleaning-product selection belongs inside the contamination program rather than outside it.

4.3 Zinc whiskers and metallic filaments

Electroplated zinc surfaces can spontaneously grow fine conductive filaments known as zinc whiskers. Growth is slow, taking years to decades, which is precisely why the risk concentrates in legacy facilities being retrofitted for AI loads. Disturbed during floor maintenance or cable work, whiskers become airborne, are drawn into equipment, and can bridge conductors to cause intermittent faults and shorts. They are a textbook example of internally generated contamination: no outdoor air handling addresses them, because the source is inside the room. NASA's metal-whisker program at Goddard Space Flight Center remains the recognized technical authority noting that the risk to modern electroplated products may be overstated is worth reading alongside it and to assess legacy surfaces rather than either panic or dismiss.



Figure 1- NASA Goddard Tin Whiskers

4.4 Electrostatic discharge (ESD)

Not a particle, but a contamination-adjacent hazard that shares the same control zones. Low humidity and insulative surfaces let static charge build on people and materials and

discharge into sensitive electronics. In a steady-state hall running within the ASHRAE humidity band, ambient ESD events are relatively contained. It's the real exposure that happens in handling and maintenance with activities like de-boxing, board swaps, cable work, and retrofit which is why staging areas warrant the same ESD discipline as the hall itself. ASHRAE research has tied ESD failure risk directly to humidity level, flooring system, and footwear, which is what puts it in the same conversation as particulate and gaseous control: it's governed at the same thresholds and entry points, and the flooring decision frequently serves both purposes at once.

4.5 Moisture and deliquescent humidity

Humidity is a more of a multiplier rather than a contaminant. This is why ASHRAE guidance couples humidity limits with particulate and gaseous limits as the categories interact, and controlling one without the others leaves the interaction unmanaged. ASHRAE TC 9.9 recommends operating around 40–60% relative humidity. The consequence of either high or low RH can both affect contamination. Humidity kept too high activates settled hygroscopic dust and accelerates corrosion, while humidity kept too low elevates ESD risk. ASHRAE additionally ties the humidity ceiling to measured corrosivity: sites with low copper/silver reactivity can run up to 70% RH, while sites showing corrosive gases (copper above ~300 Å/month or silver above ~200 Å/month) should hold below 50% RH and consider gas-phase filtration. Dust and gas load, determine how much margin the ideal range really has.4.6 Coolant-side contamination

Direct-to-chip liquid cooling existed in high-performance computing for years, but it became standard for mainstream AI racks only with the 2024–2025 hardware generation. Direct-to-chip loops introduce particulate, corrosion products, and biological growth on the fluid side. As cooling moves into the rack, a share of contamination risk moves with it. This is a domain most facilities teams have never had to monitor before.

5 Mechanisms of harm

Different contaminants damage hardware through different physical pathways. Mapping each contaminant to mechanism and each mechanism to an observable symptom is what turns a general concern into a specific, testable claim to work backwards from a symptom to a probable cause.

Contaminant	Primary mechanism	Observable symptom
Fine particulate	Insulating film on heat-transfer surfaces; airflow obstruction; filter loading	Rising component temperatures; throttling; higher fan/pump energy; shorter filter life
Sulfur/chlorine gases	Creep corrosion and connector corrosion on copper/silver	Intermittent connector faults; discolored boards; latent field failures

Contaminant	Primary mechanism	Observable symptom
Zinc whiskers	Conductive bridging across adjacent traces or pins	Intermittent shorts; unexplained transient faults after floor work
Hygroscopic dust + humidity	Electrolyte formation; leakage current; accelerated corrosion	Leakage-related faults; corrosion appearing faster than gas levels alone predict. ASHRAE TC 9.9 recommended range 40–60%
Coolant particulate	Microchannel obstruction; localized hot spots; erosion-corrosion	Cold-plate hot spots; degraded thermal performance; downstream fouling
Static charge (ESD)	Discharge into sensitive junctions	Latent or immediate device damage; elevated infant-mortality rates

Two themes run through the table. First, most mechanisms show up as efficiency loss before they show up as failure. Second, several of the most damaging mechanisms (whiskers, creep corrosion, coolant fouling) originate inside the controlled space, which means air handling alone cannot address them. That observation is the basis for the layered strategy in Section 9.

6 How contamination moves: the cycle and the reservoirs

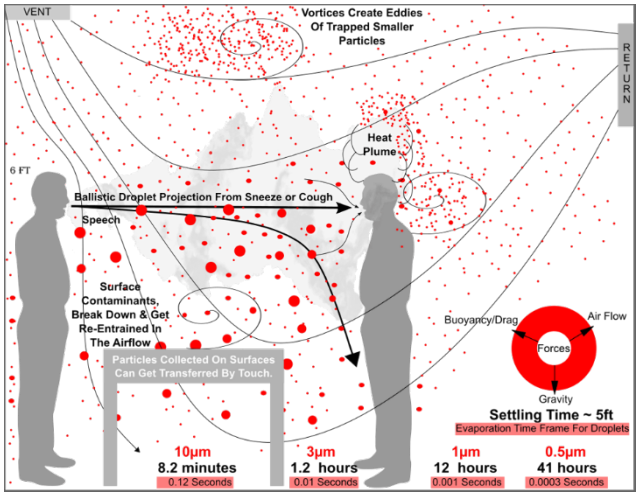
Understanding where contamination comes from is half the picture. The other half is understanding how it behaves once inside because its movement pattern is what defeats single-layer control strategies. This matters more under the AI buildout than it used to: capacity is being added so fast, and fit-out so often runs alongside live operation, that construction-grade contamination and normal operation now happen in the same space at the same time. Every vector in this section is amplified when the room is never allowed to fully settle, which is exactly when holding an ISO 14644-1 Class 8 environment gets hardest.

6.1 The entry vectors

Contamination reaches the data hall two ways: it comes in from outside, and it's generated by activity within. The inbound path is the one most controls with the data center target. External air enters through the air handlers and through infiltration at doors, penetrations, and pressure imbalances. Materials enter as packaging and the equipment itself. People and wheeled carts carry soil, fibers, salt, and fine grit directly past the air handling directly and onto the floor. In cleanroom practice, this floor-level channel of foot and wheel traffic is consistently treated as the dominant controllable entry vector. Every person and every wheel that crosses an uncontrolled threshold is a delivery mechanism that bypasses the air-handler entirely.

6.2 Resuspension: Why settled is not gone

Once inside, contamination does not stay put. Foot traffic and high-velocity cooling airflow lift settled particles back into the air by a process known as resuspension where server fans can pull them directly through the equipment. The more CFMs within the white space, the more readily settled material re-enters circulation. Fine particles are the persistent problem: once airborne, they're small enough that random collisions with air molecules (Brownian motion) keep them suspended far longer than gravity would settle them, so they stay in circulation until something captures them. This is why having a flooring type that merely relocates dust, or a cleaning process that agitates more than it removes, can leave the air measurably dirtier than before the “cleaning” began.



6.3 The hidden reservoirs

Raised-floor plenums, ceiling voids, cable trays, and the undersides of tiles accumulate contamination over years and release it during the very maintenance activities meant to help. These types of events could be tile lifts, cable pulls, or retrofit construction. A data hall can

pass a surface inspection while sitting on a reservoir that redistributes contamination into the air every time the floor is opened. The operational rule should be followed is that the cleaning scope must include the voids, and any planned floor disturbance should be treated as a contamination event with monitoring before and after.



Figure 2- Credit to Steven yates

6.4 The self-sustaining loop

Entry, settlement, resuspension, and reservoir release form a loop. Left unchecked, the loop sustains itself: you remove visible contamination, and the cycle refills the air from the reservoirs while new material walks in at the threshold. This is the structural reason you cannot clean your way out of a contamination problem once it is established and is the beginning of the loop. Stop material crossing the threshold, and every downstream stage has less to fight; the reservoirs stop growing, resuspension has less inventory to work with, and routine cleaning becomes maintenance instead of triage.



The leverage point is the threshold: stop material entering, and every downstream stage has less to work with.

