Power and Cooling for Ultra-High Density Racks and Blade Servers

White Paper #46

APC®
Legendary Reliability

Introduction

The Problem

- **Average rack in a typical data center is under 2 kW**
- **Dense deployment of blade servers (10-20 kW per rack) would greatly exceed the power and cooling ability of the typical data center**
- **Providing 10-20kW of cooling per rack is technically infeasible using conventional methods**

Introduction

The Solution

- **There are practical strategies for installing, powering, and cooling high density racks either singly or in groups**
- **Some of these strategies challenge prevailing industry thinking on high-density deployment**

Introduction

The Risk

- **Wrong choices in designing for high density can increase Total Cost of Ownership for NCPI *by many times***

Introduction

The Surprise

- **Extreme compaction of data centers (over 6 kW per rack) creates the need for extreme cooling infrastructure, which …**
  - ***Negates*** **the space savings of high-density IT equipment**
  - ***Increases*** **data center TCO**

High Density’s Challenge to Conventional Cooling

APC® Legendary Reliability

The Cooling Challenge

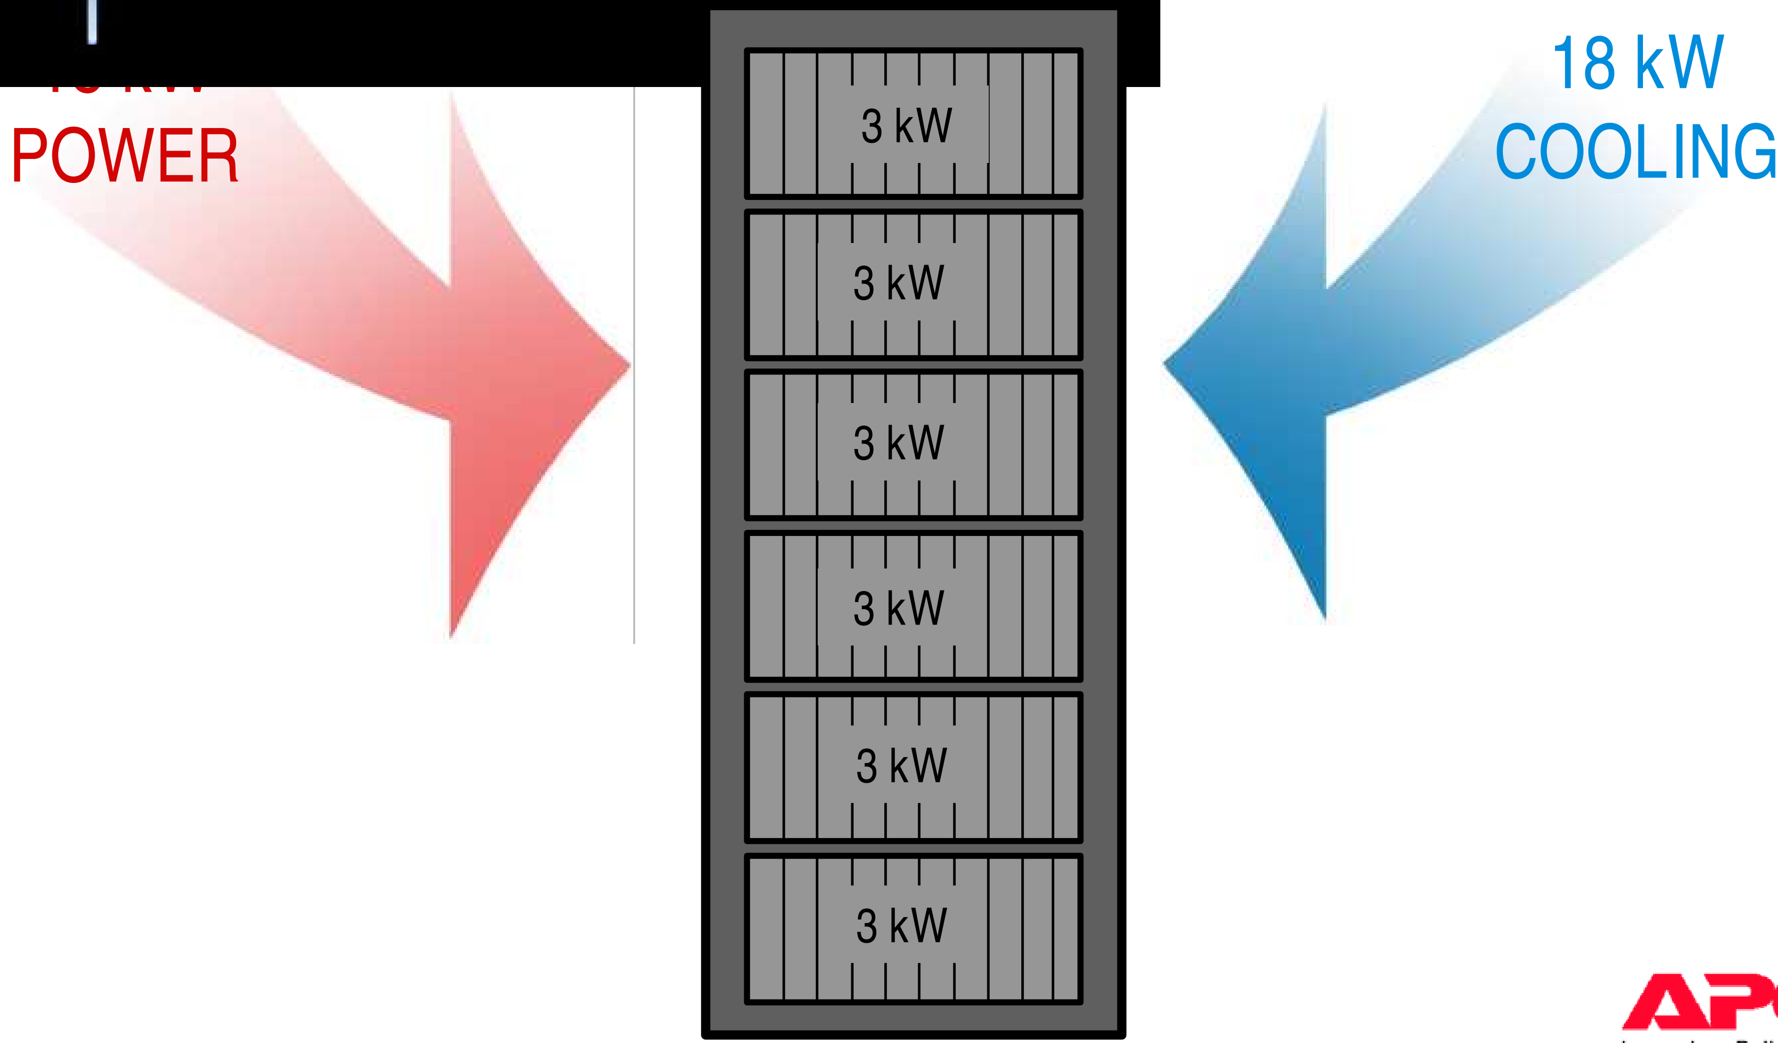

The Cooling Challenge

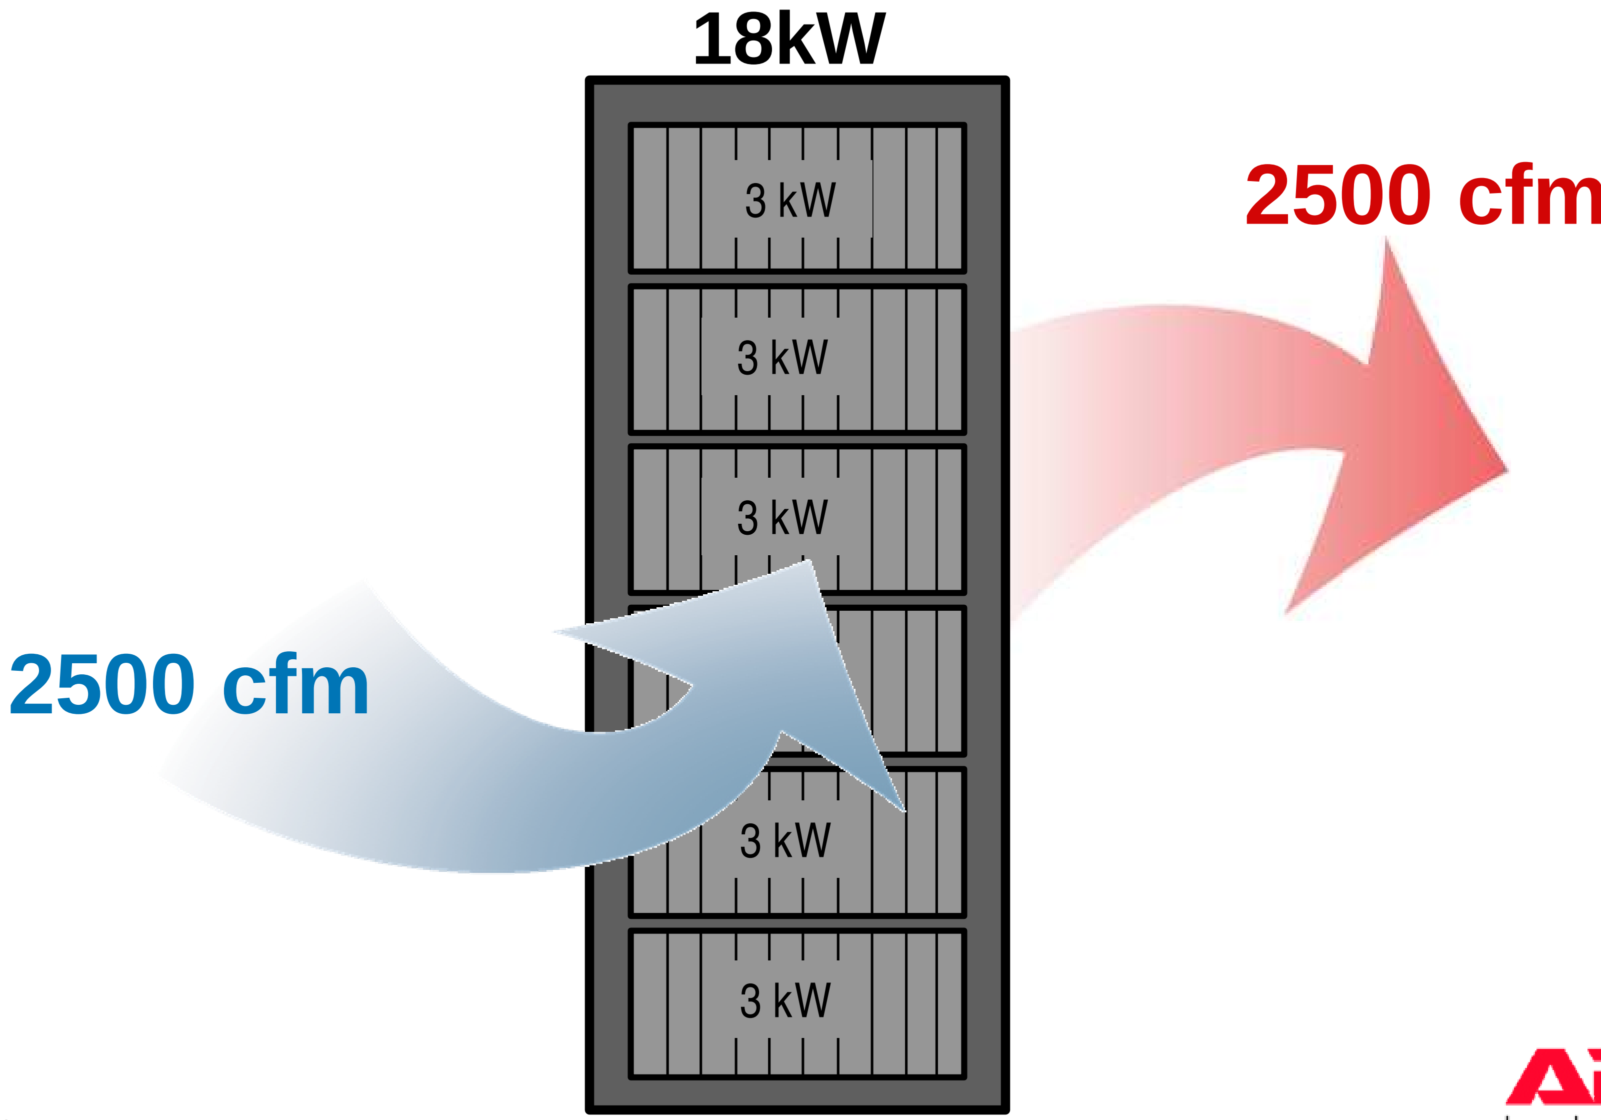

The Cooling Challenge

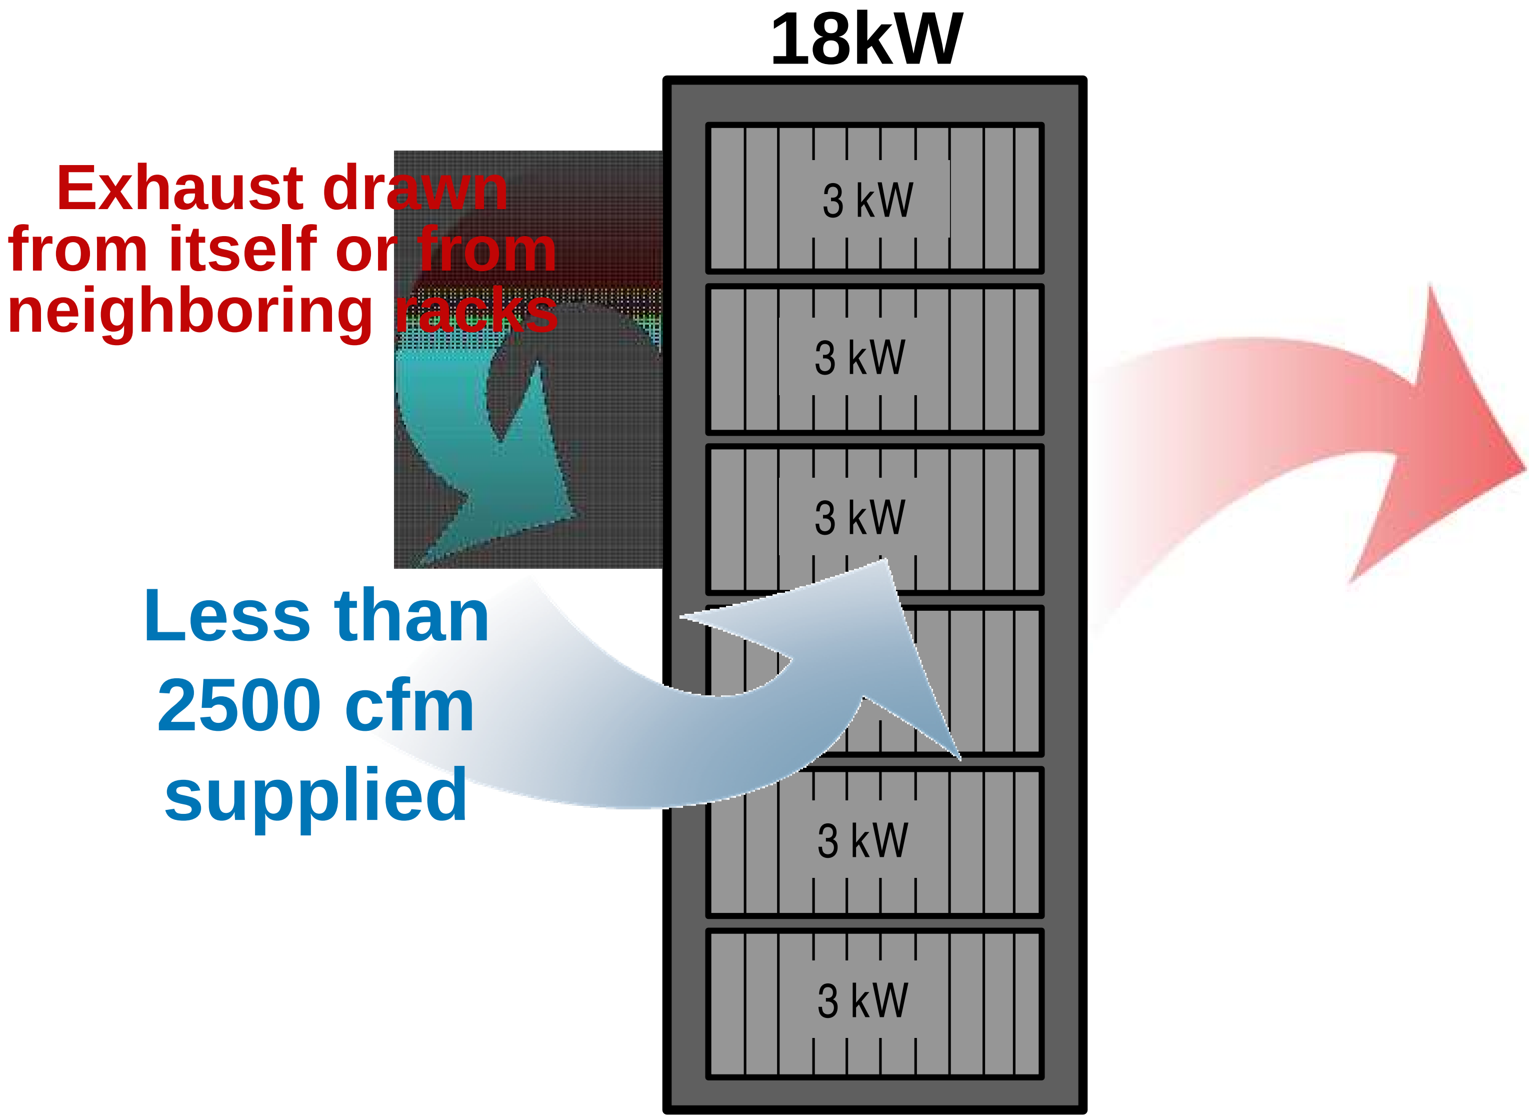

The Cooling Challenge:

(For this 18kW rack)

- **Supply 2500 cfm of cool air to the rack**
- **Remove 2500 cfm of hot exhaust air from the rack**
- **Keep the hot exhaust air away from the equipment intake**
- **Provide all these functions in a redundant and uninterrupted manner**

CHALLENGE # 1:

Supply 2500 cfm of Cool Air to the Rack

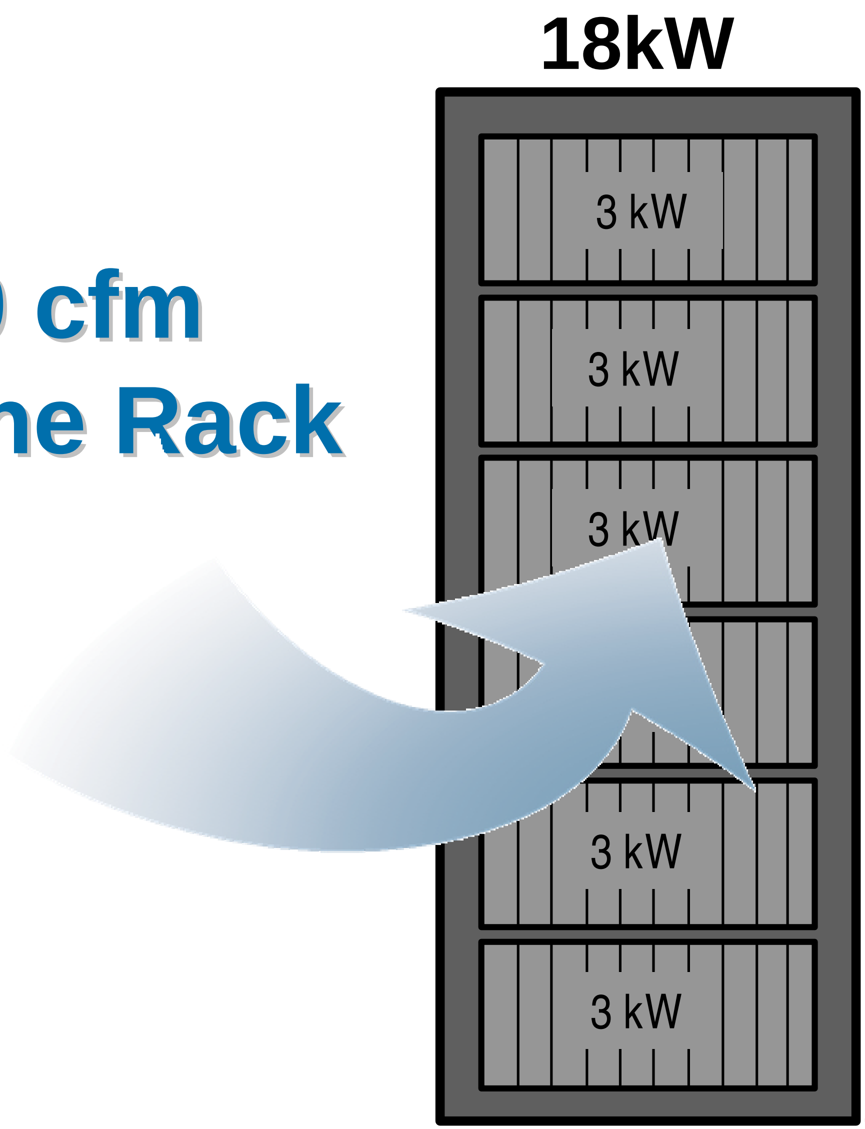

Typical Raised-Floor Airflow

One 300 cfm vented tile per rack

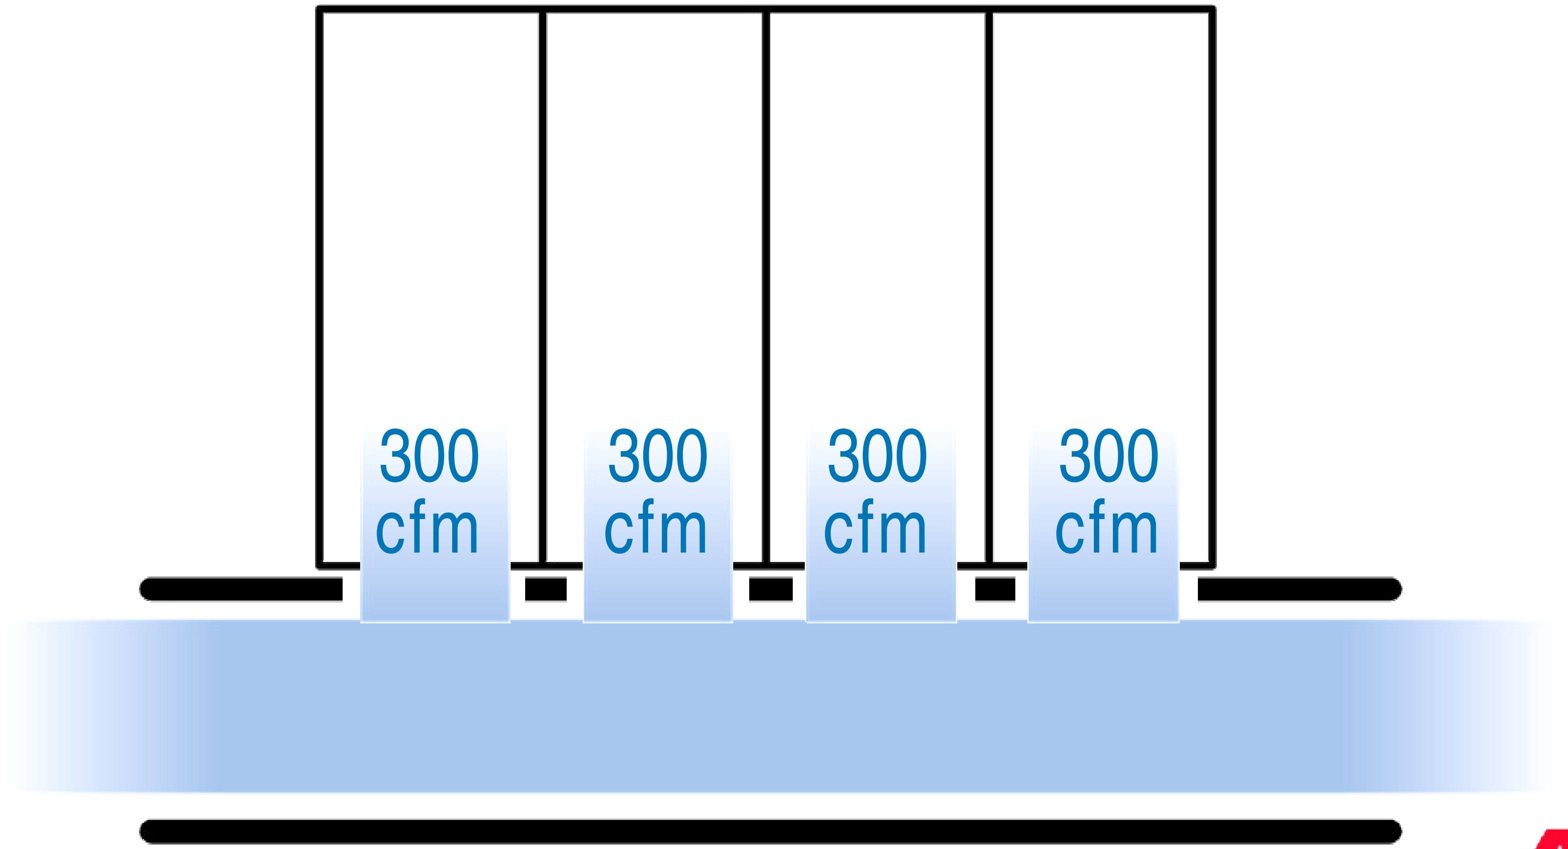

Perforated Tiles?

- **18kW rack would require 8 perforated tiles**
- **Aisle width would need to be substantially increased**
- **Spacing between racks would need to be substantially increased**

Perforated tiles cannot cool an 18 kW rack in a typical data center

APC® Legendary Reliability

Floor Tile Cooling Ability

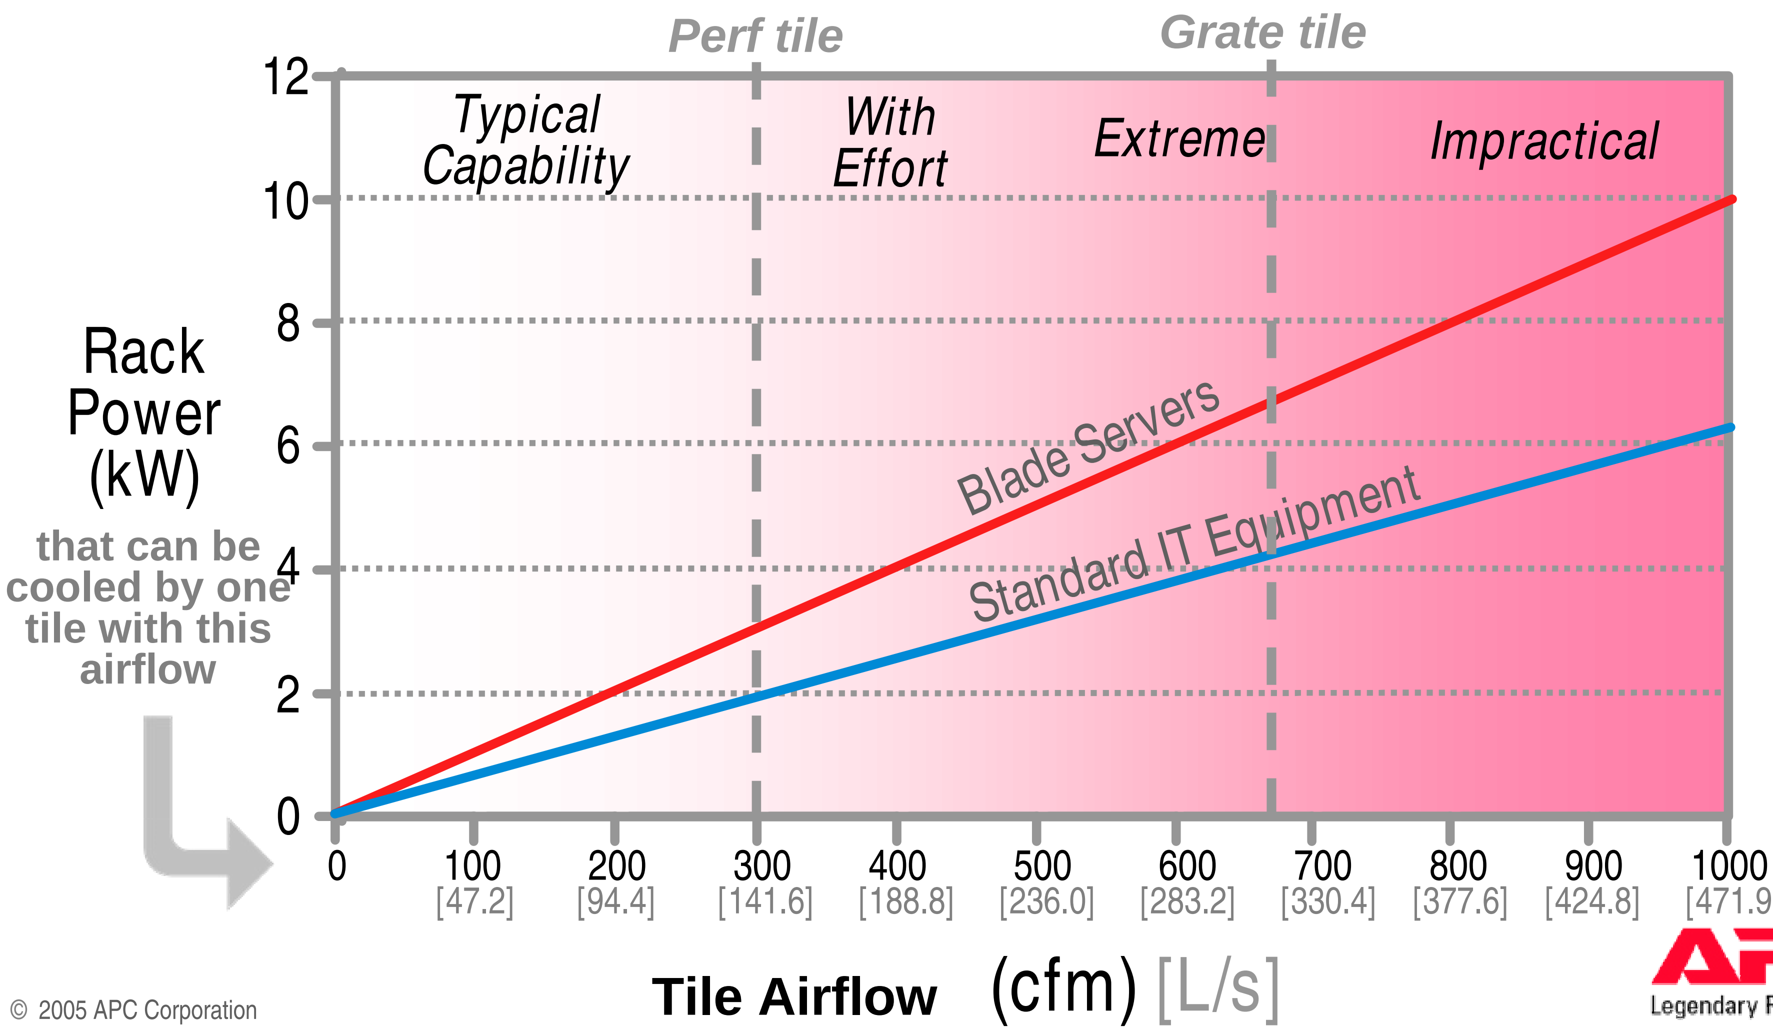

Legendary Reliability

Floor Tile Cooling Ability

Requires careful raised floor design, careful CRAC placement, and control of under-floor obstacles (pipes/wiring)