PART III

The Standards

7. The standards and frameworks landscape

Seven frameworks cover the territory. Three set measurable environmental targets (ISO 14644-1, ISA-71.04, ASHRAE TC 9.9). Two govern adjacent disciplines that share the same control zones (ANSI/ESD S20.20, EN 50600). One is an operational program framework (the DCA Anti-Contamination Guide). And two connect the program to sustainability practices (GHG Protocol, ISO 14044). Knowing which is which prevents the most common misuse: citing a classification standard as though it were a control method, or a program guide as though it were a measurable target.

7.1 ISO 14644-1 — air cleanliness classification

Originally written for cleanrooms, ISO 14644-1 classifies air by particle concentration across size thresholds. Class 8 is the widely referenced benchmark for data halls and is cited in major server-installation documentation. Classes with lower particulate counts apply to semiconductor and advanced-electronics manufacturing, not typical data centers. The standard's value is that it converts “clean” from a subjective impression into a measurable, auditable particulate target with a defined test method. This method takes sampling counts of defined particle sizes, at defined sample points, with defined protocol.

ISO Room Class Particle Count

Class	Maximum Particles/m³						FED STD 209E equivalent
	>0.1 um	>0.2 um	>0.3 um	>0.5 um	>1 um	>5 um	
ISO 1	10	2					
ISO 2	100	24	10	4			
ISO 3	1,000	237	102	35	8		Class 1
ISO 4	10,000	2,370	1,020	352	83		Class 10
ISO 5	100,000	23,700	10,200	3,520	832	29	Class 100
ISO 6	1,000,000	237,000	102,000	35,200	8,320	293	Class 1,000
ISO 7				352,000	83,200	2,930	Class 10,000
ISO 8				3,520,000	832,000	29,300	Class 100,000
ISO 9				35,200,000	8,320,000	293,000	Room Air

However, there are two important notes to mention with this standard. First, Class 8 is a condition of the air, not of the surfaces. Meaning that a data hall can meet Class 8 while carrying heavy settled contamination, which is why operational cleanliness programs pair the airborne target with a surface and reservoir cleaning scope. Second, classification is a

snapshot taken at the time the samples were taken. The key is to continue monitoring over time in these critical spaces.

7.2 ASHRAE TC 9.9 — thermal and contamination guidance

The **Center of Expertise for Energy Efficiency in Data Centers** published *Gaseous and Particulate Contamination Guidelines for Data Centers*, drawing on ASHRAE Technical Committee 9.9 guidance. The paper helped establish the now-common pairing of an **ISO 14644-1 Class 8 particulate target** with limits for gaseous contamination, and it also connected the transition to lead-free solder with the rise in observed creep-corrosion failures.

7.3 ANSI/ISA-71.04 — corrosion severity

ISA-71.04 classifies an environment's corrosivity by measuring metal reactivity, typically using copper and silver coupons. The 2013 revision added silver reactivity as a required metric, specifically because copper alone misses chlorine-driven attack and over-weights. The severity levels below are the framework operators use to decide whether gas-phase filtration is warranted. G1 is the reliability target embedded in ASHRAE and major OEM guidance.

Severity	Copper reactivity (approx.)	Interpretation
G1 — Mild	< ~300 Å/month	Reliability target for data centers; contamination not normally a factor in equipment failure
G2 — Moderate	~300–1,000 Å/month	Corrosion measurable; may affect sensitive equipment; monitoring advised
G3 — Harsh	~1,000–2,000 Å/month	High probability of corrosive attack; gas-phase filtration typically required
GX — Severe	> ~2,000 Å/month	Only specially designed/packaged equipment expected to survive

7.4 ANSI/ESD S20.20 — electrostatic discharge control

The recognized framework for building and auditing an ESD control program: grounding, flooring and footwear systems, personnel practice, and verification. This is discussed in this document because ESD control shares entry points, flooring, and handling discipline with particulate control, and is most efficiently managed as part of the same zoned strategy. When you're building an AI data center, decide on flooring, footwear, and staging areas with both static and dust control in mind from the start. Handling them together usually costs about the same as handling either one alone and clearly going back to add the one you skipped always costs more.

7.5 EN 50600 and the DCA Anti-Contamination Guide — the operational frameworks

These two frameworks address the program rather than the measurement. The EN 50600 series is the European data-center design and operations standard, and it is increasingly referenced in European procurement. The Data Centre Alliance (DCA) Anti-Contamination Guide is the most directly practical document in the set mentioned in this guide. It frames a complete program around nine areas, all underpinned by accurate site documentation and audit records:

- geographic location and building design
- floors, ceilings, walkways, and layout
- air filters and air-handling units
- waste and packaging management
- cleaning processes
- floor and ceiling voids
- flora and fauna
- monitoring
- contractor selection

7.6 GHG Protocol and ISO 14044 — where sustainability and contamination control meet

Sustainability and contamination control pull in the same direction: keeping a facility clean is also a way of keeping it efficient. Every measure that reduces cooling energy or extends hardware life is also a measure that lowers emissions.

The goal is those gains can be expressed in the frameworks organizations already report against, the GHG Protocol for emissions accounting and ISO 14044 for life-cycle assessment. Cleaner heat-transfer surfaces and less-loaded filters mean less cooling energy, a direct reduction in the electricity-related footprint; longer hardware life spreads each device's manufacturing emissions over more years of use, easing the value-chain footprint.

8. Measurement, monitoring, and documentation

A standard is only useful if you measure against it. Each contamination category has an established measurement method tied to a recognized standard; together they turn contamination from an opinion into a documented, auditable metric. This section covers the methods and why the records matter — it stops short of telling you how to interpret a borderline result, which is site-specific engineering judgment best made with a specialist.

8.1 The measurement toolkit

No single instrument covers contamination, because the categories behave differently with each contamination source requiring their own method, tied to its own standard. The tools below map to the categories in Section 4: particle counting for particulate, reactivity coupons for corrosive gases, humidity trending for the moisture multiplier, and conductivity for the coolant loop. The ideal strategy is to perform them on a defined schedule and keeping the records.

Particle counting. Optical particle counters measure airborne particle concentration against the size and count limits defined in ISO 14644-1. Spot checks capture events such as construction or commissioning; continuous monitoring captures drift over time.

Reactivity (corrosion) monitoring. Copper and silver coupons are exposed in the environment for a defined period per the ISA-71.04 method, then laboratory-analyzed to classify the environment's corrosivity. Real-time reactivity monitors provide a continuous alternative. Silver is included because copper alone misses chlorine-driven attack and is over-sensitive to humidity.

Humidity and temperature trending. Already instrumented in every modern facility. The contamination-relevant step is analytical, not additional hardware: trend humidity against the deliquescence thresholds discussed previously, and treat sustained excursions as contamination-activation conditions, not just comfort deviations.

Coolant quality (liquid-cooled systems). In-line conductivity is the fastest real-time indicator of ionic contamination and is already present in most CDUs; particulate is managed with fine filtration. Both should be monitored to the thresholds set in your OEM's specification rather than to any universal figure, since liquid-cooling guidance is still consolidating.

Surface and reservoir inspection. Periodic inspection with photographic records of subfloor plenums, cable trays, and tile undersides tracks the accumulation points that airborne measurement never sees.

Documentation and audit trail. A defined measurement plan and retained records give operators an auditable trail for warranty, insurance, and environmental reporting.

8.2 Why the records matter

Measurement is not paperwork for its own sake. Three practical payoffs justify it:

- **Warranty protection.** Where OEM support terms require a documented environment (commonly ISO 14644-1 Class 8 and ISA-71.04 G1), monitoring records are what substantiate a claim if contamination-related damage is ever disputed.
- **Change detection.** Continuous or periodic trending flags drift before it becomes damage such a filter loading faster than usual, corrosivity creeping toward G2, or coolant conductivity rising. The value is in the trend, not any single reading.
- **Baseline for incident response.** When something goes wrong (Section 14), a pre-incident baseline is what lets you verify that remediation returned the environment to its prior state. Without a baseline, "clean again" is a matter of opinion.