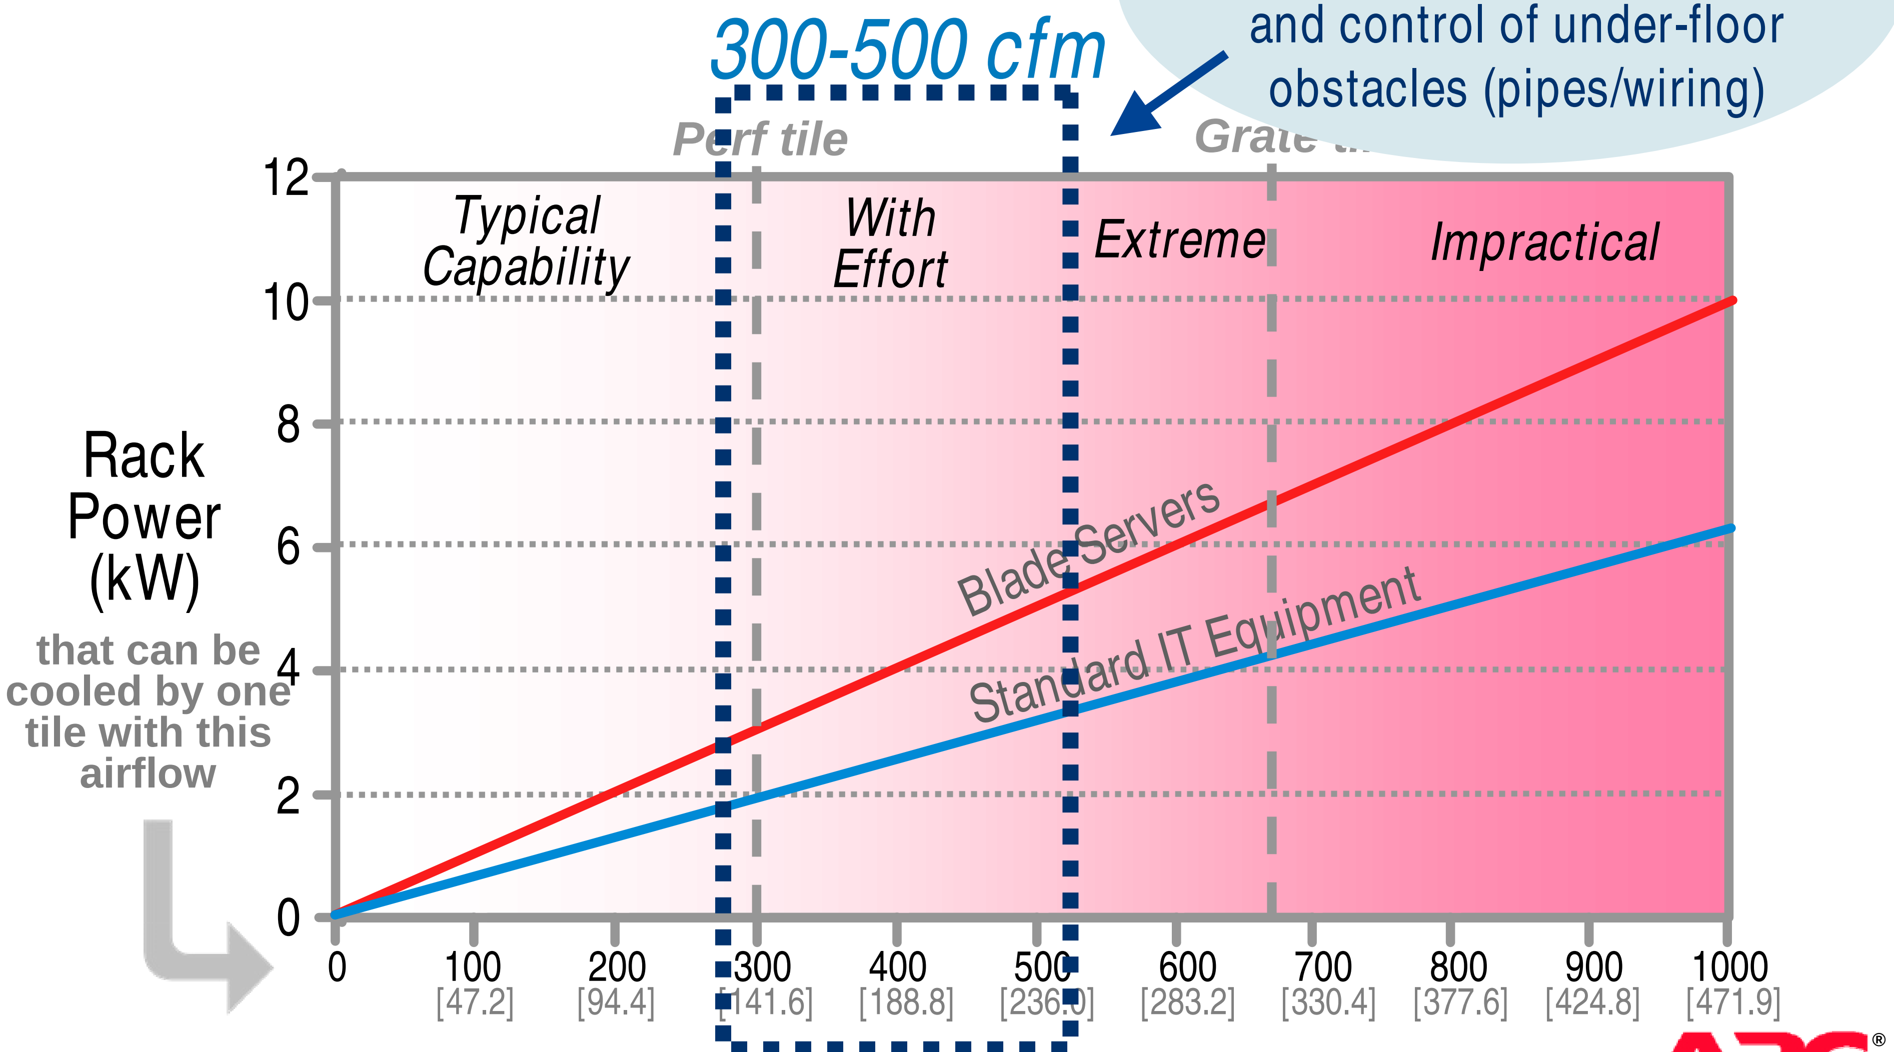

APC® Legendary Reliability

Floor Tile Cooling Ability

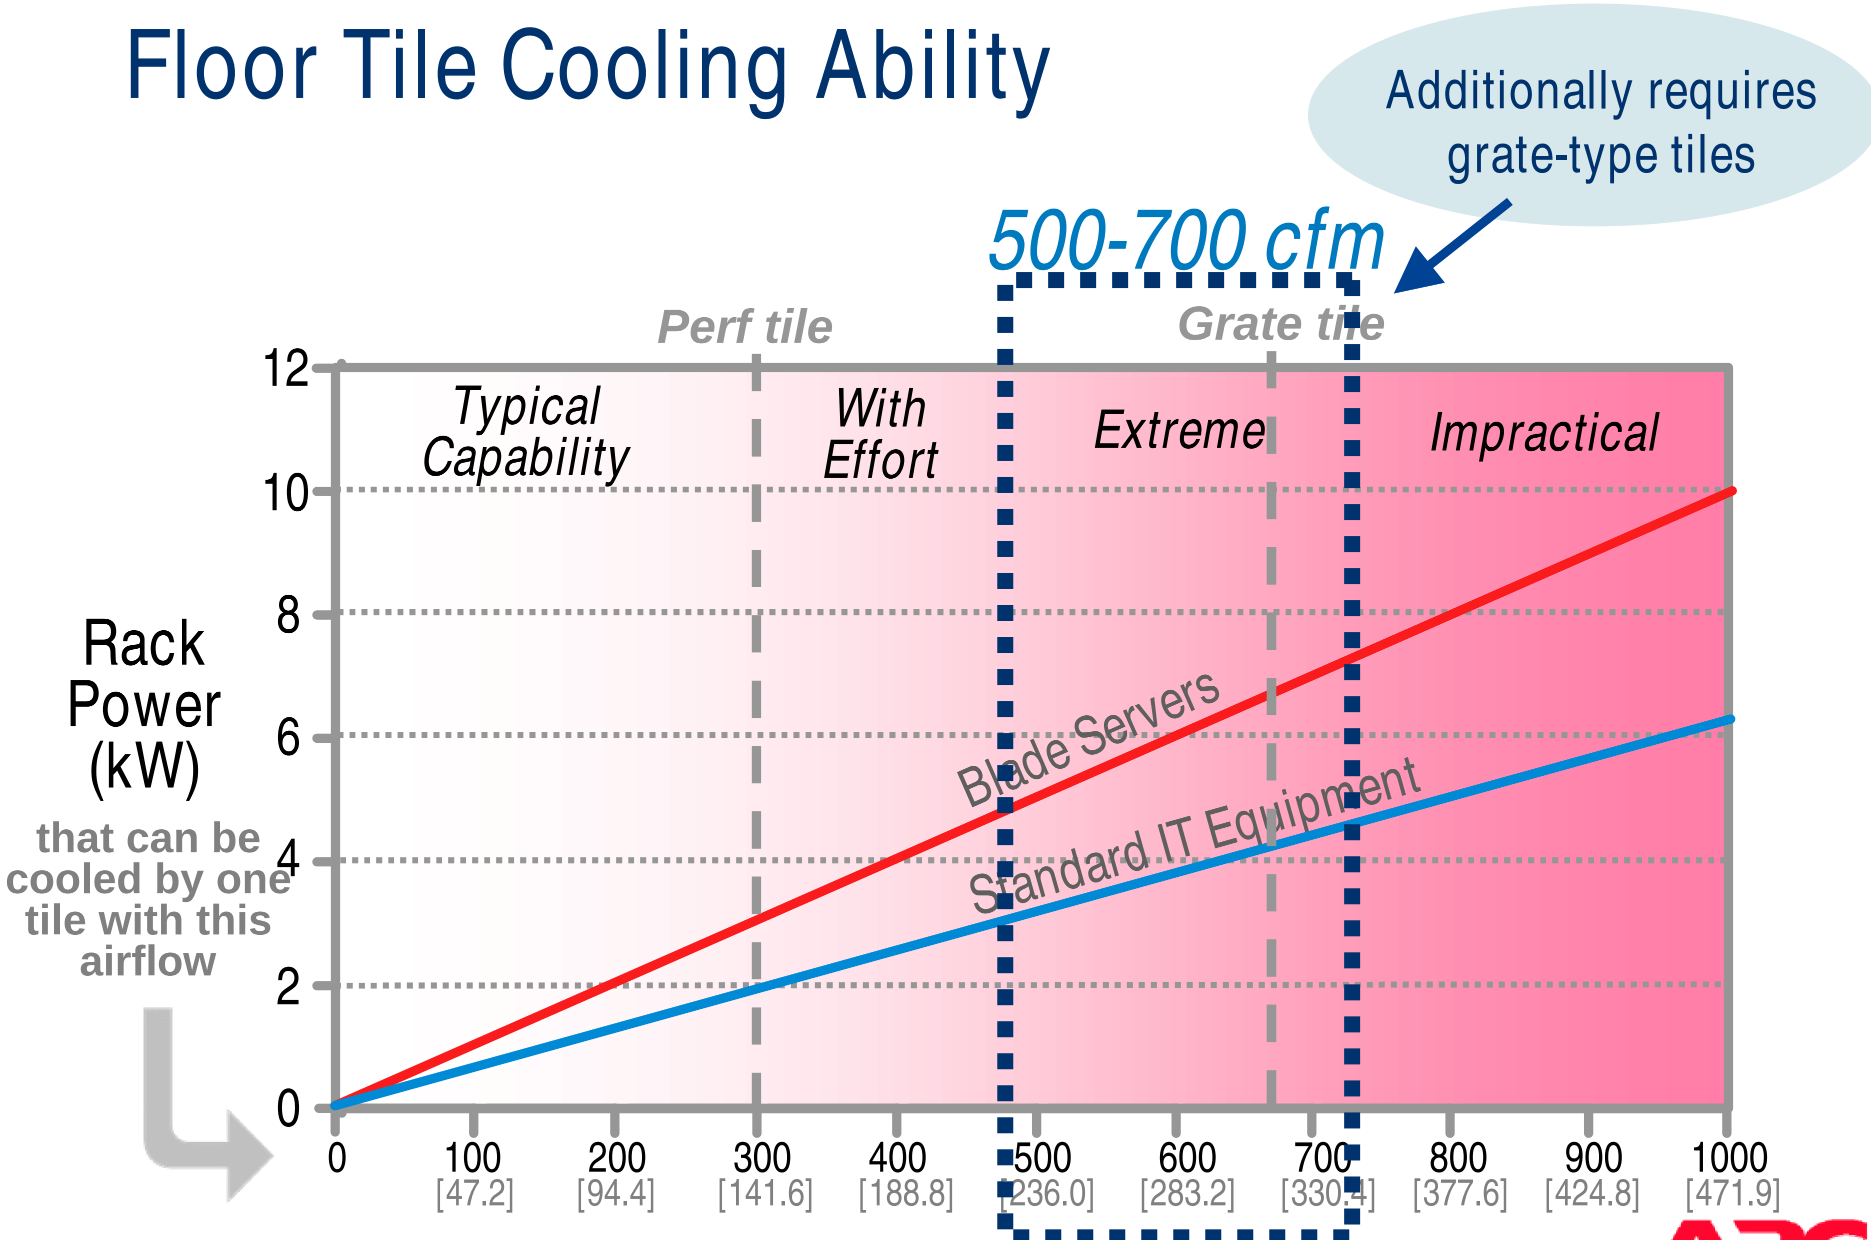

APC® Legendary Reliability

Increased Floor Depth?

Airflow variation decreases
as floor plenum depth increases

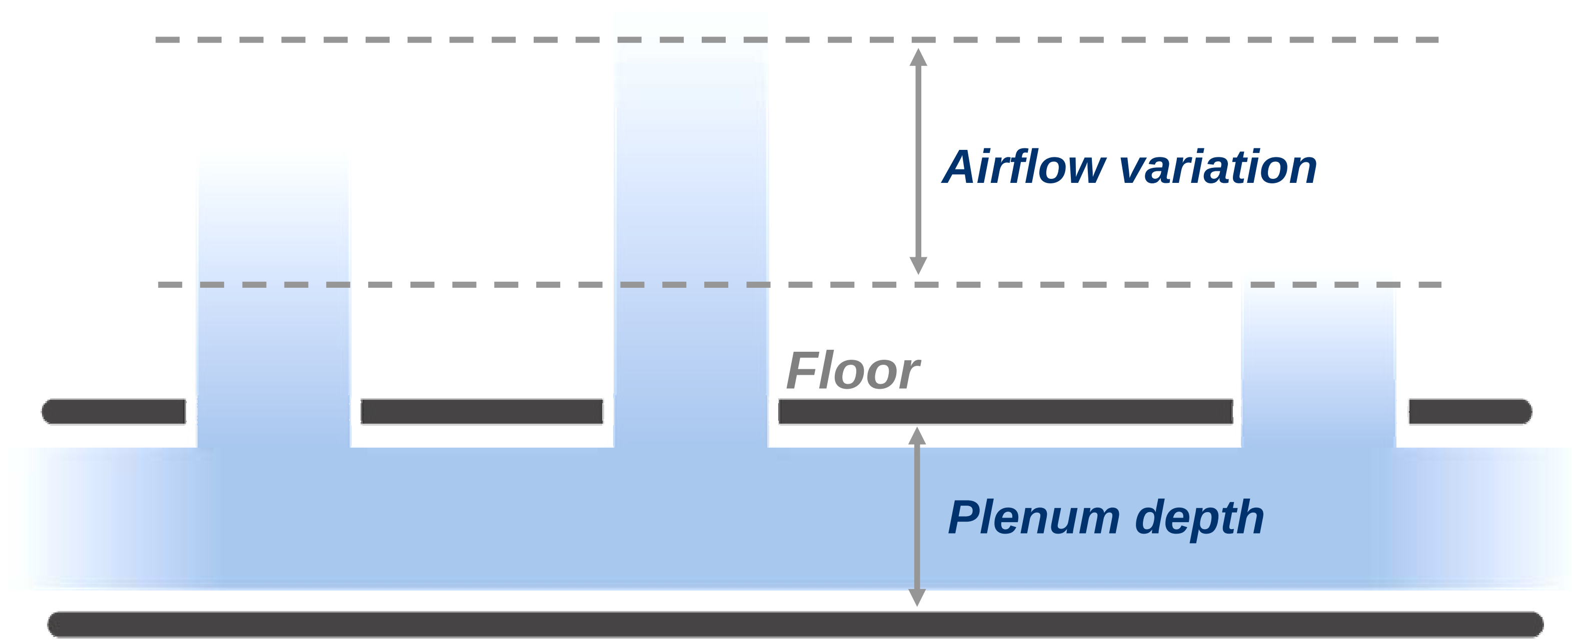

Increased Floor Depth?

Airflow variation decreases
as floor plenum depth increases

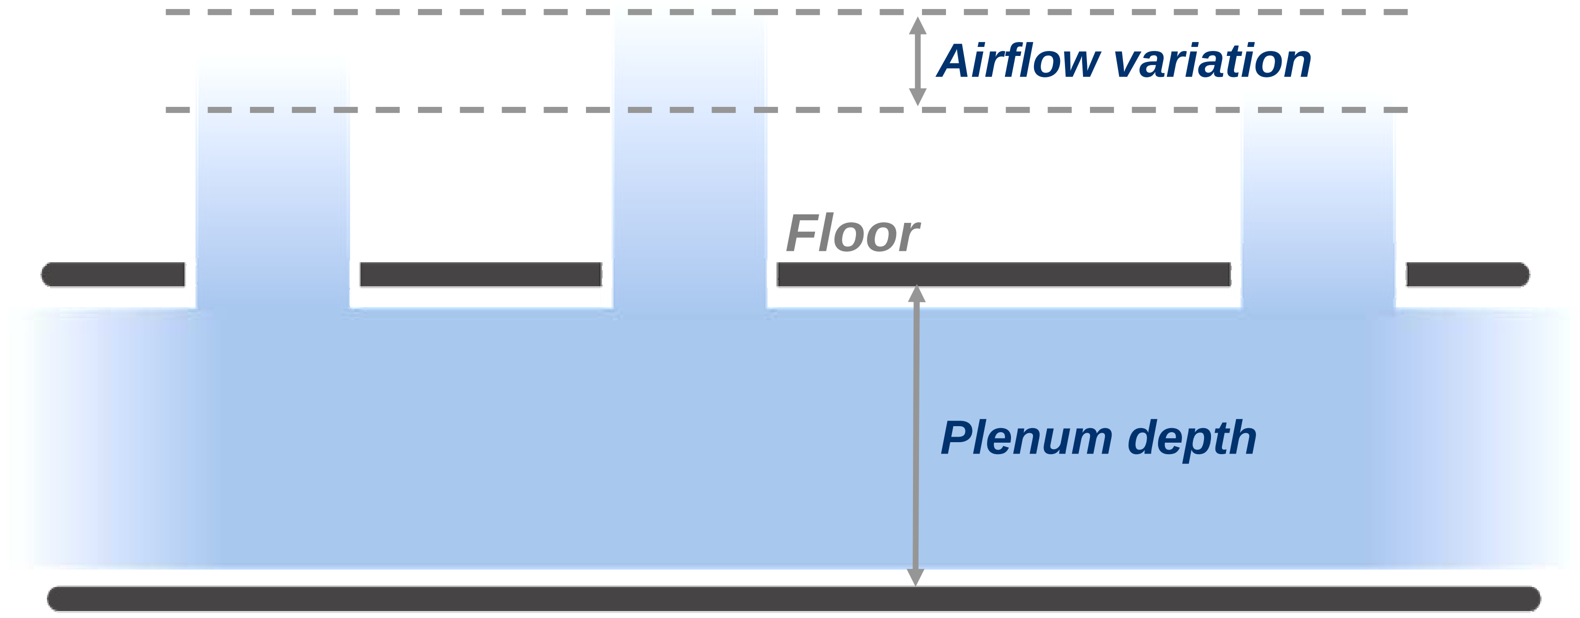

Increased Floor Depth?

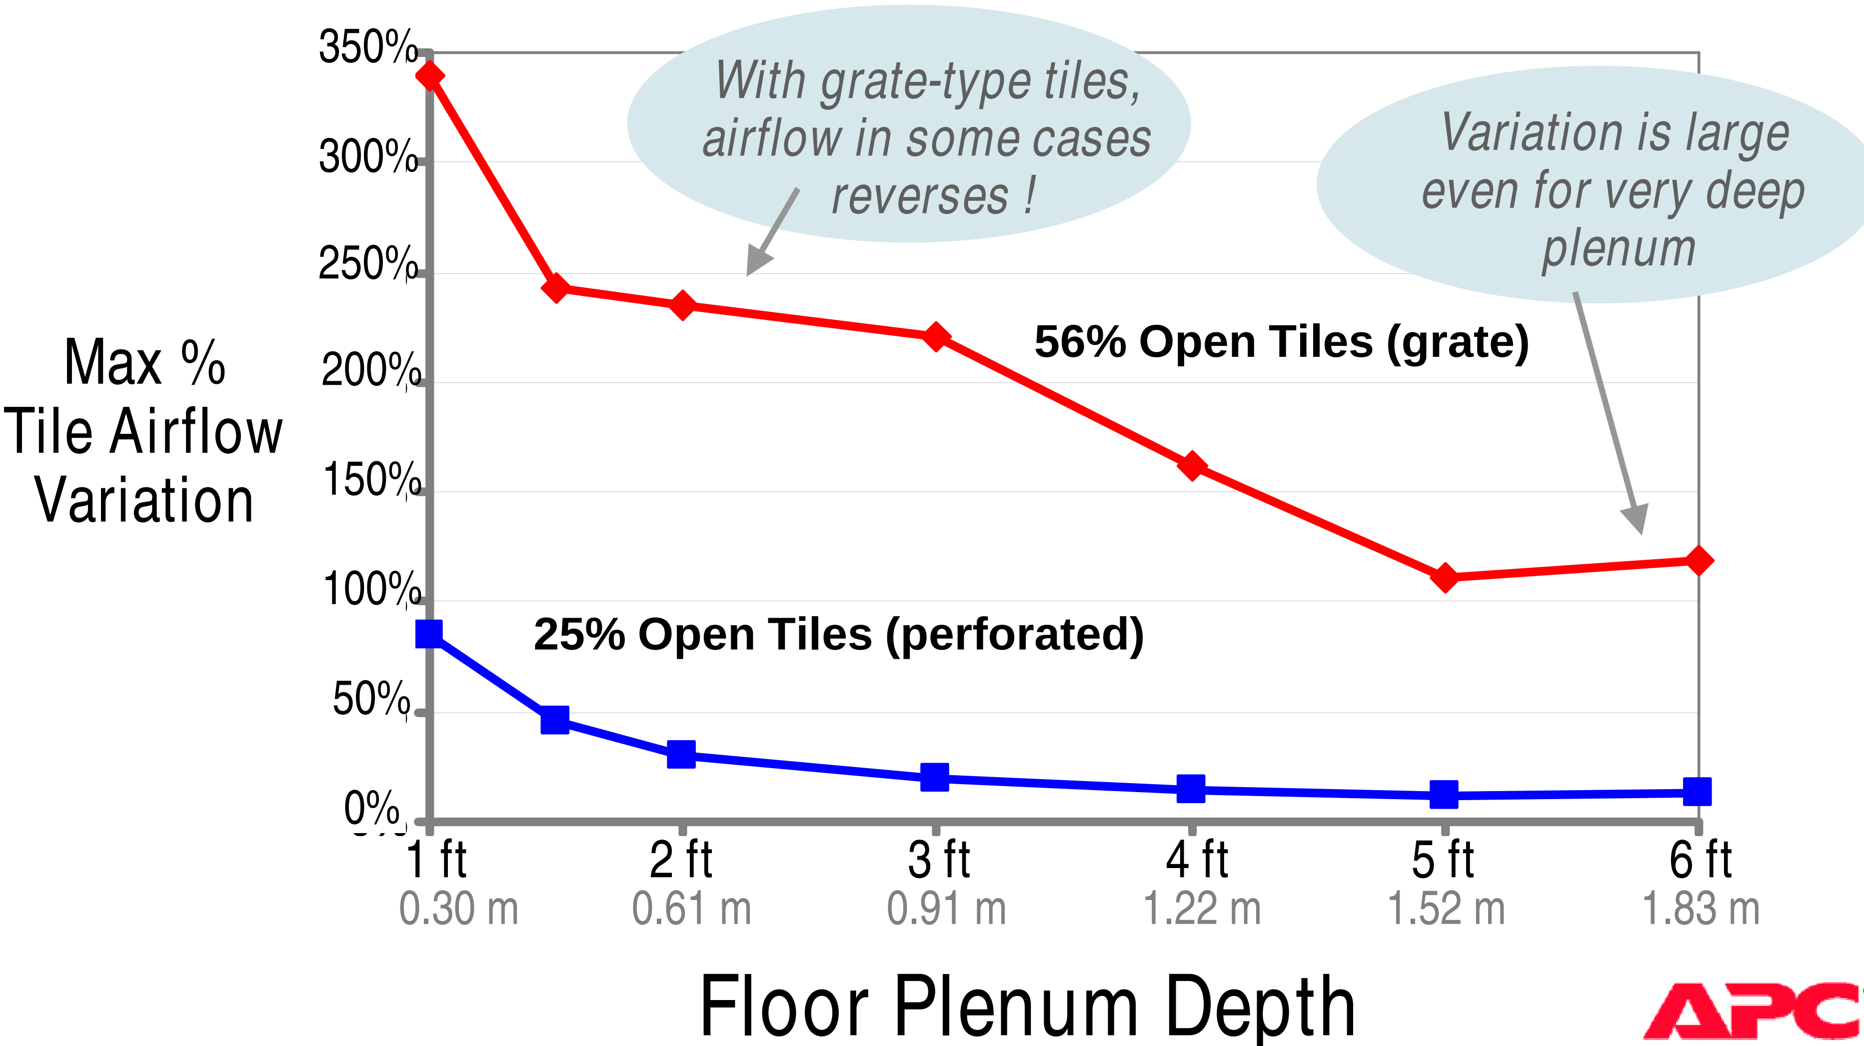

APC® Legendary Reliability

Grate-Type Tiles?

- **Grate-type tiles dramatically alter under-floor pressure gradients, making cooling non-uniform and unpredictable**
- **Grate-type tiles in one area impact airflow in neighboring areas**
- **Large airflow variations when using grate-type tiles mean some locations will NOT receive enough cooling**
- **Even if an "extreme" cooling design could solve these large airflow variations, it would still take 3-4 grate-type tiles to cool one 18kW rack**

APC® Legendary Reliability

Conventional Cooling Won’t Work

A conventional data center layout with one vented tile per rack simply cannot cool racks over approximately 6 kW per rack over a sustained area

CHALLENGE # 2:

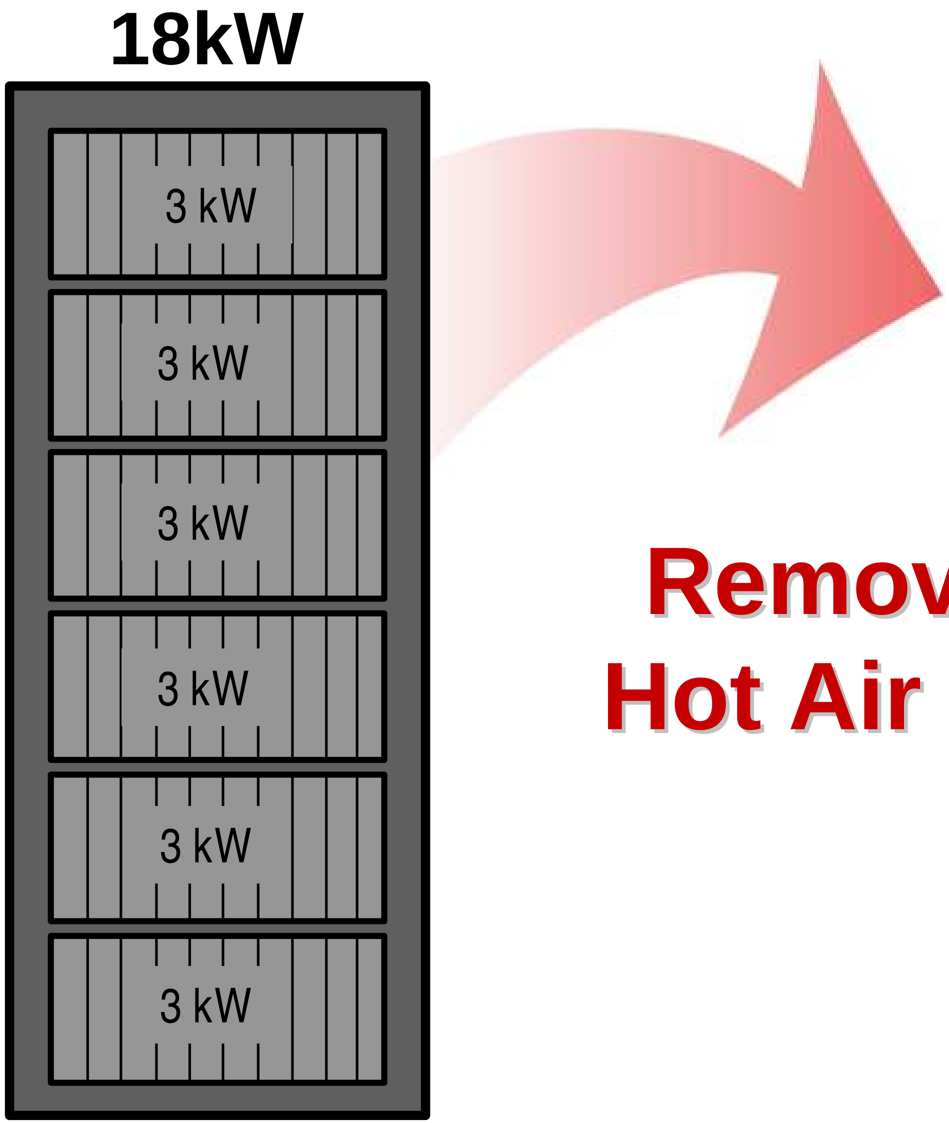

Remove 2500 cfm of Hot Air From the Rack

Removing Heat

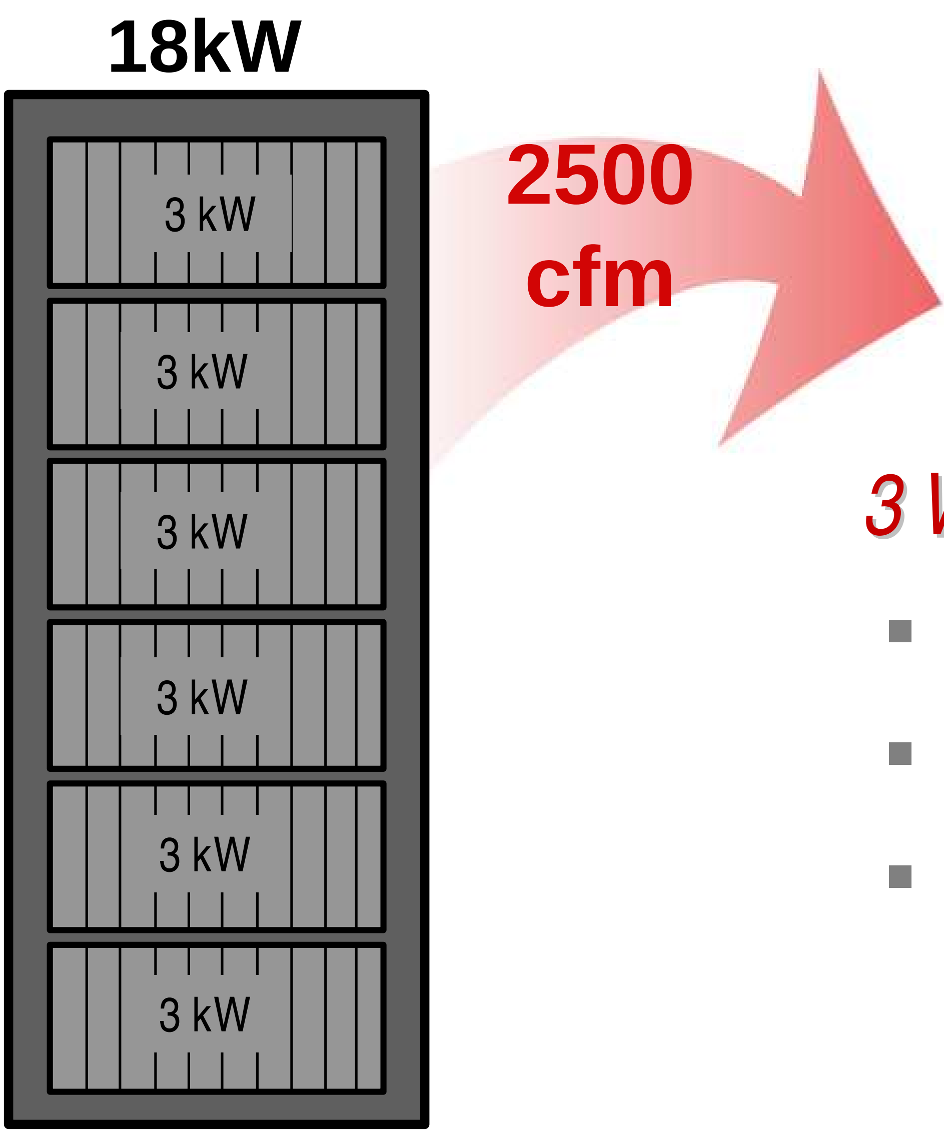

3 Ways to Remove Heat:

- **Through the room**
- **Through a duct**
- **Through ceiling plenum**

APC Legendary Reliability

Removing Heat

The ideal:

Hot exhaust air from equipment is taken directly back to the cooling system

- **Unobstructed, direct return path**
- **No chance to mix with surrounding air**
- **No chance to be drawn into equipment intakes**

2500 cfm through a 12" round duct goes 35 mph

1180 L/s through a 30 cm round duct goes 56 km/h

Removing Heat

Typical methods used in data centers

- **High, open ceiling with bulk air return at a central high point**
- **Return ductwork**
- **Suspended ceiling plenum**
- **Bulk air return across the room, under ceiling that is just a few feet above the racks**

These methods present high-density design challenges due to high air velocity

Removing Heat

More than 400 cfm (189 L/s) **airflow per rack – either supply or return – over a sustained area requires specialized engineering to ensure performance and redundancy**

CHALLENGE # 3:

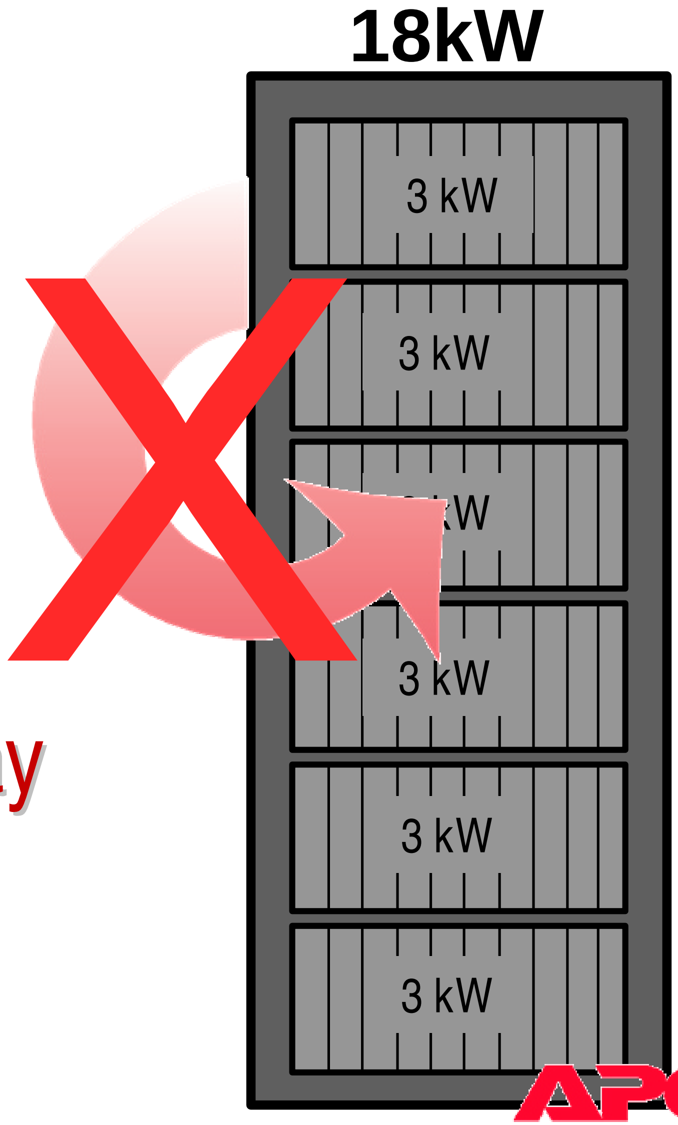

Keep Hot Exhaust Air Away From Equipment Intake

Preventing Recirculation

- **The shortest path for air to reach the equipment intake is recirculation from the equipment's own exhaust**

- **In high density environments, high airflow velocities become subject to resistance in ductwork, which degrades airflow patterns**

Supply and return paths must dominate airflow near potential recirculation paths to keep equipment from ingesting its own hot exhaust

Cooling the data center

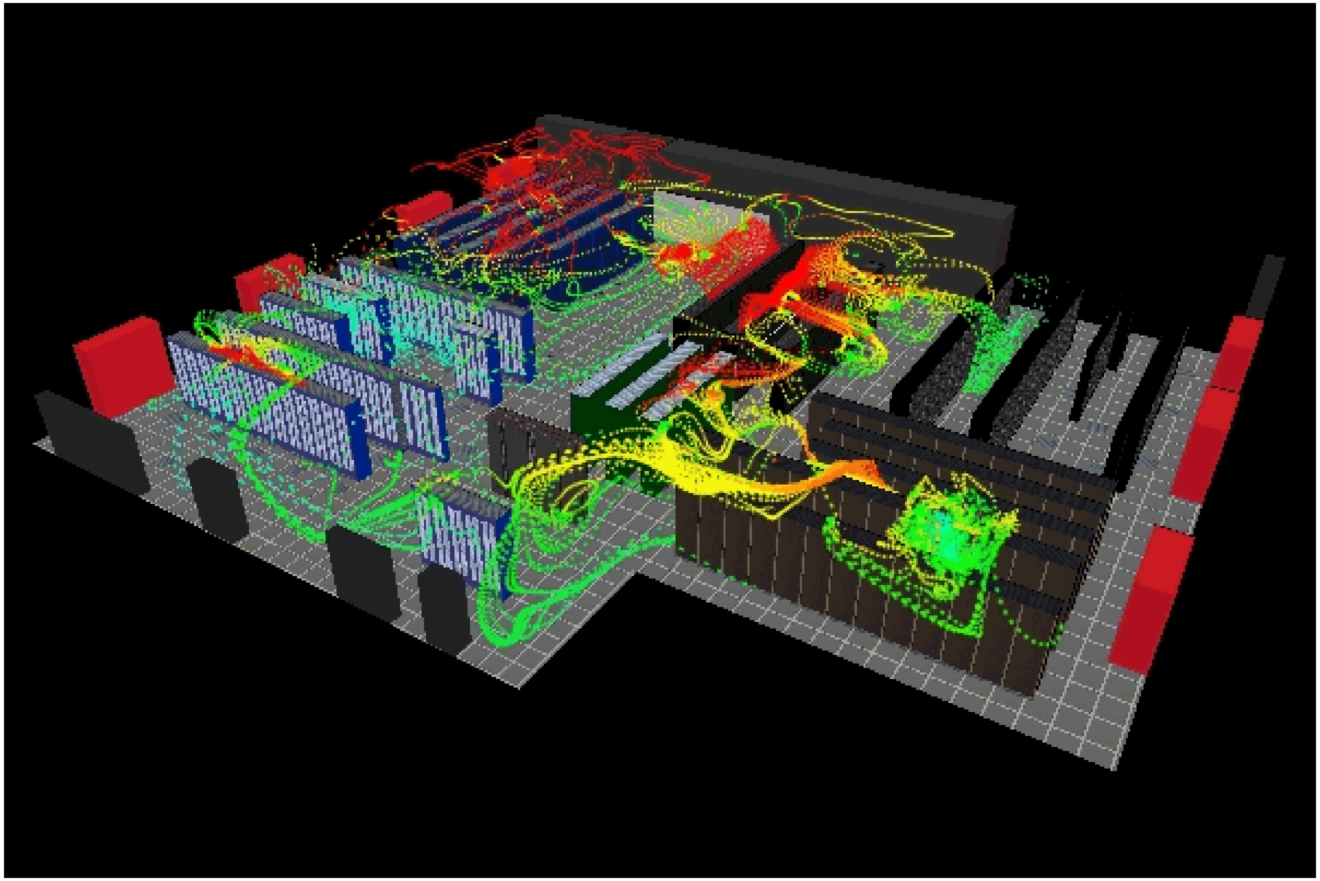

CHALLENGE # 4:

Provide Cooling in a Redundant and Uninterrupted Manner

Redundant

Cooling must continue during downtime of a CRAC unit

Uninterrupted

Cooling must continue during failover to generator backup

Redundant Cooling: Conventional Solution

- **Multiple CRAC units feed a shared raised floor or overhead plenum**
- **Plenum is assumed to sum all CRAC outputs and provide consistent pressure throughout**
- **System is designed to meet airflow and cooling requirements when any one CRAC unit is off**

Redundant Cooling: High Density Challenge

- **Airflow in cooling plenum increases**
- **Fundamental assumptions about shared-plenum system begin to break down**
- **With one CRAC unit off, local airflow velocities in the plenum are radically altered**
- **Airflow at an individual tile *may even reverse*, drawing air down into floor from venturi effect**

➡ *Under fault conditions, cooling operation becomes unpredictable*

Uninterrupted Cooling: Conventional Solution

- **Conventional system puts CRACS on generator, not UPS**
- **Temperature rise during 5-20 second generator startup is acceptable: 1°C (1.8°F)**

Uninterrupted Cooling:
High Density Challenge

- **In high density environment, temperature rise during uncooled 5-20 second generator startup can be 8-30℃ (14 -54℉)**
- **CRAC units may additionally have up to 5-minute "settle" time after power outage before they can be restarted**

With high power density, CRAC fans and pumps (CRAC units, in some cases) must be on the UPS to provide continuous cooling during generator startup

CRACs on UPS are a major cost driver and a major barrier to HD deployment

Five Strategies That Work

5 Strategies for Deploying High-Density Racks and Servers

1. **Design room for peak rack density**

Or design room BELOW peak rack density, and…

2. **Provide supplemental cooling for high-density racks**
3. **Establish rules for interspersed high-density racks to borrow cooling from adjacent racks**
4. **Spread out high-density equipment among multiple racks**
5. **Create a separate highly cooled area for high-density equipment**