PART IV

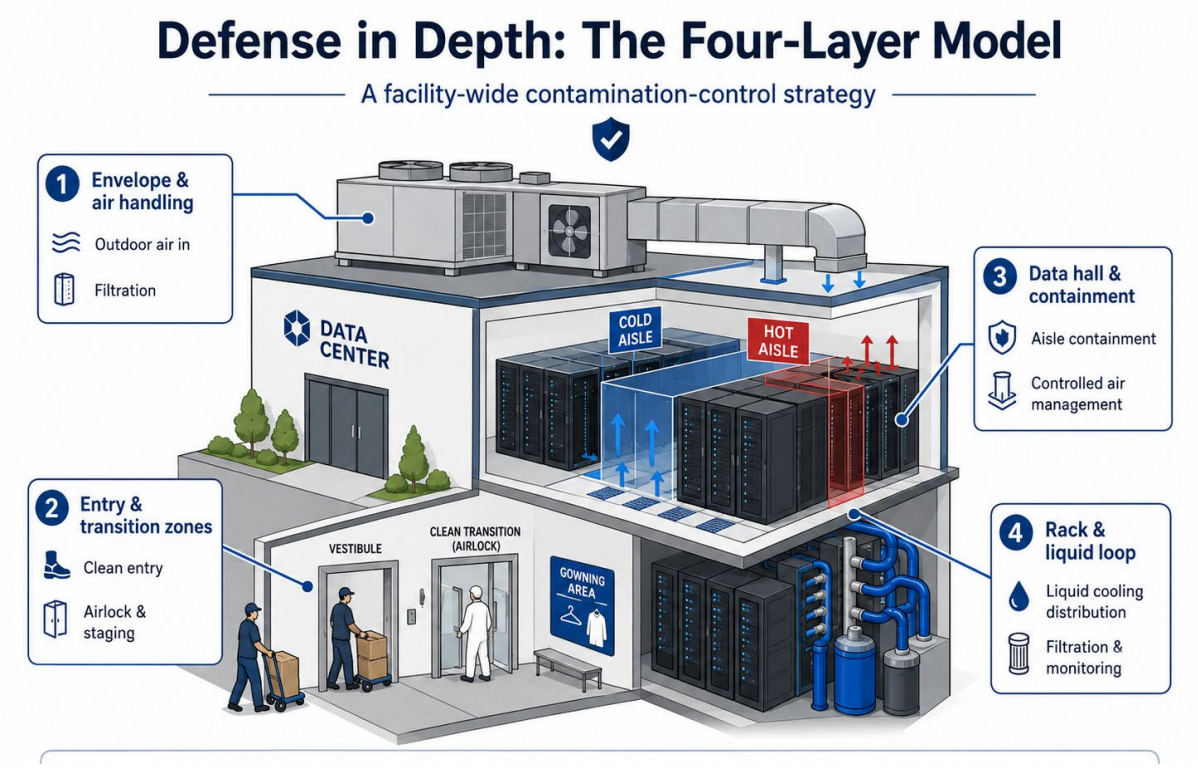
The Contamination Control Program

This part turns the science into a program. It builds the layered control model, walks it out zone by zone, and places floor-level capture it. It then extends to the newer and harder cases: liquid cooling, construction and retrofit, cleaning operations, and contamination incident response.

9. Defense in depth: the four-layer model

Contamination enters and develops within a data center through multiple pathways. Some contaminants originate outside the facility and are introduced through ventilation, personnel, equipment, packaging, and wheeled traffic. Others are generated inside the white space through maintenance activity, material degradation, corrosion, or liquid-cooling systems. Because these risks occur at different points in the environment, no single control measure can address them all.

The most effective approach is therefore **defense in depth**: a layered contamination-control strategy in which each layer addresses a specific pathway and compensates for the limitations of the others. The objective is not to rely on one barrier but to create multiple independent controls from the building envelope through to the equipment itself.

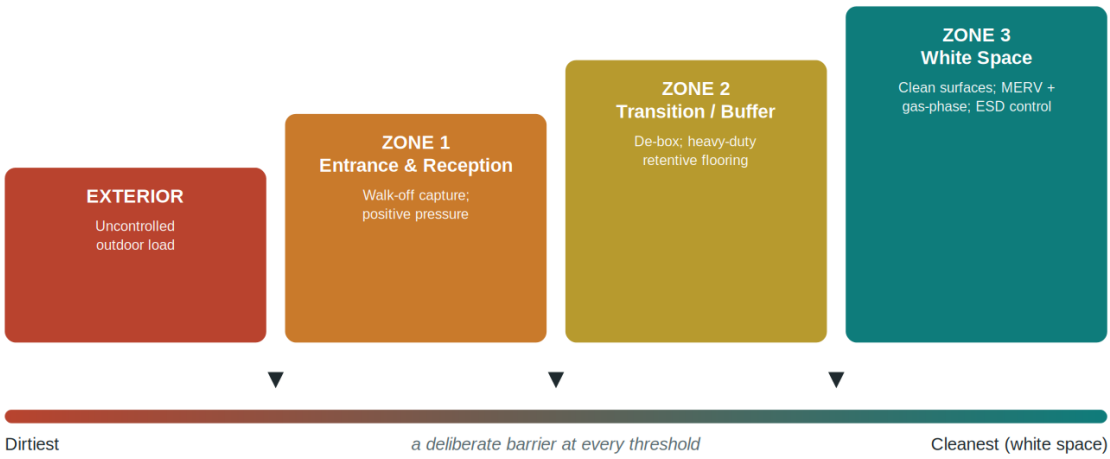


Layer	Scope	What it covers	Doesn't cover
1 · Envelope & air handling	Facility	Bulk-air particulate, positive pressurization, humidity control	Anything generated or carried in downstream of the air handlers and filters
2 · Entry & transition zones	Thresholds	Foot- and wheel-borne particulate captured at every doorway; packaging discipline; gowning where used	Airborne fine particulate, gases, coolant-side contamination
3 · Data hall & containment	White space	Aisle discipline, sealed penetrations, blanking panels, active void/plenum management	Material entering at thresholds; coolant loop
4 · Rack & liquid loop	Equipment	Connector and optics cleanliness; coolant filtration, chemistry, and monitoring	Room-level air quality and traffic-borne load

Layer 1 cleans the air coming into the room, but it can't touch contamination that's created inside or walked in after the air is filtered. Layer 2 handles the biggest blind spot by addressing everything that walks and rolls through the door. Layer 3 keeps whatever gets generated or slips through from circulating freely and cleans out the hidden buildup in floor and ceiling voids. Layer 4 protects the equipment itself including the coolant loop, which has become a reliability concern.

10. The zoned program

The proven implementation model is a cascade of zones, each cleaner than the last, with a deliberate barrier at every threshold. This model has been used for many years within the cleanroom and semiconductor industries. Contamination is stripped away in stages so that very little reaches the white space. The cascade works because it matches how contamination travels: most of the controllable load moves at floor level, threshold by threshold, so that is where the barriers go.



ZONE 1 · ENTRANCE & RECEPTION

- Install a high-performance walk-off / contamination-control zone at every external entrance where it's long enough that each shoe or wheel makes at least three full contacts before the edge.
- Place retentive contamination-control flooring at the first main shared entry point or threshold to capture foot- and wheel-borne particles on contact.
- Control how people and goods enter; remove outdoor footwear or apply overshoes where the protocol calls for it.
- Keep doors closed and maintain positive pressure so air leaks out, not in; treat propped-open doors as the envelope breach they are.

ZONE 2 · TRANSITION / BUFFER

- De-box and unwrap everything in this zone. Cardboard is a major fiber and dust source, and it is the single most commonly violated rule in otherwise well-run facilities.
- During retrofit or construction periods, apply booties and a clean-footwear discipline appropriate to the facility's protocol before the hall.

- Use heavy-duty retentive flooring where wheeled objected cross as pallet and large carts carry loads ordinary walk-off measures cannot capture.
- Stage and inspect deliveries; wipe down equipment surfaces before they proceed. Staging areas should meet the same ESD requirements as the hall.

ZONE 3 · WHITE SPACE

- Maintain contamination-control flooring at aisle entries and, on raised access floors, in formats that lift out for under-floor work.
- Clean all accessible surfaces free from settled dust on a defined schedule.
- Air filtration: MERV 8 on recirculated air and MERV 11–13 on make-up air are commonly referenced baselines; add gas-phase filtration where corrosivity exceeds ISA-71.04 G1.
- Manage ESD: static-dissipative flooring and anti-static tools for all cleaning work, per ANSI/ESD S20.20.