Strategy #1
Design Room For Peak Rack Density

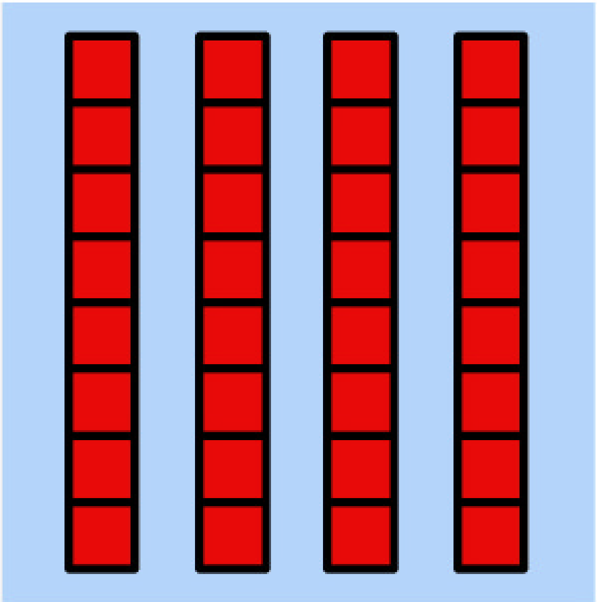

- **Handles wide range of future high-density scenarios**
- **Requires very complex analysis and engineering**
- **Capital and operating cost up to 4x alternative methods**
- **Risk of extreme underutilization of expensive infrastructure**
- **For rare and extreme cases of large farms of high-density equipment and limited space**

APC® Legendary Reliability

Strategy #2
Supplemental Cooling for HD Racks

Types of supplemental cooling:

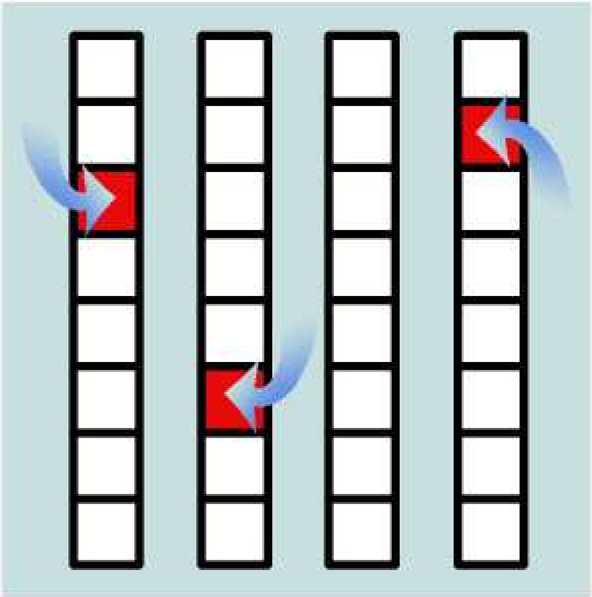

- **Specialty floor tiles or fans to boost cool air supply to rack**
- **Specialty return ducts or fans to scavenge hot exhaust from racks for return to the CRAC**
- **Special racks or rack-mounted cooling devices to provide cooling directly to the rack**

APC® Legendary Reliability

Supplemental-Cooling: Considerations

Strategy #2

- Provides high density when and where needed
- Defers capital cost of upgrading to high-density infrastructure
- High efficiency
- Optimal use of floor space
- Limited to about 10 kW per rack
- Room must be designed in advance to allow it

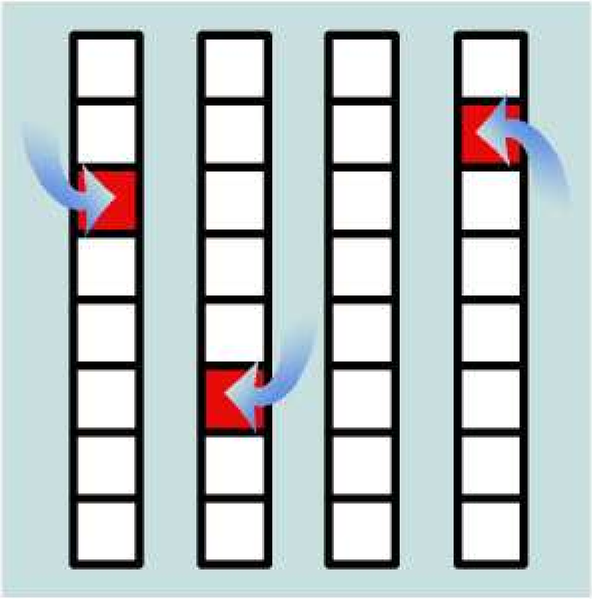

Supplemental-Cooling: Applications

Strategy #2

- **New construction or renovation**
- **Mixed environment**
- **When location of high-density equipment is not known in advance**

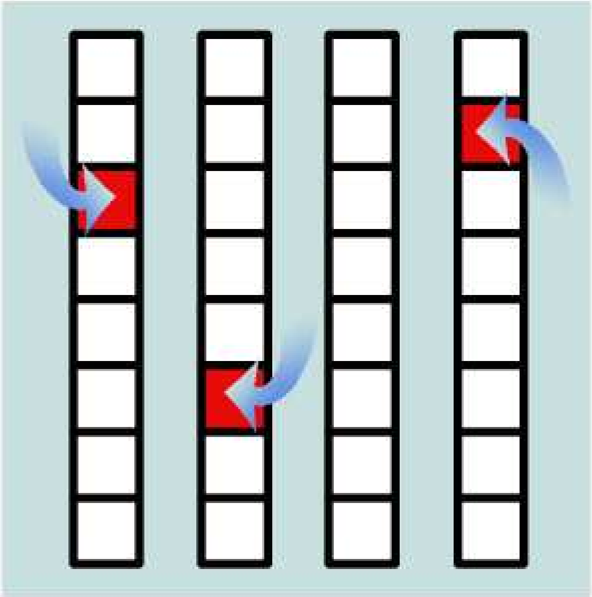

Rack Air Distribution Product Design Benefits

Dual Fans

Increases airflow from top to bottom of rack

A-B Power Input Feeds

Redundant power, Maximizes uptime

Filtered, Conditioned Air

Increases life of equipment by supplying cool clean air

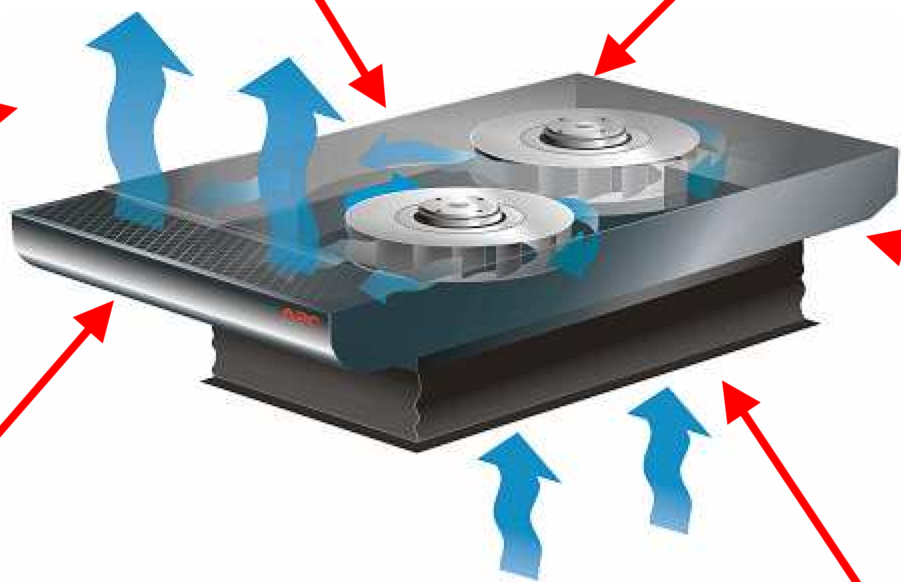

Independent Fan Control Switch

Vary amount of airflow to equipment

Air Filter

Removes airborne particles from the rack

Raised Floor Duct

Allows air to be pulled into the rack directly from the raised floor

Rack Air Distribution Unit

Rack Air Distribution Unit Airflow Diagram

- **Conditioned Room Air Is:**
  - **Pulled in from underneath raised floor**
  - **Delivered from bottom to top of rack by dual fans**
  - **Drawn in by the IT equipment**
- **Provides Cooler Air to the Rack**
  - **Provides better cooling for IT equipment reducing heat related failures**
  - **Extends the life of equipment in the rack**

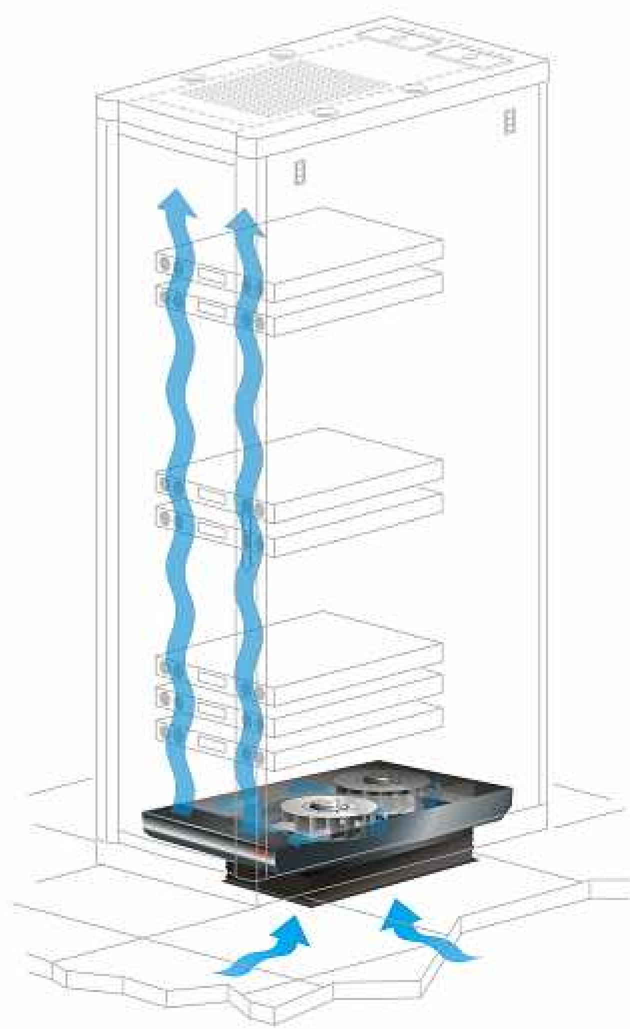

APC® Legendary Reliability

ARU AIR FLOW

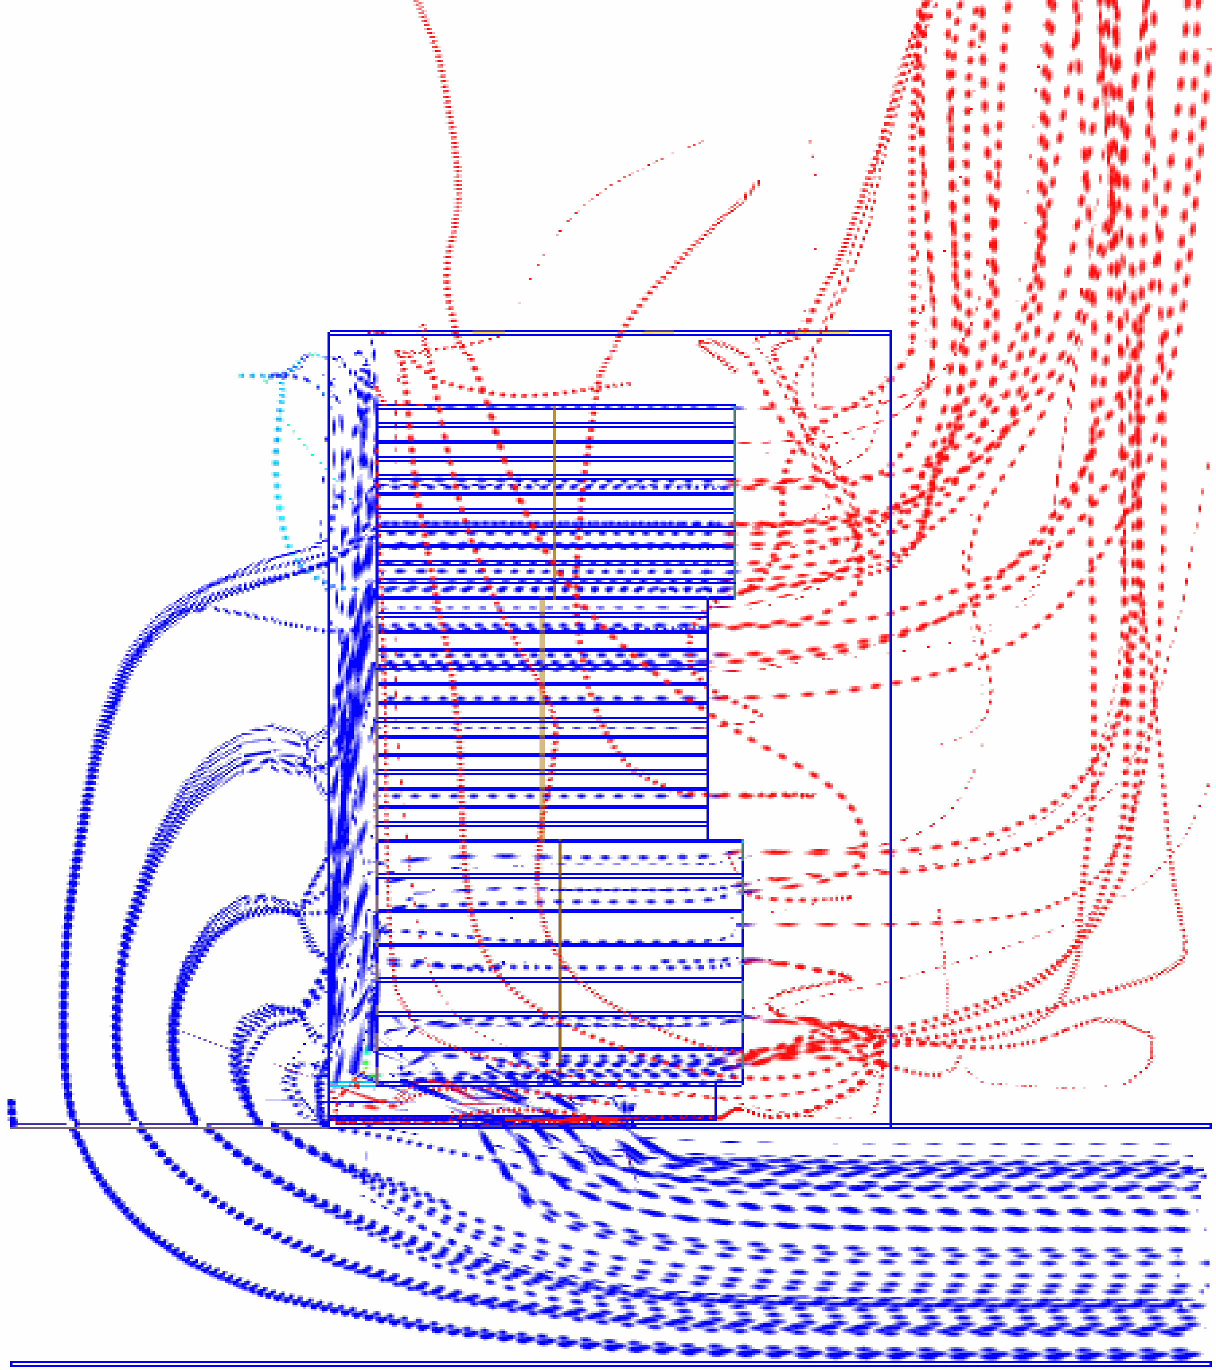

Rack Air Removal Unit Airflow Diagram

- **Fans pull in rack equipment exhaust air**
  - **Cable impedance is overcome by high powered fans**
- **Ducted exhaust system (optional) delivers hot air to plenum**
  - **Eliminates hot air from mixing with room air**
- **Proper airflow through the enclosure is ensured**
  - **Cool inlet air moves freely to equipment in the rack**

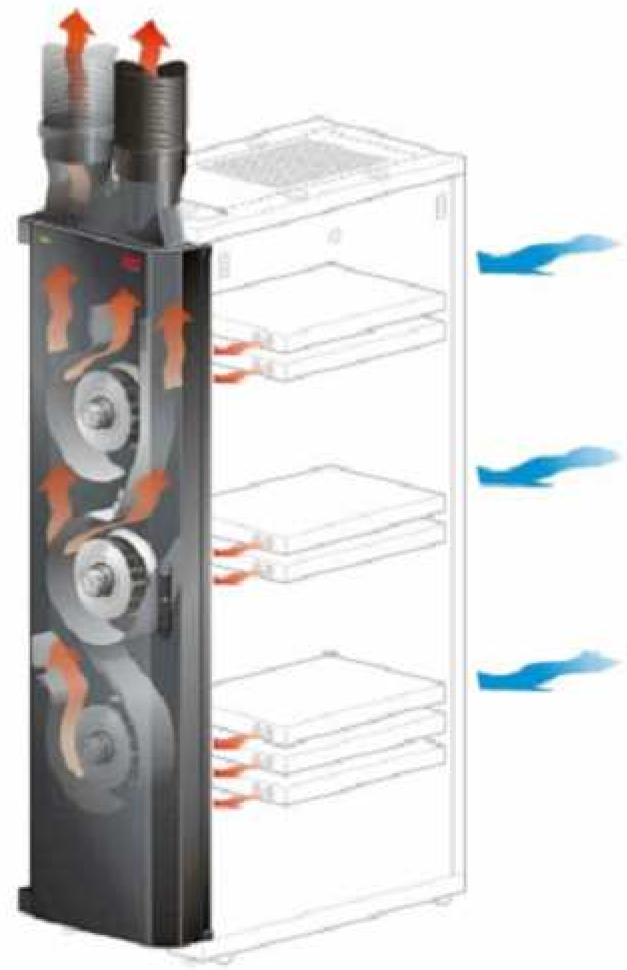

APC® Legendary Reliability

ADU Air Flow

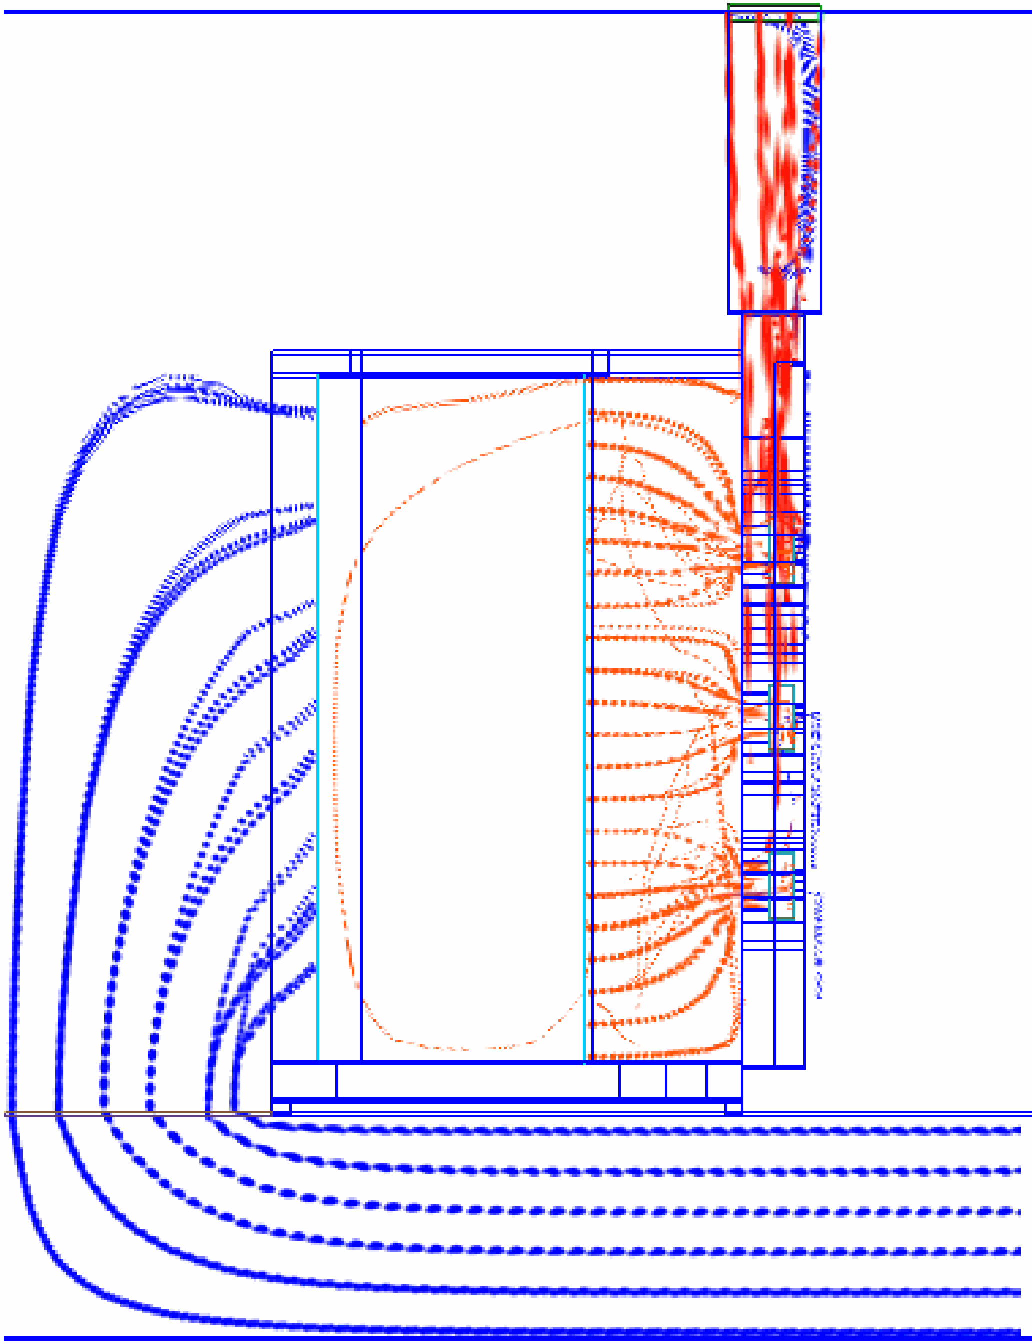

ARU with Fully Ducted Return

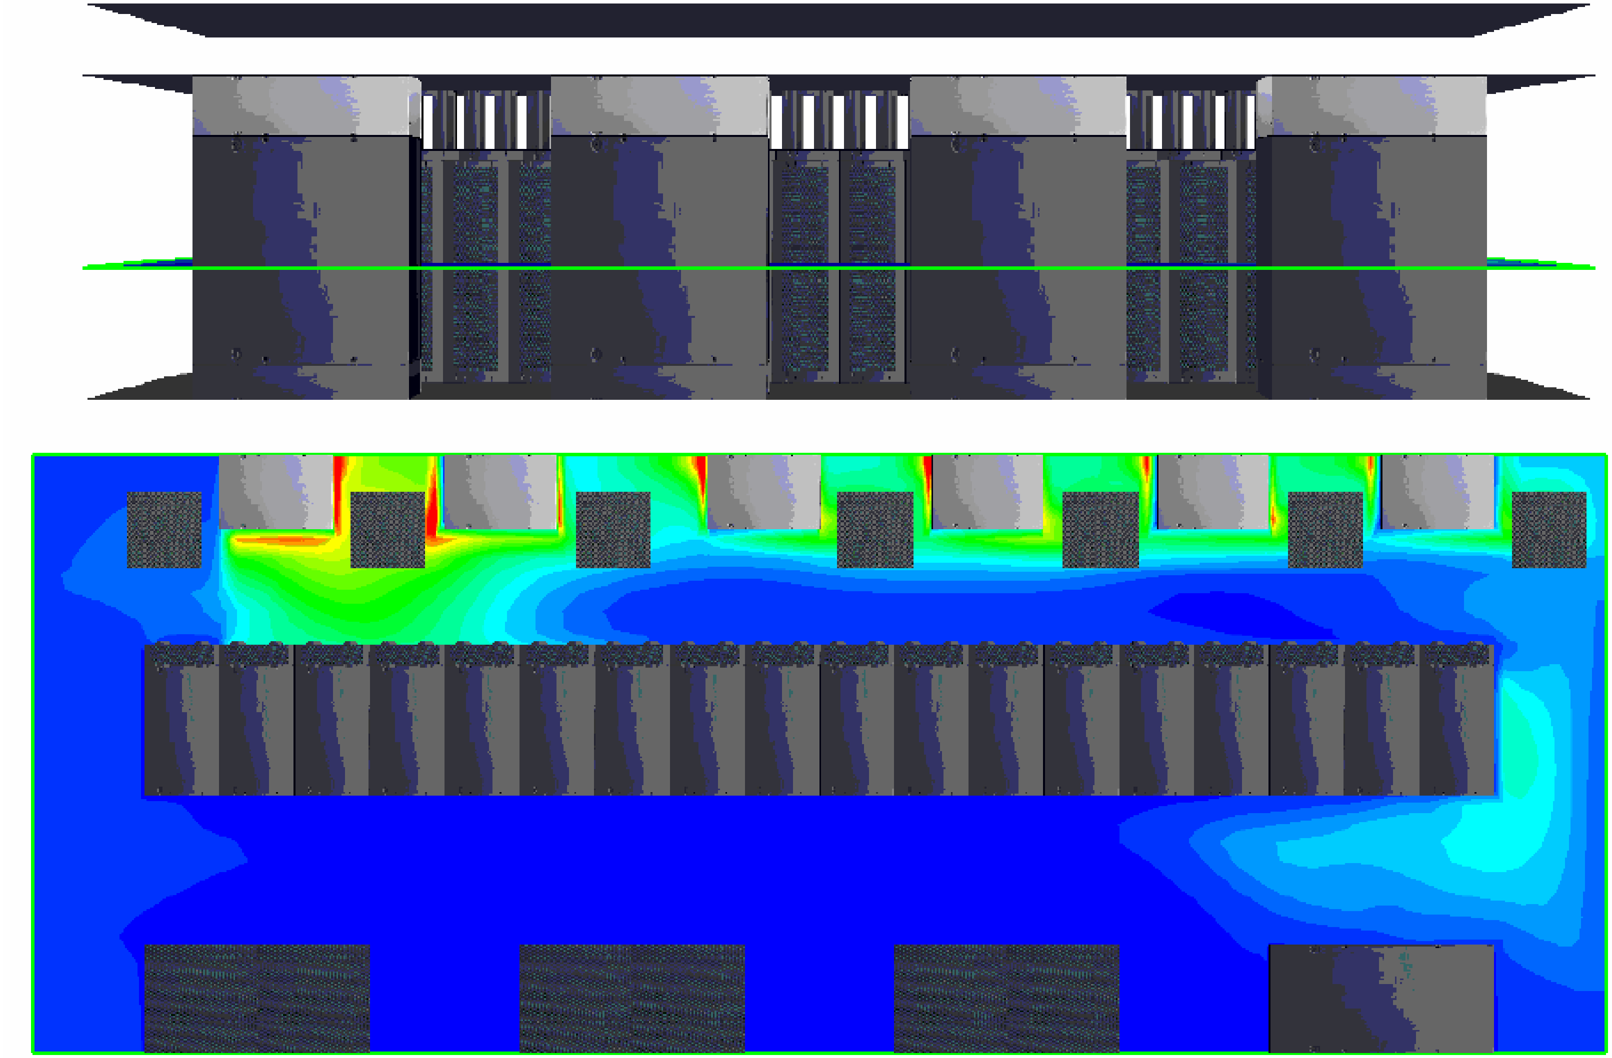

Strategy #3
Borrow Cooling From Adjacent Racks

- **Common and effective strategy in typical data centers**
- **Unused cooling capacity from neighboring racks can be used for up to 3x design density**
- **Uses rules for locating high-density racks to avoid creating hot spots**

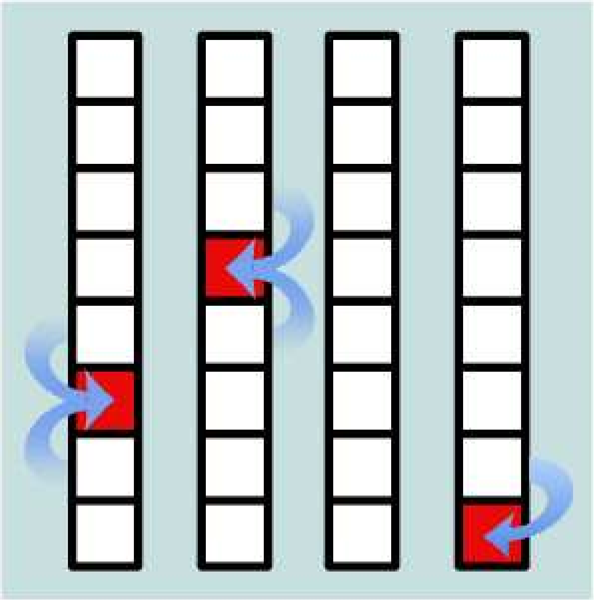

- **Compliance can be verified by monitoring power consumption at the rack level**

APC® Legendary Reliability

"Borrowed Cooling" Considerations

Strategy #3

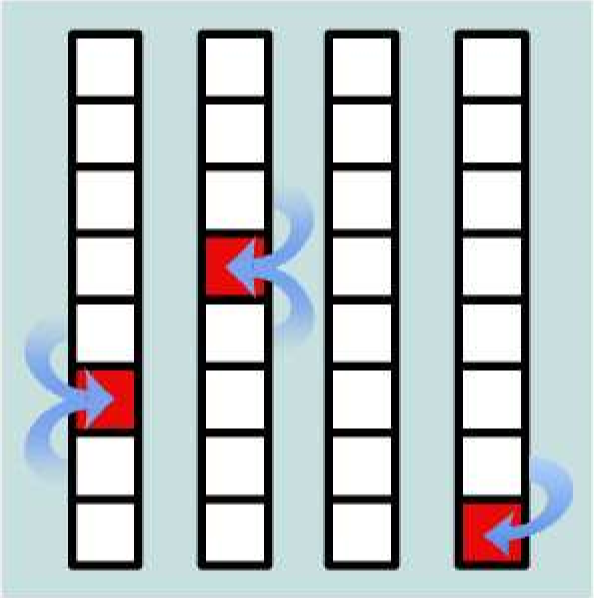

- No new equipment needed
- Essentially "free" in many cases
- High-density racks are limited to about 2x average density
- Requires more floor space than supplemental-cooling strategy (lower density)
- Requires enforcement of deployment rules

APC® Legendary Reliability

"Borrowed Cooling" Applications

Strategy #3

- **For existing data centers, when high-density equipment is a small fraction of total load**

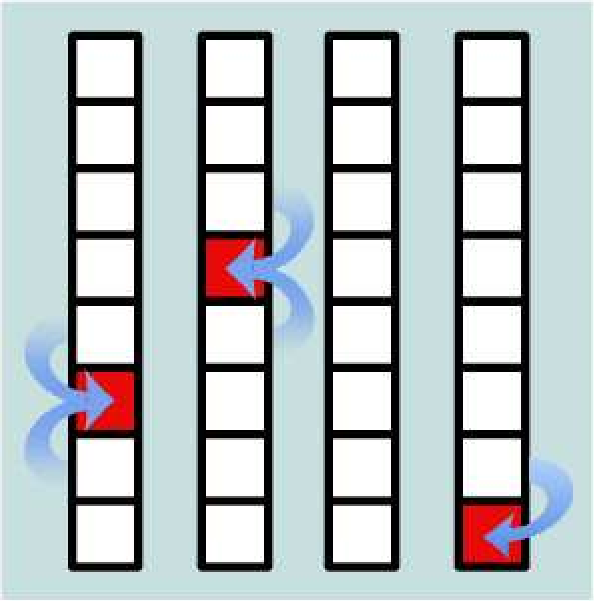

"Borrowed Cooling" Example of Deployment Rules

Strategy #3

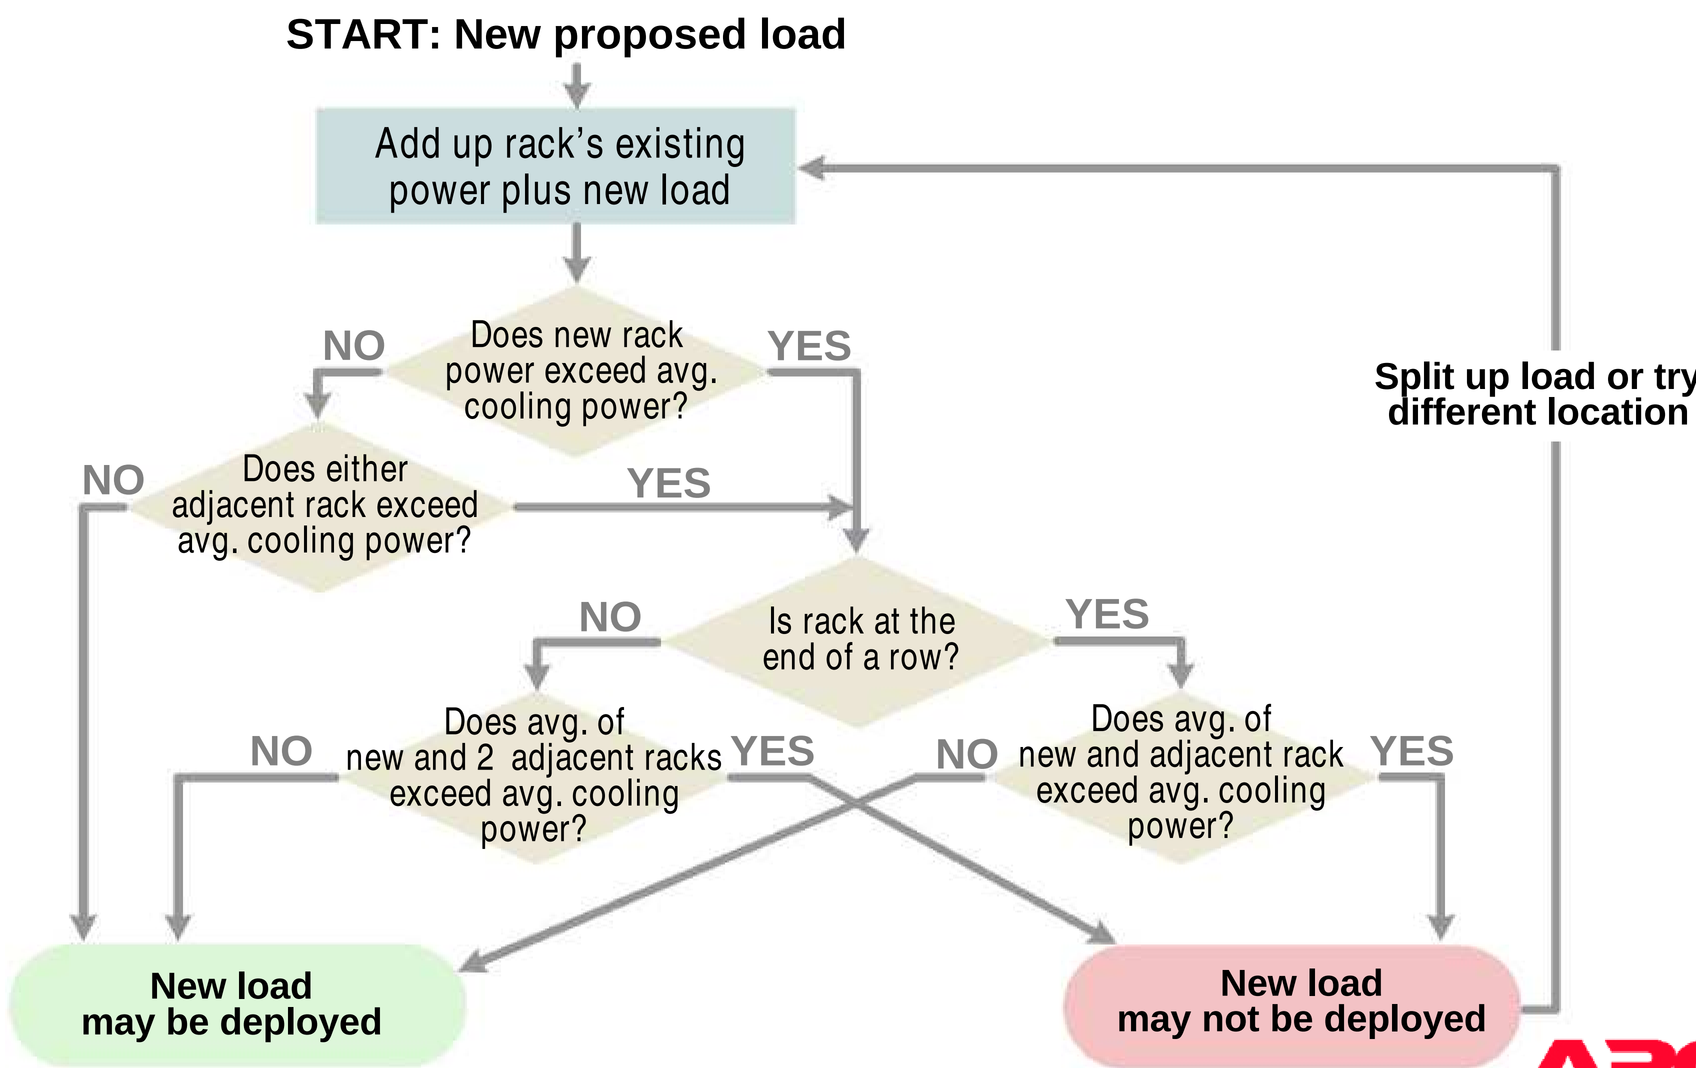

"Borrowed Cooling" Example of Deployment Rules

Strategy #3

"If the proposed rack averaged with its neighbors (only one neighbor if at the end of a row) does not exceed design average power, AND neither neighbor is already a high-density rack, then it's OK to put it there"