WHOLE FACILITY · ONGOING

- Clean the floor and ceiling voids, including the hidden reservoirs as part of the routine scope, not as an afterthought.
- Maintain AHUs and ductwork; verify filter seating (a bypassing filter performs like no filter) and change schedules.
- Monitor with particle counts and corrosion coupons.
- Use a qualified data-center cleaning contractor as improper cleaning re-suspends dust and does active harm to the white space.

11. Floor-level control: mechanism, fit, and scope

Layer 2 — the entry and transition zone — is where floor-level contamination control does its work. This section describes the vector it addresses, the mechanism behind the more effective forms of it, and, just as importantly, what it does not do.

11.1 The foot and wheel traffic vector

People and wheeled logistics carry contamination directly past the air handling and onto the raised floor: soil and salt on soles, fibers from clothing and packaging, fines picked up from loading docks and corridors, and everything a pallet truck's wheels have rolled through since they last left the building.

This pathway is continuous (it operates every shift), it scales with activity (commissioning and refresh periods generate the most traffic and most importantly, it's entirely untouched by Layer 1 as no air handler can filter a boot. The combination of continuous, activity-correlated,

and structurally invisible to the air-side investment is what makes the threshold the highest-leverage single point of control in the cascade.

11.2 Capture versus retention

Not all threshold controls work the same way, and the difference is mechanical. Conventional approaches like peel-off adhesive ("tacky") mats typically capture and release. They collect particulate but shed it back into traffic as they load, saturate, or dry out, and peel-off mats add a consumable-waste stream and a per-peel labor cost.

Polymeric contamination-control flooring works by a different mechanism: particulate is drawn to and held at the surface by **Van der Waals attraction** and retained until the floor is wet-cleaned, at which point retention capacity is restored and the floor is reused. The distinction is capture-and-release versus capture-and-retain, and it determines whether a threshold is a barrier or merely a delay.

This is where Dycem's polymeric flooring sits within the model. The defensible claim is specific and mechanism-based: at the entry and transition zone, a retentive polymer floor captures and holds floor-level particulate from foot and wheeled traffic more durably than release-prone alternatives, and does so as a cleanable, reusable system rather than a consumable.

Quantified performance figures are the manufacturer's published, independently tested values and are presented — attributed as such — in the closing section of this guide rather than woven through its technical argument.

11.3 What floor-level contamination control does not do

Credibility depends on being explicit about boundaries. Floor-level contamination control addresses particulate tracked in at the threshold. It does not filter the bulk air; it does not scrub corrosive gases; it does not manage coolant chemistry; and it does not, on its own, deliver an ISO Class 8 or ISA G1 environment. This solution complements air-side filtration, gas-phase control, containment, and coolant management. It does not replace any of them. Positioned this way, floor-level control is a strong contributor to Layer 2 of the defense-in-depth model, and its benefits. Reduced particulate load reaching hardware, lighter downstream filter duty, and cleaner subfloor reservoirs are real and substantiable precisely because they are not overstated.

12. Liquid cooling: the new contamination frontier

As direct-to-chip liquid cooling becomes standard for AI racks, a contamination regime that barely existed a few years ago becomes first-order. The fluid side deserves the same rigor operators already apply to the air side because the discipline, the instrumentation habits, and the institutional experience haven't caught up with the hardware. This section maps the fluid-side risk categories, the monitoring practices, and the operational disciplines that close the gap.

12.1 Why the fluid side is unforgiving

Three properties make coolant-side contamination categorically harder than air-side contamination.

First, tolerance: cold-plate microchannels are engineered to fine dimensions for turbulence and heat transfer, so the particle sizes that matter are far smaller than anything the air side worries about. Second, coupling: the loop is a closed, shared circulatory system meaning contamination introduced anywhere reaches everywhere. Third, opacity: you cannot see into a loop. Air-side contamination eventually announces itself as visible dust; fluid-side contamination announces itself as a thermal anomaly on your most expensive silicon, usually well after the causative event.

12.2 The four fluid-side risk categories

- **Particulate in microchannels.** Assembly debris, corrosion products, and precipitated solids partially obstruct flow channels, create localized hot spots, and drive erosion-corrosion downstream. System-level filtration around sub-50 µm is a commonly referenced target in vendor documentation, with finer bypass filtration where mixed metals can shed corrosion products.
- **Coolant chemistry drift.** Rising conductivity signals ionic contamination and accelerating galvanic corrosion at dissimilar-metal interfaces (copper cold plates, aluminum manifolds, brass fittings, stainless piping — most loops are mixed-metal whether designed to be or not). Inhibitor packages deplete over time; makeup water introduces ions; every service event is a chemistry event.
- **Biological growth.** Warm-water loops are hospitable to biofilm, which fouls heat-transfer surfaces, degrades flow, and can accelerate corrosion under the film (microbially influenced corrosion). Fluid selection, biocide management, and loop hygiene move from optional to essential as inlet temperatures rise — and the industry is deliberately raising inlet temperatures to improve plant efficiency.
- **Materials compatibility and leaching.** Hoses, gaskets, and seals can leach plasticizers and particulate into the fluid, especially early in life. Wetted materials review belongs in procurement, not in post-incident analysis.

12.3 Monitoring the loop

In-line conductivity is the fastest and cheapest real-time check, and the sensors are already built into most coolant distribution units (CDUs). For many operators the gap isn't hardware, it's actually watching the data. A practical baseline is needed where you look at: trend conductivity, filter pressure differential, flow per rack or cold-plate group, and supply/return temperature differences; send fluid out periodically for lab analysis (particle counts, metals, inhibitor levels, biological activity); and inspect filter elements at each service as a diagnostic, not just a swap where you look at what the filter caught tells you what the loop is doing.

12.4 Operational discipline

- **Commission clean.** Flushing, fill-water quality, and cleanliness verification at commissioning determine the loop's starting condition. Loops are rarely dirtier than the day they were filled unless they were filled improperly.
- **Treat every connection event as a contamination event.** Quick-disconnect operations, cold-plate swaps, and hose replacements each introduce risk to the process.
- **Specify coolant quality as a parameter, not a default.** Fluid specification, monitoring thresholds, and maintenance chemistry should be documented, owned, and audited — not inherited silently from whatever the integrator shipped.
- **Plan the leak response before the leak.** Leak detection, isolation, cleanup materials, and refill procedure belong in the incident playbook from day one of liquid operation.

Operators operating liquid-cooled AI capacity now should treat coolant quality as a monitored, specified parameter from day one and plan the standards to firm up over the next several years.

13. Construction, commissioning, and retrofit

The dirtiest days in a data hall's life are the days it is built, fitted out, and upgraded. Construction generates gypsum and concrete fines, metal shavings, and fiber loads orders of magnitude beyond anything in normal operation. A contamination program that starts at turnover or go-live has already missed the largest single loading event the facility will ever experience.

13.1 New builds: cleanliness as a commissioning deliverable

- **Schedule dirty work before cleaning work.** Dust-generating work (cutting, drilling, drywall) should be complete and verified clean before sensitive fit-out begins in a space; where overlap is unavoidable, physical segregation and temporary barriers are the minimum.
- **Protect the voids before they close.** Subfloor plenums and ceiling voids should be cleaned and inspected with photographic records before tiles and panels seal them.
- **Control construction traffic like operational traffic.** Threshold captures at construction entrances, designated dirty/clean routes, and keeping cardboard and packaging out of the clean area during fit-out cost little and prevent the site from pre-loading its own reservoirs.
- **Make cleanliness verification a handover deliverable.** Particle counts to the ISO 14644-1 method (sampling airborne particles at defined points and sizes to confirm the space meets its cleanliness class), a documented void inspection, and for liquid-cooled capacity loop flush and fill-water verification with conductivity records.

13.2 Liquid-loop commissioning

Loop cleanliness is determined at commissioning more than at any other moment. The discipline: flush per the OEM procedure until returned fluid meets specification; verify fill-water quality before it enters the loop, not after; confirm filtration elements are installed and seated (loops have been commissioned with filter housings empty); record baseline conductivity, flow, and pressure-differential values. These baselines are what every future anomaly will be measured against; and retain samples and records in the commissioning package.

13.3 Retrofit in live facilities: the hard case

Most AI capacity in the next several years will land inside existing facilities, which means construction-grade contamination events inside operating white space. Three disciplines carry most of the weight.

- **Contain the work zone.** Temporary barriers with sealed penetrations, negative pressure inside the work zone relative to the operating hall, and dedicated entry points with threshold capture. The goal is that the work zone behaves like a separate, dirtier building.

- **Treat floor disturbance as a monitored event.** In legacy facilities, inspect tile undersides for zinc plating and whisker evidence before large-scale changes. Schedule particle monitoring before, during, and after and lastly stage under-floor work to limit how much reservoir is open at once.
- **Protect the hardware paths.** New equipment should not travel through active construction routes. De-box in the buffer zone and during retrofit periods, enforce it twice as hard, because delivery volume and construction dust peak together.

Retrofit is where contamination programs are won or lost. The facility that treats an AI retrofit as a construction project with a cleanliness plan inherits a clean data hall.

14. Cleaning operations and contractor qualification

Cleaning is where the contamination control program is put into practice, but an unqualified approach can do more harm than good. Dry sweeping re-suspends more than it removes; household vacuums without HEPA-grade filtration exhaust fine particulate directly into the airstream; wrong chemistry leaves residues or attacks floor finishes; and careless work around equipment causes the ESD and physical-contact incidents the program exists to prevent.

14.1 Scope of a proper cleaning program

- **Surfaces and floors** should be cleaned on a defined schedule to maintain critical-clean conditions, with all accessible areas free from settled dust and particulate.
- **The Reservoirs.** Subfloor plenums and ceiling plenums on a periodic cycle, and always ahead of planned floor disturbance. This is the item most commonly missing from incumbent cleaning contracts as this takes the most amount of coordination and trust in the vendor.
- **Retentive flooring maintenance.** Wet-clean polymeric contamination-control flooring on the schedule its loading demands. This cleaning restores retention capacity, so the mat program and the cleaning program are part of one contamination control program.
- **Equipment-adjacent work** with anti-static tools, HEPA-filtered vacuums, and approved chemistry only; no liquids near energized equipment beyond the approved method.
- **Documentation.** Scope, method statement, chemistry list, and a certificate of cleanliness with test results after each major cleaning. This audit trail helps supports warranty, insurance, and reporting positions.

14.2 Qualifying a data-center cleaning contractor

In cleaning a data center, look beyond familiarity with industry standards and evaluate whether the company has a documented cleaning program, proven experience in controlled environments, and technicians trained to work safely in a live data hall: Ask a prospective contractor:

- Do they have a formal critical-cleaning program with a detailed, step-by-step approach? This is not just a general commercial-cleaning checklist.
- What's their track record in critical and controlled environments? Do they also clean cleanrooms or semiconductor fabs? Experience in those spaces is a strong signal they understand the discipline a data hall needs.
- Are they certified at ISO 9001? It's a good proxy for whether they'll follow a defined process and leave you with proper documentation at the end.
- What training do their technicians go through before they set foot in the space? Do they know what they can and can't touch, how to behave in a live data hall (no photography, no contact with equipment or cabling they aren't there to clean), and what to do if a spill or accidental contact occurs?

- What HEPA-filtration standard does their vacuums meet, and how is it verified? How often are they changing the HEPA Filters on the vacuums?
- Have they cleaned raised-floor plenums, and can they show references?
- What chemistry do they use, and what's excluded? (Chlorine-based products near electronics are a classic exclusion.)
- Do they provide post-clean particle verification and a certificate?
- How do they control their own contamination like equipment, packaging, and footwear on the way in?

How they answer these tell you everything: the right contractor is protecting your hardware; the wrong one is putting it at risk.

Get a printable checklist to walk through with your facility cleaning team here.

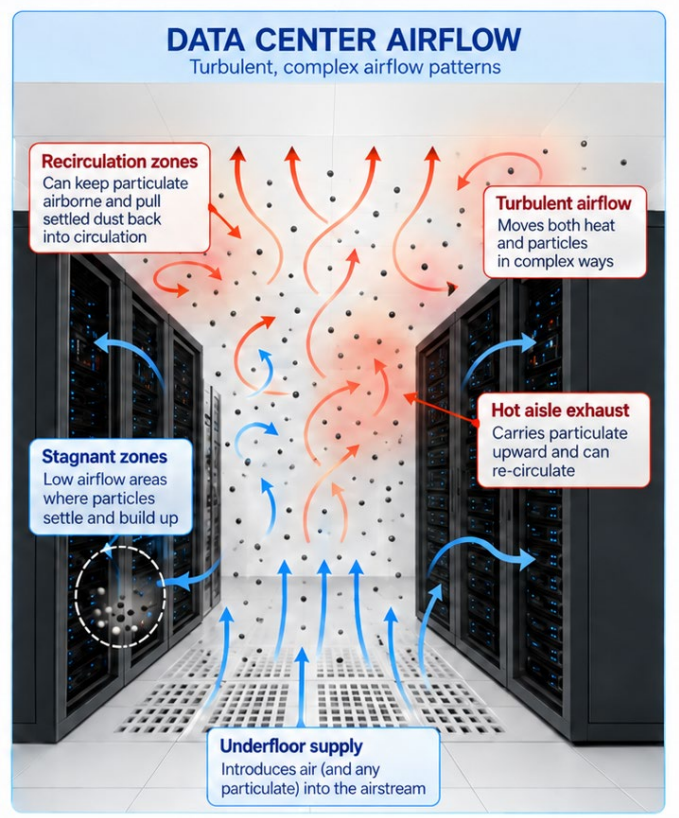
15 Airflow behavior: the hidden variable in contamination control

Airflow is not just what cools the room. It is what distributes contamination through it. In a modern data hall, airflow is neither uniform nor intuitive: it channels, recirculates, stagnates, and accelerates in ways invisible to the eye and rarely matches the design intent. Two halls with identical filtration and cleaning schedule can end up with very different contamination outcomes purely because of how their air moves. Understanding that movement is one of the first steps in prevention, and the step most programs skip.

15.1 Why airflow is a contamination risk, not just a cooling concern

Airflow analysis in data centers is almost always done for cooling where it tries to find the thermal hot spots and cut the cooling energy. While that work is valuable it also maps the forces that govern contamination, because the same air patterns that move heat move particles.

Two features of that map matter most. Stagnant zones (dead pockets where air barely moves) are where suspended particulate loses velocity and settles out, quietly building up that later re-entrain during maintenance or even during regular activity. High-velocity airflows loops can be the opposite problem. These airflows can keep particulate airborne and pull settled particulate back into circulation before it can be cleaned. A data hall's contamination behavior is a map of where these two conditions sit relative to the hardware. A rack downstream of recirculation airflow is re-presented with the same dust. You cannot see that map by walking the floor.



15.2 Making airflow visible: CFD

Computational Fluid Dynamics (CFD) builds a three-dimensional digital model of a room and simulates how air, heat, and particulate move through it. By using the equipment layout and heat loads, the cooling supply, and the floor and ceiling details, it solves the physics and renders the result visually. CFD is usually run to find thermal problems before they cause throttling. The same model can also track particles: where they travel, where they settle, where they linger.

A CFD model answers questions a facility manager, or designer, otherwise can only guess at:

- Where does particulate entering a given doorway end up?
- Which racks sit downstream of a loop that keeps recirculating the same dust?
- Where are the stagnant zones that build reservoirs fastest, and which aisles carry the most particulate past open equipment?

These are contamination questions and often they can be answered by the same model the cooling team is already building. That makes airflow one of the cheapest contamination analyses a facility can run.

A CFD map also shows where contamination gets into the room and where it travels once inside. This is exactly where floor-level capture and entry controls (Section 11) do the best. Air flow analysis doesn't just find the problem; it shows you where to put the controls.

15.3 Where CFD falls short

CFD is a powerful tool, not a magic one, and a credible program is clear about its boundaries. Full high-fidelity CFD is computationally expensive and slow where modeling a single room can take hours to days, and it needs an expert to set up and interpret correctly. A model is also a snapshot of one specific configuration: move servers, open different floor tiles, shift the workload, or change a cooling setpoint, and the airflow pattern changes. This is why CFD has historically lived in the design and major-change phases rather than in daily operation or even as part of the contamination control program.

The field is moving toward faster methods with calibrated models, reduced-order surrogates, and sensor-fed digital twins that approximate CFD closer to real time. However, these faster methods are still maturing. Understanding how the air behaves is a prevention step; CFD is simply the clearest way to get that understanding when the stakes justify it.