Strategy #4
Split Up High Density Equipment

- **Most popular solution**
- **High-density equipment is spread out among many racks**
- **No rack exceeds design power density**
- **Predictable cooling performance**

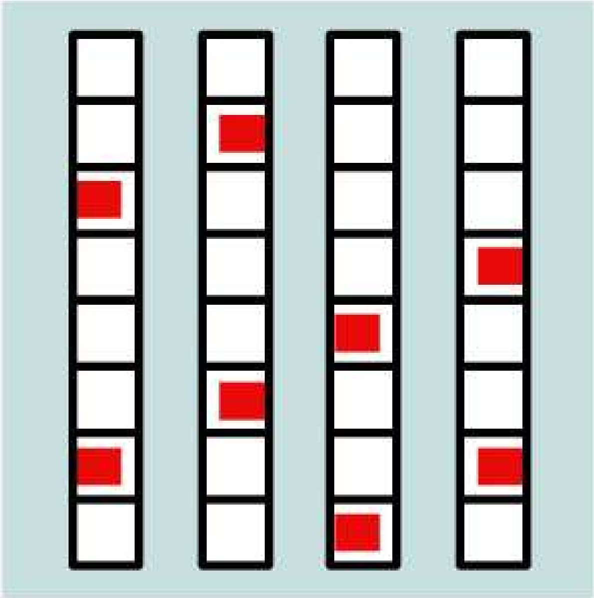

Splitting Up Equipment: Considerations

Strategy #4

- **No new equipment needed**
- **Essentially “free” in many cases**
- **High-density equipment must be spread out even more than in “borrowing” strategy**
- **Uses more floor space than full high-density racks**
- **Can cause data cabling problems**
- **Empty vertical space must be filled with *blanking panels* to prevent in-rack recirculation of hot exhaust air**

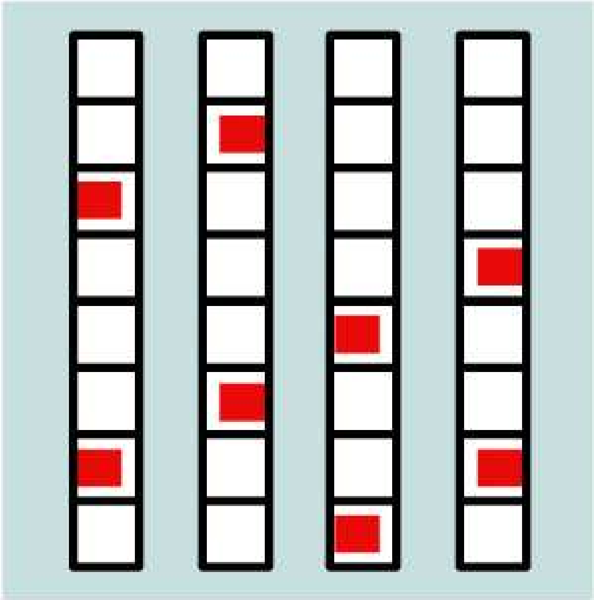

APC® Legendary Reliability

Splitting Up Equipment: Applications

Strategy #4

- **For existing data centers, when high-density equipment is a small fraction of the total load**