16. When something goes wrong: contamination incident response

Contamination incidents such as a construction breach, filter failure, or smoke or ash intrusion from a nearby fire are rare enough that most facilities have no rehearsed response, which converts a manageable event into an extended one. A lightweight playbook covers most cases in four moves.

- **Characterize before you clean.** Deploy particle counts and, where corrosion is suspected, expedited reactivity monitoring; photograph and sample before remediation destroy the evidence. What entered, how far it spread, and whether it is conductive or corrosive determine everything downstream including whether equipment can safely keep running.
- **Stop the source and contain the spread.** Close the breach, isolate the affected air path or loop segment, restrict traffic through the affected zone, and put threshold capture at its boundary so response traffic doesn't spread the event.
- **Remediate to a measured endpoint.** Clean the contamination control program standard and verify by measurement against the pre-incident baseline. This is why routine monitoring matters as without a baseline, “clean again” is a matter of opinion.
- **Document for downstream conversations.** Maintain a time-stamped record of the incident, findings, corrective actions, and verification results. Clear documentation supports warranty discussions, insurance claims, customer communications, and internal reporting. Conclude with a formal review to determine which control failed, why it failed, and what corrective or preventive measures are required.

PART V

The Business Case

Turning the technical case into a funded program to control a preventable source of operational risk around extremely valuable infrastructure: *sustainability accounting, and a practical roadmap.*

17. Building the internal case

Contamination control loses more budget battles to framing than to evidence. Presented as janitorial services, it's the first line cut in a tight quarter. However, presented as what it is as a reliability-and-efficiency investment protecting the most valuable assets in the building, it competes on equal footing with any other infrastructure spend.

This section is about those terms: which costs to put forward, in what order, and which arguments to avoid because they invite the wrong kind of scrutiny.

17.1 Lead with the cost nobody is counting

The instinct is to lead with the outage. This can be a dramatic failure or a scary number. But to clearly show the case, resist that urge. Everyone in the room already knows and believes an outage is bad. Leading down that path doesn't show them something they already don't know, and it frames the whole program as insurance against a low-probability event.

Lead instead with the cooling-energy penalty, because it is the cost the organization is *already paying and not counting*. Contaminated heat-transfer surfaces and loaded filters make the cooling plant work harder for the same result, every hour of every day. Because it never surfaces as an event, it disappears into the facility's power draw, unattributed. In a facility where cooling is one of the largest energy lines, a continuous penalty is not a rounding error; it is real money leaving the building silently. The power of leading here is that it reframes the conversation from "protect us against something that might happen" to "stop paying for something we're losing right now."

This is the core of the case: most of the recoverable money is in the chronic costs, not the rare failure. A program sold as outage insurance gets funded grudgingly and cut first; a program sold as efficiency recovery gets funded for what it is.

17.2 The five costs, in the order they persuade

Present the costs in order of how often they go unattributed. The cost the audience already recognizes adds little; a cost they hadn't been counting is what shifts the conversation. Contamination imposes five distinct costs:

Cooling energy - Covered above; it leads because it's continuous, large, and currently invisible.

The acute outage — A contamination-driven failure carries a catastrophic price tag, and it belongs in the case. Present it with the organization's own cost-of-downtime figure and a credible, conservative frequency, and let the audience see you being measured rather than alarmist. An honest low-probability, high-cost line earns more trust than a dramatic one, and trust is what the whole case runs on.

Hardware life — Corrosion, thermal stress, and wear will shorten effective equipment life, but the acceleration hides inside a replacement cycle everyone treats as fixed. At AI-class hardware values, even a modest shortening of service life moves serious capital where some fraction of it is preventable.

Lost computing from throttling — Thermal penalties push accelerators into throttling, so the facility draws full power and delivers less work. For anyone selling or consuming GPU hours, that's purchased capacity not delivered — and it's almost never traced to environment, because throttling data and environmental data live in different systems (Section 8). Naming it connects a performance problem the operator already feels to a cause they hadn't considered.

Warranty and support exposure — the cheap protection. Where OEM terms require a documented environment, contamination damage in an unmonitored one can be contested. This is contingent rather than continuous, but it reframes monitoring from a cost into an insurance policy the facility holds the paperwork for.

17.3 Anchor the case to standards

Frame every target as a recognized standard. This does more than add credibility; it changes the *type* of conversation. A case built on "the vendor says" invites a vendor's skepticism. A case built on "the OEM's own installation requirements specify" and "the recognized industry standard is" is an engineering conversation, and engineers approve engineering cases. A measured baseline (Section 8) turns the whole proposal from a request into a gap analysis: here is the standard, here is where we measure against it, here is what closes the gap. It is worth making that case well, because contamination-control programs are almost always won or lost in the skeptical engineering review.

18. Sustainability accounting done honestly

The chronic-cost framing has a direct environmental expression: energy not wasted and hardware not prematurely manufactured are both carbon lines. Contamination control participates in a facility's sustainability story through two channels — the operational-energy channel (reduced cooling energy is reduced Scope 2 emissions under the GHG Protocol) and

the embodied-carbon channel (extended hardware life amortizes each device's manufacturing footprint over more service, assessable under ISO 14044 life-cycle methods).

The claim discipline matters more here than anywhere else in the guide, because sustainability claims attract scrutiny and greenwash accusations stick. The scope-honest posture: quantify only measured effects; attribute floor-level benefits to floor-level mechanisms (particulate load, filter duty, reservoir condition) rather than to whole-facility outcomes; state plainly that contamination control complements rather than replaces airside and cooling-plant efficiency measures; and let PUE trends, filter-consumption records, and fleet-life data carry the argument.

19 Implementation roadmap: the first 90 days and the first year

A pragmatic sequence for a new build, a retrofit, or a high-density deployment within an existing facility. The sequence is deliberately measured first: it produces a defensible baseline before it asks for control spending, which makes every subsequent request easier.

19.1 Phase 0 — Design and planning (before construction, and up to ~90 days before turnover)

The highest-leverage contamination decisions are made before any measurement is taken. A control designed at this stage costs a fraction of the same control retrofitted into a running hall. Skipping this phase doesn't remove its costs; it converts them into expensive rework later.

Whatever the site, contamination-control planning at this stage has four jobs:

Design control into the layout. Entry and transition zones sized for threshold capture, logistics routes separated from personnel routes, a de-boxing and buffer zone ahead of the white space, and for raised floors; a plan for void access and cleaning that doesn't fight the cooling design.

Set the environmental targets and the monitoring that will prove them. Decide the cleanliness and corrosivity targets the facility will be held to (typically ISO 14644-1 Class 8 and ISA-71.04 G1) and specify the baseline and ongoing monitoring that will demonstrate compliance at turnover and beyond. Targets without a measurement plan are aspirations; targets with one are acceptance criteria.

Plan the construction and commissioning sequence for cleanliness. Decide how dust-generating work is phased and contained, how voids get cleaned and inspected before they're sealed, and how cleanliness verification becomes a handover deliverable rather than an afterthought (Section 13). For contamination, the dirtiest days of a facility's life are the days it is being built.

Assign ownership and write it into the contracts. Name who owns the contamination program, who signs off acceptance at handover, and how that connects to the commissioning authority. Put cleanliness acceptance criteria, contractor cleanliness obligations, and monitoring deliverables into the specifications and contracts. A program with no named owner and no contractual teeth is one that lapses the first busy week.

What differs by facility type. For a new build, this phase is where cleanliness is established. For a live retrofit, the biggest planning question is containment: how the work zone will be isolated, pressurized, and routed so construction contamination never reaches the operating hall. For a high-density deployment into an existing room, planning focuses on what the new load changes whether it is added airflow, added heat, or added liquid cooling and whether the existing environment and monitoring were ever validated against hardware this sensitive.

19.2 DAYS 0–30 · BASELINE

- Deploy copper/silver coupons and particle counters walkthrough; begin humidity trending against deliquescent thresholds.
- Inspect and photograph reservoirs: subfloor plenum, ceiling voids, cable trays, tile undersides (with zinc-whisker assessment in legacy facilities).
- For liquid-cooled capacity: record baseline coolant conductivity, flow, and filter pressure differentials; confirm filtration specification against OEM guidance.
- Audit thresholds and traffic: map all personnel and logistics entry to the hall and note current packaging discipline.
- Score the facility against the Appendix A self-assessment.