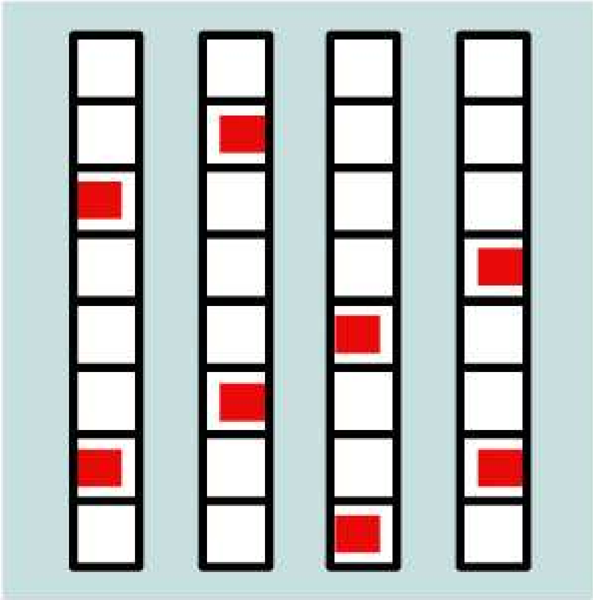

Blanking Panels

Strategy #4

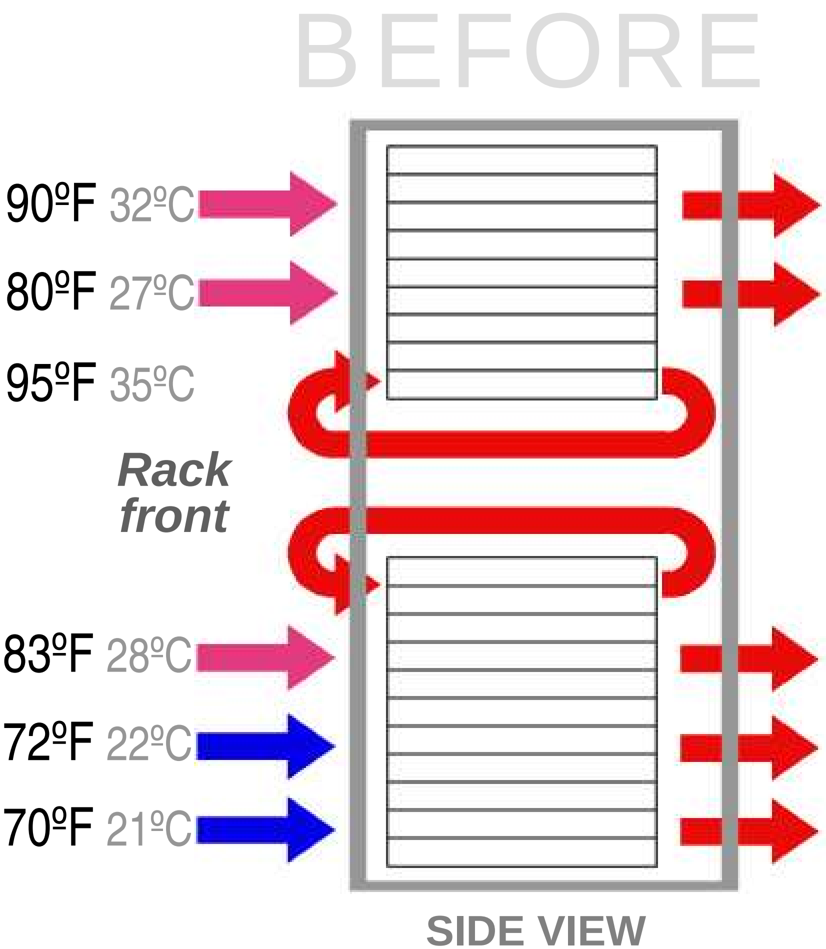

Air Flow with no blanking panels

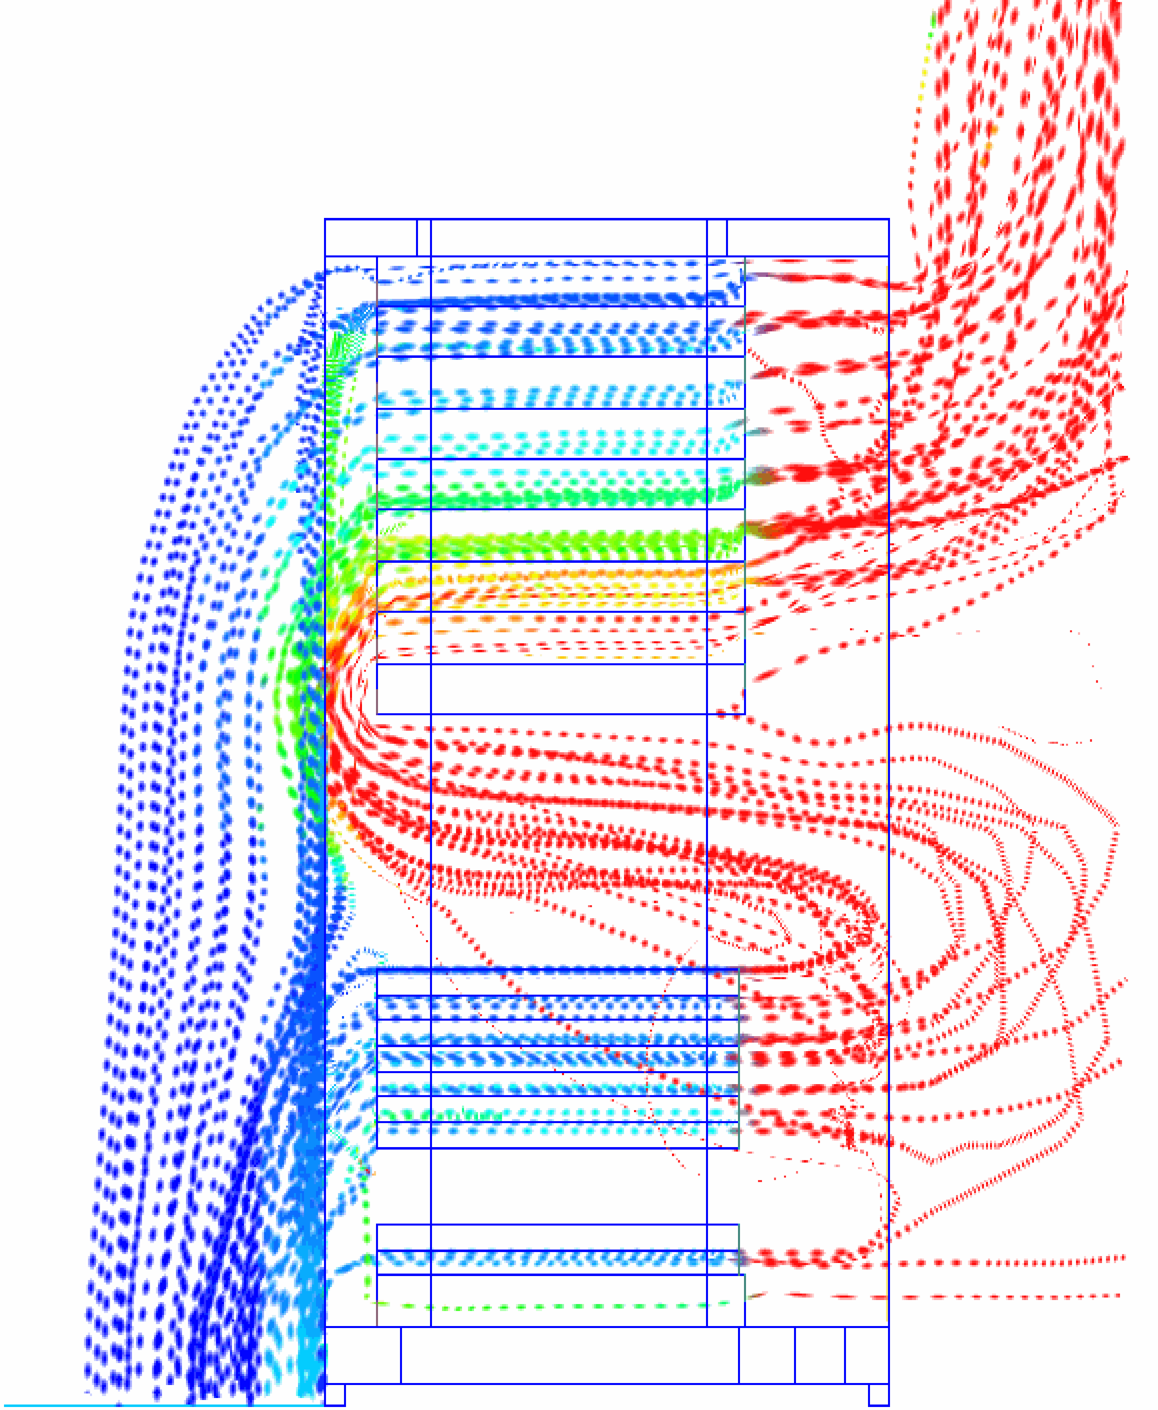

Blanking Panels

Strategy #4

Blanking panels block internal recirculation

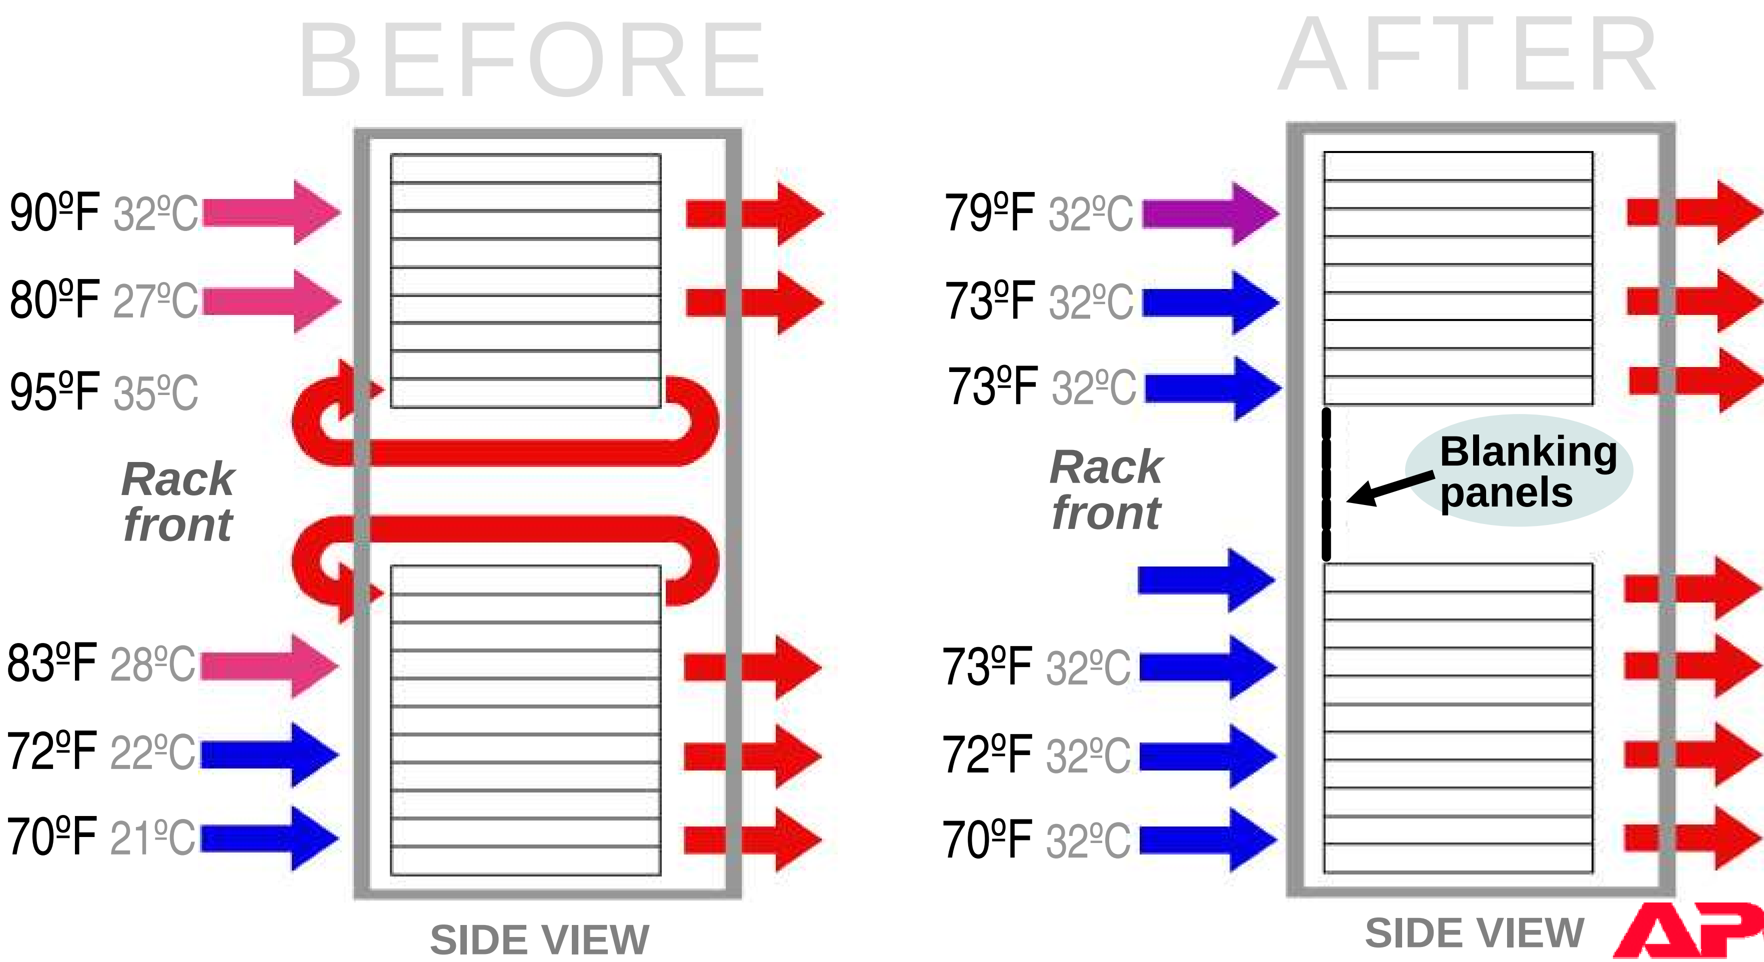

APC® Legendary Reliability

Air Flow with blanking panels

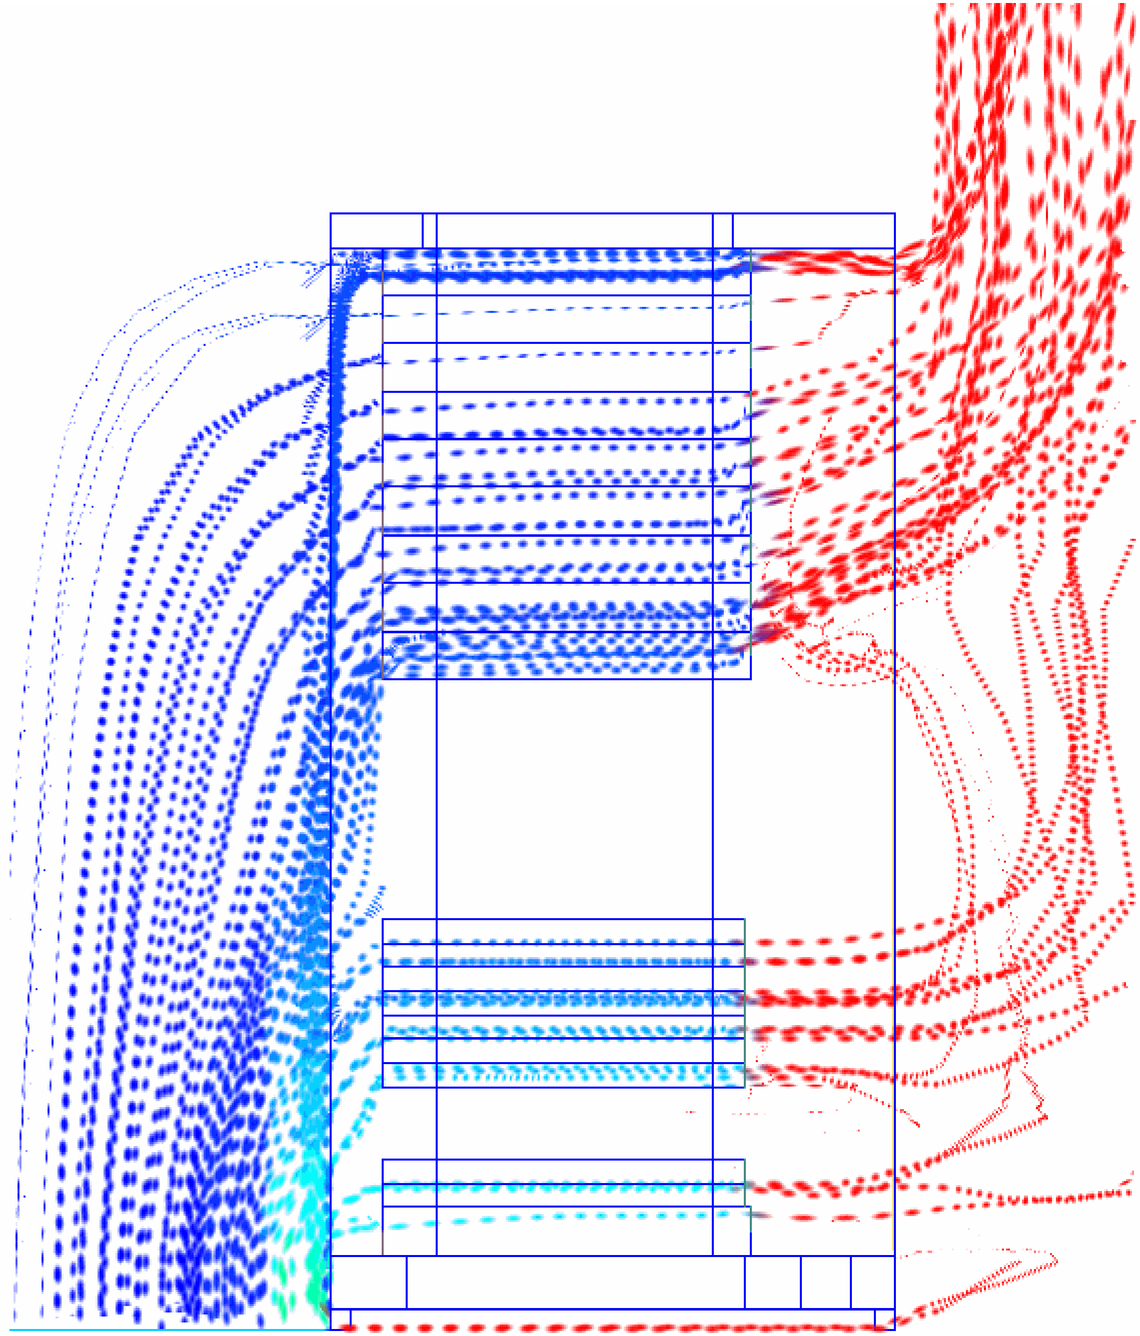

Blanking Panels

Strategy #4

Snap-in blanking panel

Strategy #5
Dedicated High-Density Area

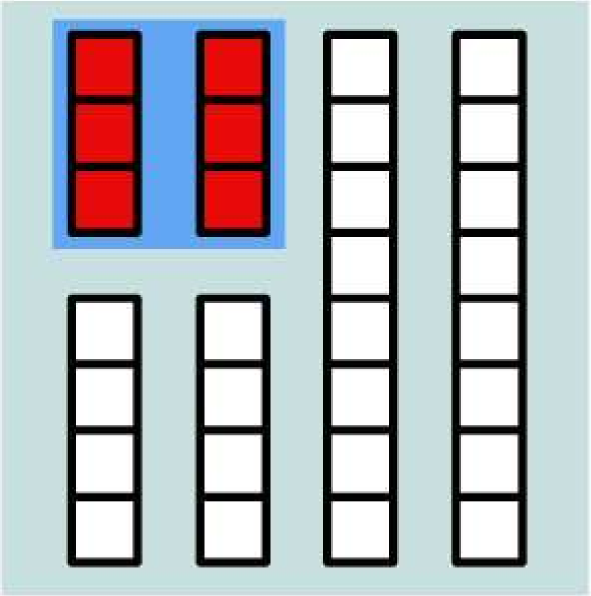

- **Supports maximum-density racks**
- **Doesn’t require spreading out of high-density equipment**
- **Optimal floor space utilization**
- **New technologies can deliver predictable, highly efficient cooling**

Dedicated High-Density Area: Considerations

Strategy #5

- **Requires prior knowledge of number of high-density racks**
- **Need to plan high-density area in advance and reserve space for it**
- **Requires ability to segregate high-density equipment**

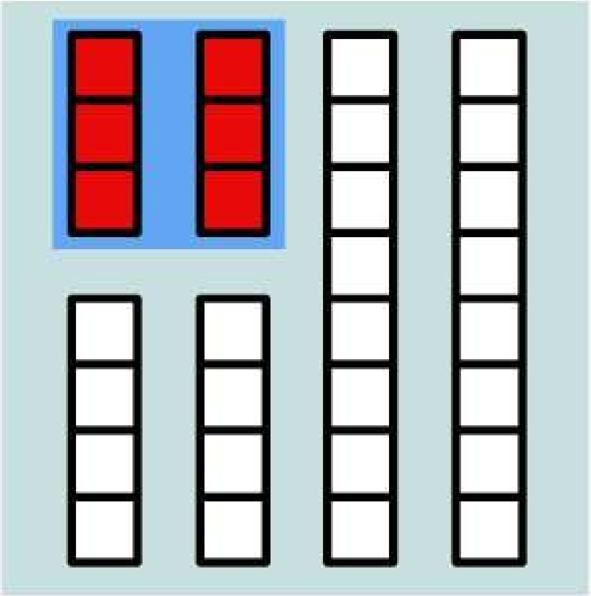

Dedicated High-Density Area: Applications

Strategy #5

- **New construction or renovations**
- **Density of 10-25 kW per rack**
- **High-density co-location**

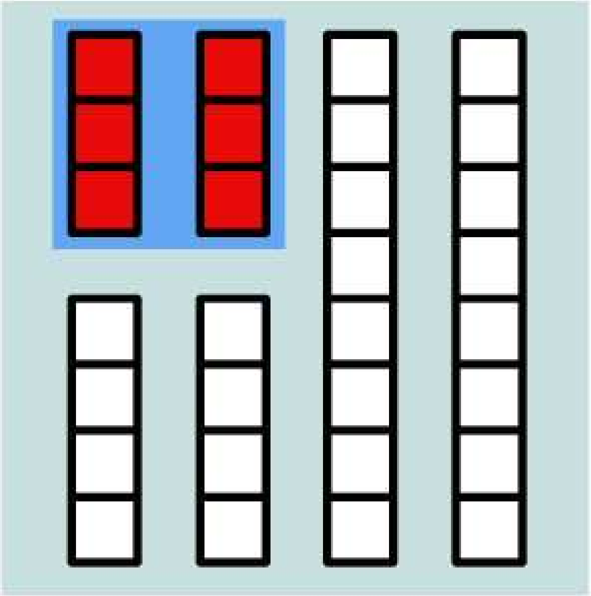

APC® Legendary Reliability

Dedicated High-Density Area: Power / Cooling Technology

Strategy #5

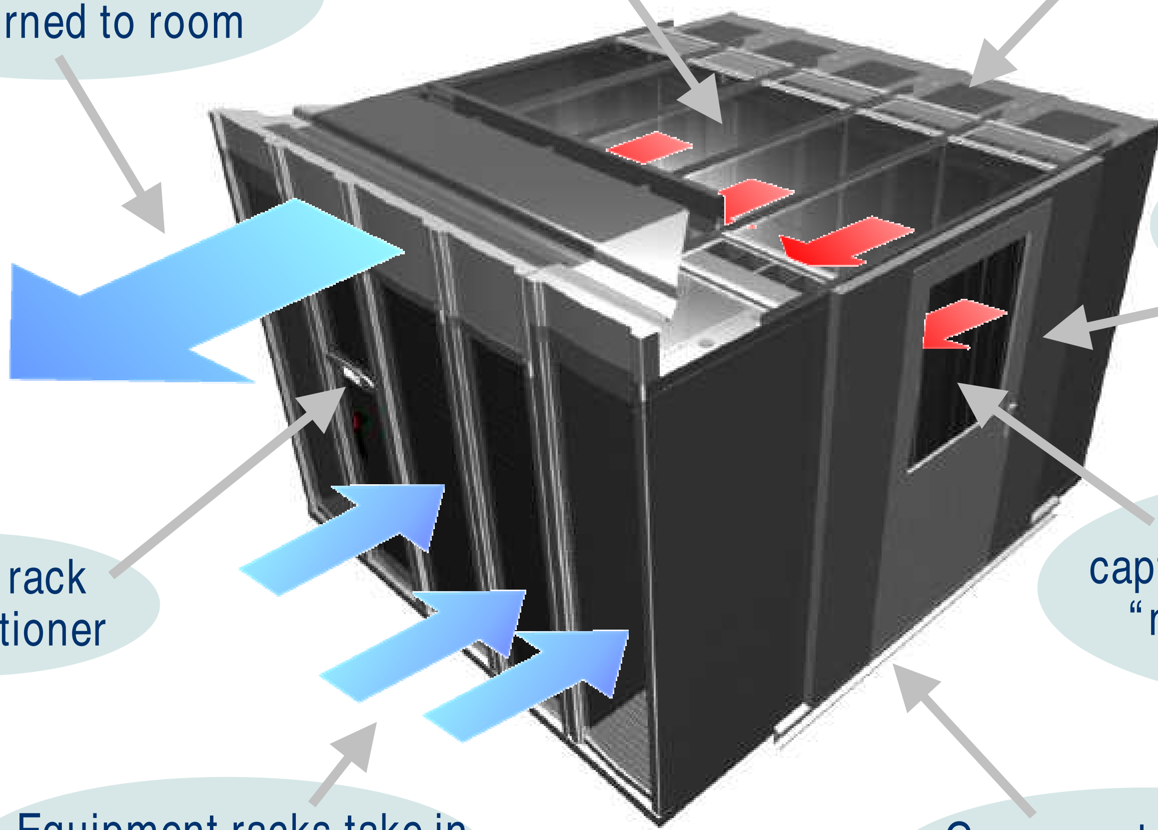

Legendary Reliability

ISX High Density Air Flow Pattern

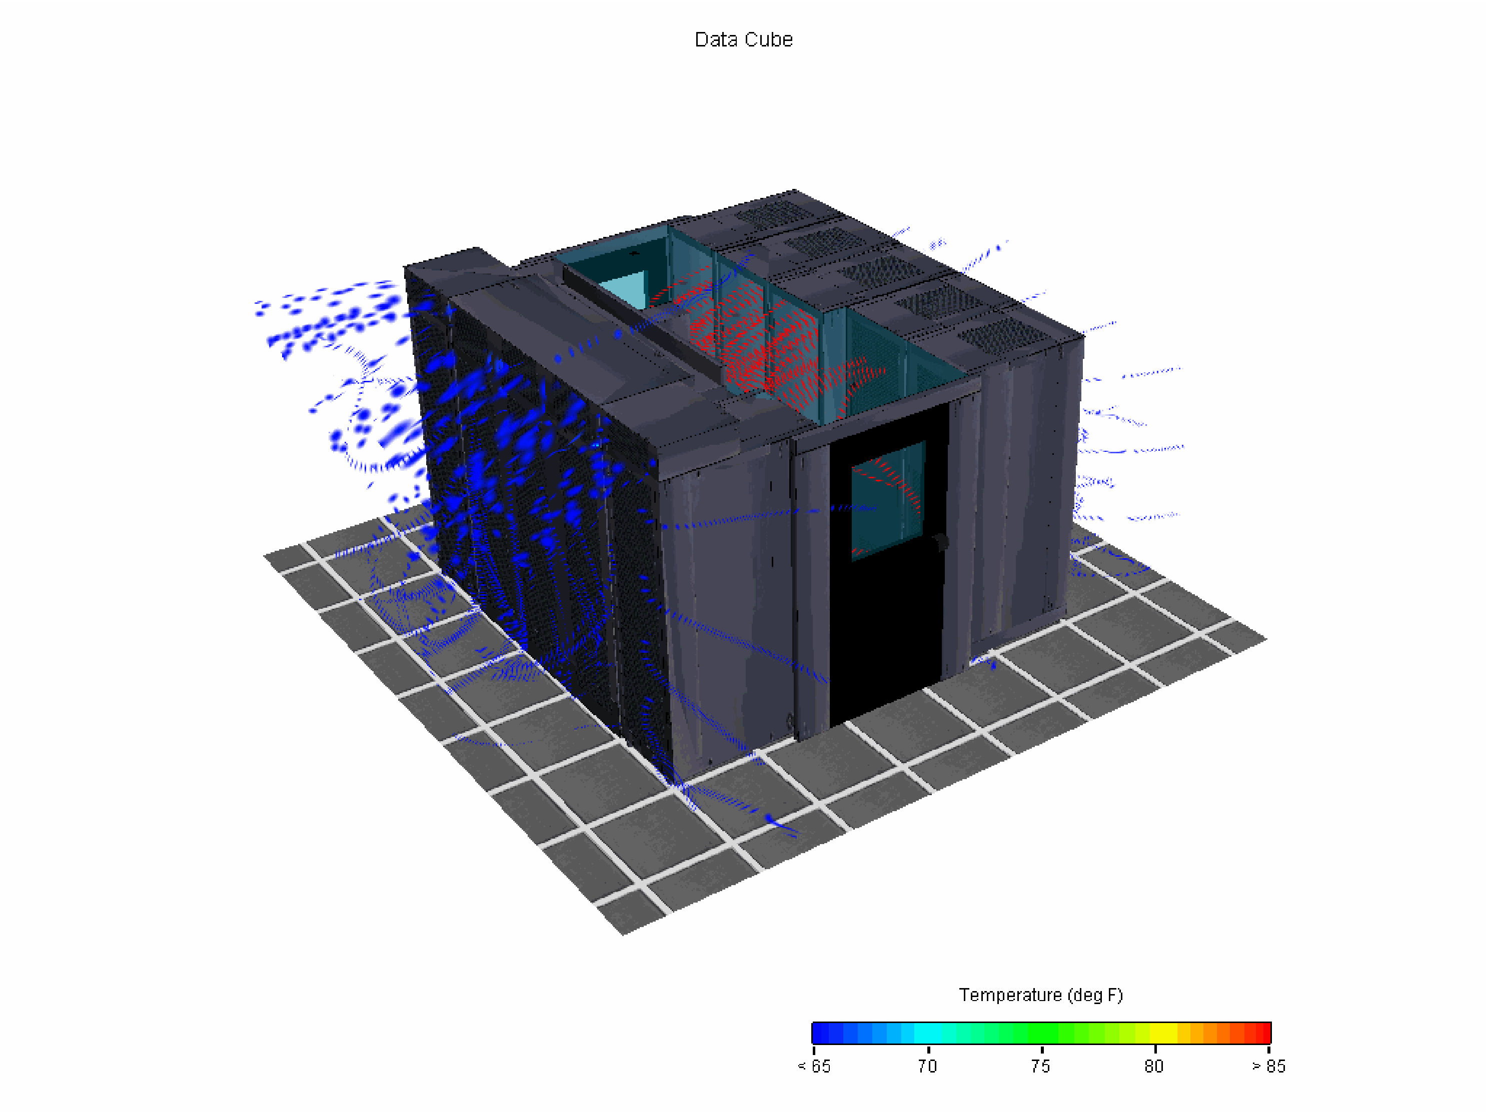

Summary and Recommendations

Seven Elements of an Optimal Cooling Strategy

Elements of Optimal Cooling Strategy

1

Ignore physical size of equipment and focus on functionality per watt consumed

Minimizes area and TCO

Elements of Optimal Cooling Strategy

2

Design the system to permit later installation of supplemental cooling devices

Allows for future supplemental cooling equipment where and when needed, on a live data center, in the face of uncertain future requirements

Elements of Optimal Cooling Strategy

3

Choose a room average power density between 40 and 100 W / ft^2

0.4 – 1.1 kW / m2

Practical for most new designs:
80 W / ft2 or 2.8 kW / rack
.9 kW / m2

- *Avoids waste due to oversizing*
- *Keeps both routine and redundant performance predictable*

APC® Legendary Reliability

Elements of Optimal Cooling Strategy

4

If the fraction of high density loads is high and predictable, establish high-density areas of 100-400 w / ft^2 (3-13 kW per rack)

1.1 – 4.3 kW / m^2

- *Requires planning ahead for specially equipped areas*
- *Adds significant cost, time, and complexity*
- *These areas do not use raised-floor cooling*

Elements of Optimal Cooling Strategy

5

Establish policies/rules for allowable rack power based on location and adjacent loads

- *Reduces hot spots*
- *Ensures cooling redundancy*
- *Increases system cooling efficiency*
- *Reduces electrical consumption*
- *More sophisticated rules and monitoring can enable even higher power density*

APC®
Legendary Reliability

Elements of Optimal Cooling Strategy

6

Use supplemental cooling devices where indicated

Boosts local cooling capacity up to 3x room design to accommodate high-density equipment

Elements of Optimal Cooling Strategy

7

Split up equipment that cannot be installed to meet the rules

- *Lowest cost, lowest risk option*
- *Consumes considerable space if there is more than a small fraction of high-density loads*
- *Chosen as primary strategy by users without significant area constraints*

Questions?

APC®

Legendary Reliability®