19.3 DAYS 30–90 · DESIGN AND CLOSE THE URGENT GAPS

- Read coupon results; classify the site and set the filtration decision.
- Install retentive floor-level capture at every threshold identified in the traffic audit; establish the no-cardboard line and the buffer-zone protocol.
- Remediate loaded reservoirs before any planned floor disturbance; fix filter seating, penetration sealing, and pressurization deficiencies found in the baseline.
- Qualify (or re-qualify) the cleaning contractor against Section 14.2; put void cleaning into the contracted scope.
- Draft the incident playbook (Section 15) while the baseline data is fresh.

19.4 MONTHS 3–12 · OPERATE, VERIFY, INSTITUTIONALIZE

- Run the monitoring program continuously where drift matters (particulate, reactivity, coolant conductivity); spot-check around construction and commissioning events.

- Wet-clean retentive flooring on the schedule its loading demands; clean reservoirs on cycle, not on discovery.
- Maintain the audit trail: testing records and certificates of cleanliness for warranty, insurance, and sustainability reporting.
- After two full quarters of data, replace the placeholder percentages with measured values and re-present the business case with your own numbers.
- Review the program each hardware generation — density, cooling, and OEM environmental requirements are all moving, and a program validated against last generation's hardware has not been validated at all.

FROM THE PUBLISHER

20. Where Dycem Fits

This guide has been vendor-neutral by design. Dycem provides floor-level contamination control for entrances, transitions, and traffic routes within data center environments. It is intended to complement air filtration, controlled cleaning, material-handling procedures, ESD controls, and environmental monitoring as part of a layered contamination-control program.

20.1 The Contamination Pathway Dycem Addresses

Air filtration is an essential component of data center environmental control, but it does not capture all particulate entering the facility. Contamination may also be carried ac into a facility on:

- Footwear
- Cart and trolley wheels
- Toolboxes and service equipment
- Pallet trucks and material-handling equipment
- Equipment casters
- Construction equipment
- Packaging and delivery materials

These sources may introduce particulate downstream of the air-handling and filtration system. Once deposited inside the facility, particulate may be redistributed through personnel movement, equipment movement, cleaning activities, and airflow.

20.2 Published Performance Claims

The figures below are based on Dycem-published testing under defined conditions. Actual performance will vary according to product selection, installation size, traffic levels, contamination load, cleaning frequency, and site conditions.

Claim	Figure	Basis
Foot- and wheel-borne particle capture	Up to 99.9%	Dycem published, independently tested
Reduction of airborne particulate	Up to 75%	Dycem published, independently tested

Claim	Figure	Basis
Static dissipation	**10^8 ohms**	Supports ESD-controlled zones per ANSI/ESD S20.20 program design
Antimicrobial performance	**BioMaster additive**	Helps inhibit microbial growth on the flooring surface
Comparison point: peel-off tacky mats	**<30% capture**	Published comparison testing; performance degrades further as mats load

20.3 Product Selection

Dycem is available in four formats designed for different project phases, traffic patterns, and installation requirements. Selection should consider whether the area is temporary or operational, how the traffic route may change, and whether access to raised-floor systems must be maintained.

Project phase	Entry and transition points	Staging and preparation rooms	Corridors and maintenance routes	Raised-floor zones
New construction	Dycem Lite	Dycem Lite	Dycem Lite	Access Panels
Commissioning and fit-out	Floating Mats	Floating Mats	Dycem Lite	Access Panels
Operational	Fixed Solution	Floating Mats	Floating Mats	Access Panels
Retrofit	Fixed Solution	Floating Mats	Floating Mats	Access Panels

Product Overview

Dycem solution	Description and application
Dycem Lite	A low-profile, reusable mat for fast deployment in high-traffic transition areas. Suitable for new construction, fit-out, server rooms, data halls, airlocks, staging areas, and corridors where a flexible solution is required.
Floating Mats	Portable, freestanding contamination control for temporary entry points, preparation areas, maintenance routes, and locations where traffic patterns or protection requirements may change.

Access Panels	Contamination-control material inlaid into raised-floor panels, providing long-term floor-level protection while preserving access to the underfloor plenum.
Fixed Solution	Permanently installed contamination control for high-traffic entrances and critical transition zones requiring continuous, long-term protection.

Product selection should also account for traffic frequency, wheel types and loads, available floor space, door clearance, cleaning access, ESD requirements, substrate condition, and the likelihood of personnel or equipment bypassing the controlled surface. Where necessary, a placement plan can be used to determine the appropriate product, location, and installation size.

Learn more about Dycem products here.

APPENDICES

The Toolkit

A. Facility self-assessment scorecard

Score each item 0 (absent), 1 (partial/informal), or 2 (in place, documented, verified). Twenty items; 40 points maximum. The value is less in the total than in the zeros — each zero is a named gap with a section of this guide attached to it.

#	Item	Score 0–2
1	Airborne particulate measured against ISO 14644-1 Class 8 within the last 12 months	
2	ISA-71.04 corrosivity classification (Cu + Ag) current for the site	
3	Continuous or scheduled particle monitoring with trend review	
4	Humidity trended against deliquescent-activation thresholds	
5	Retentive floor-level capture at every personnel threshold	
6	Heavy-duty capture at wheeled-logistics thresholds	
7	No-cardboard/packaging line enforced at the buffer zone	
8	Positive pressurization verified; doors and penetrations sealed	
9	Filter seating and change schedule verified (not just scheduled)	
10	Gas-phase filtration decision made from measured severity (installed, or documented as not required)	
11	ESD program per ANSI/ESD S20.20 covering hall and staging areas	
12	Subfloor/ceiling voids inspected and cleaned on a defined cycle	
13	Legacy zinc-plated surfaces assessed before floor disturbance	
14	Qualified data-center cleaning contractor; void scope in contract	
15	Certificates of cleanliness and test records retained (audit trail)	
16	Coolant quality specified with monitored conductivity and filtration (liquid-cooled capacity)	
17	Loop commissioning records: flush, fill-water quality, baselines	
18	Construction/retrofit cleanliness plan with containment and monitoring	
19	Contamination incident playbook drafted and known to the team	

#	Item	Score 0–2
20	Program reviewed against OEM environmental requirements at each hardware refresh	

Reading the total:

32–40 — mature program; focus on trend analysis and per-generation review.

20–31 — sound foundation with named gaps; sequence fixes via the Section 19 roadmap.

Below 20 — the facility is relying on luck and legacy thermal margin; start with the Days 0–30 baseline, because measurement is the cheapest item on the list and unlocks all the others.

Download the facility scorecard

B. Glossary

Term	Definition
ASHRAE TC 9.9	Technical committee that publishes the thermal and contamination guidance most data centers design to.
CDU	Coolant distribution unit — the interface between facility water and the rack-level liquid-cooling loop; typically hosts pumps, heat exchangers, filtration, and sensors.

Term	Definition
Creep corrosion	Migration of corrosion products across a surface (notably immersion-silver finishes) that can bridge adjacent conductors.
Deliquescent RH	Relative humidity above which a hygroscopic particle absorbs moisture and becomes an electrolyte, accelerating corrosion and leakage.
Direct-to-chip (DTC)	Liquid cooling delivered via cold plates mounted directly on processors and accelerators.
ESD	Electrostatic discharge - sudden transfer of accumulated static charge into electronics, causing immediate or latent damage; governed by ANSI/ESD S20.20.
ISA-71.04	Standard classifying environmental corrosivity by measured copper/silver reactivity into severity levels G1–GX.
ISO 14644-1	Cleanroom air-cleanliness classification by particle concentration; Class 8 is the common data-hall benchmark.
MERV	Minimum Efficiency Reporting Value - the filter-efficiency rating scale used for air-handling filtration.
PUE	Power Usage Effectiveness — ratio of total facility energy to IT energy; the headline efficiency metric contamination quietly taxes.
Resuspension	Re-entrainment of settled particulate into the air by traffic or airflow — the mechanism that makes settled contamination a recurring airborne load.
Van der Waals retention	Short-range molecular attraction by which polymeric contamination-control flooring holds captured particulate until wet-cleaned.
White space	The productive data-hall floor area housing IT equipment as distinct from support and plant space.
Zinc whiskers	Fine conductive filaments that grow from electroplated zinc surfaces; disturbed and airborne, they can bridge conductors.

View more data center definitions here.

C References and sources

Operators building a formal program should consult the current editions directly.

STANDARDS AND TECHNICAL GUIDANCE

- **ISO 14644-1:2015** — Cleanrooms and associated controlled environments, Part 1: Classification of air cleanliness by particle concentration. Current edition (reviewed and confirmed 2021); source of the Class 8 data-hall target. http://www.iso.org/standard/53394.html
- **ASHRAE TC 9.9 — Thermal Guidelines for Data Processing Environments, 5th ed. (2021)**, ASHRAE Datacom Series Book 1. The fifth edition adds research on high-humidity/gaseous-pollutant corrosion of IT equipment and a new high-density equipment class. ANSI webstore.
- **ASHRAE TC 9.9 — Gaseous and Particulate Contamination Guidelines for Data Centers** (white paper, 2009; rev. 2011). Established the Class 8 / modified ISA-71.04 G1 design pairing. As noted in Section 7.2, the white paper has drawn published criticism for claims stronger than its own references support — cite it as a widely accepted design target, not as settled failure-rate data. Link to the source.
- **ANSI/ISA-71.04-2013** — Environmental Conditions for Process Measurement and Control Systems: Airborne Contaminants. Classifies corrosivity G1 (mild) through GX (severe) by copper and silver reactivity; the 2013 revision added silver reactivity as a required metric and reflects RoHS-era changes. webstore.ansi.org
- **ANSI/ESD S20.20-2021** — Protection of Electrical and Electronic Parts, Assemblies and Equipment. The ESD control-program framework. esda.org
- **EN 50600 series** — Information technology — Data centre facilities and infrastructures. European design/operations standard family referenced in §7.5; environmental control expectations sit within its facility and operations parts. Link to source.
- **DCA — Data Centre Anti-Contamination Guide** (Data Centre Alliance). The operational program framework aligned with §10's zoned model and the program checklist. Link to Source.
- **GHG Protocol** (WRI/WBCSD Corporate Accounting and Reporting Standard) and **ISO 14044:2006** (Life cycle assessment — Requirements and guidelines).. ghgprotocol.org · iso.org/standard/38498.html

AI HARDWARE ROADMAP AND DENSITY

- **NVIDIA Newsroom — NVIDIA Vera Rubin platform** (16 March 2026). Primary source for platform architecture: the Vera Rubin NVL72 integrates 72 Rubin GPUs and 36 Vera CPUs and is a fully liquid-cooled, rack-scale system. nvidianews.nvidia.com
- **Rack-power figures (§2.1) are analyst estimates and vary by source and configuration.** A representative trajectory: ~40 kW (Hopper) → 120–140 kW (Blackwell GB200/GB300) → ~190–230 kW (Vera Rubin VR200 NVL72) → ~600 kW (2027 Rubin Ultra "Kyber") → megawatt-class beyond. Other coverage of the same NVL72 cites figures near 120 kW; the guide deliberately cites the range, not a point value. Lead with NVIDIA and SemiAnalysis; treat single-number blog figures as secondary.

CONTAMINATION MECHANISMS AND SPECIFIC TOPICS

- **Zinc whiskers — NASA NEPP / Goddard Space Flight Center metal-whisker program** (J. Brusse et al.), the recognized authority. Documents whisker growth on zinc-coated undersides of raised-floor tiles and shedding during maintenance. Balance with the industry counter-view (e.g., Legrand / Data Center Frontier) that the risk to modern electroplated products may be overstated and under-tested — citing both is the defensible position. nepp.nasa.gov/whisker
- **Creep corrosion, lead-free finishes, iNEMI** — After the RoHS-driven transition to lead-free PCB surface finishes, immersion-silver (ImAg) boards were found to creep-corrode in high-sulfur environments (ISA-71.04 Class G2 or higher), producing copper-sulfide bridging that short-circuits adjacent traces. The iNEMI Creep Corrosion Project — a neutral industry consortium (members including HP, Cisco, IBM, Dell, Intel, Huawei, Alcatel-Lucent, Dow Chemical) — published a 2012 position statement defining the temperature, humidity, and copper/silver corrosion-rate limits within which PCBs perform reliably, consistent with the 2011 ASHRAE white papers and anchored to modified ISA-71.04 severity level G1. Notably, humidity control alone is not sufficient: iNEMI testing found creep corrosion on some ImAg finishes at relative humidity as low as 15%, so protection depends on surface-finish selection and gas-phase filtration of sulfur-bearing contamination rather than humidity reduction. Finish behavior is not uniform — ImAg is most susceptible, ENIG showed creep corrosion only at high humidity (74–80%), and OSP suffered none across the tested range. Link to document
- **Liquid-cooling contamination** — for the §12 figures (sub-50 µm system-level filtration; conductivity ceilings near 100 µS/cm for deionized-water systems), reference NVIDIA GB200 system documentation and the Open Compute Project Cooling Environments work running alongside ASHRAE TC 9.9. opencompute.org

FLOOR-LEVEL CONTROL AND THE ENTRY VECTOR

- **Footwear and wheeled traffic as a contamination vector** — well established in cleanroom practice (IEST recommended practices and general contamination-control literature). iest.org
- **Van der Waals retentive polymer flooring** — the mechanism description is written to stand on its own terms; the quantitative performance figures in the closing section are Dycem's published, independently tested values, attributed as manufacturer claims. Dycem.com

D Disclaimers

This guide is provided for general educational and informational purposes only. It reflects the understanding of contamination control practices and data center industry developments as of its publication date and does not constitute engineering, legal, regulatory, financial, or professional advice for any specific facility, project, or circumstance.

Accuracy and Currency of Information. While reasonable efforts have been made to ensure the accuracy of the information presented, the data center industry is evolving rapidly. Statistics, technical specifications, performance figures, industry benchmarks, equipment capabilities, and other information referenced throughout this guide may change following publication. As a result, certain content may become outdated or no longer reflect current industry conditions, technologies, or best practices. Dycem does not guarantee that all information will remain accurate, complete, or current and assumes no obligation to update or revise this document. Readers are encouraged to independently verify figures, claims, and recommendations against the most recent available information before making decisions.

Site-Specific Considerations. Data center environments vary widely in design, location, equipment, and operating conditions. Nothing in this guide should be relied upon as a substitute for site-specific assessment, measurement, and professional engineering judgment. Any contamination control program, environmental target, control measure, or implementation decision should be validated against the facility's measured conditions, equipment manufacturer requirements, and the current editions of applicable standards.

Standards and Third-Party References. References to standards, including ISO 14644-1, ANSI/ISA-71.04, ASHRAE guidance, ANSI/ESD S20.20, and others, are provided for informational purposes only. Standards, regulations, and industry guidelines are periodically reviewed and revised. Readers should consult the latest official publications to confirm current requirements and recommendations. References to third-party organizations, products, hardware, or publications do not imply endorsement. All trademarks remain the property of their respective owners.

Forward-Looking and Estimated Information. Certain figures, including rack power densities, hardware roadmaps, energy consumption projections, and emerging liquid-cooling technologies, may be based on estimates, forecasts, publicly available research, or vendor-reported information. These figures are subject to change as technologies, market conditions, and industry practices develop. Forward-looking statements represent informed assessments at the time of publication and should not be interpreted as guarantees of future performance, adoption, or outcomes.

Performance Claims. Performance figures attributed to Dycem products reflect results obtained under defined testing methods and conditions and are provided for evaluation purposes. Actual performance may vary depending on installation, maintenance, traffic levels, environmental conditions, and other site-specific factors. Results are not guaranteed for any particular application. Independent test documentation is available upon request.

No Warranty; Limitation of Liability. This guide is provided "as is," without warranties of any kind, express or implied, including warranties of accuracy, completeness, merchantability, or fitness for a particular purpose. To the fullest extent permitted by applicable law, Dycem accepts no liability for any loss or damage arising from reliance on information contained in this guide, including information that may have become outdated or superseded since publication. Readers are responsible for verifying the accuracy, relevance, and applicability of the information to their own circumstances and ensuring compliance with all applicable laws, regulations, standards, and manufacturer requirements.