IT@Intel White Paper
Intel IT
IT Best Practices
Facility and IT Innovation
January 2012

Facilities Design for High-density Data Centers

We have enabled high density computing and lowered facility costs by adopting a proactive server refresh strategy and investing in the latest generation Intel® Xeon® processors.

John Musilli
Senior Data Center Architect, Intel IT

Brad Ellison
Senior Data Center Architect, Intel IT

Executive Overview

At Intel IT, we have enabled high density computing and lowered facility costs by adopting a proactive server refresh strategy and investing in the latest generation Intel® Xeon® processors. Other contributors to data center density include building an enterprise private cloud and increasing cabinet sizes, including 15 kilowatts (kW), 22.5 kW, and 30 kW. We anticipate that our data center density will continue to increase.

The ever-increasing capabilities predicted by Moore's law—basically, doing more in less space—make it imperative that we optimize data center efficiency and reliability. To do this, we focus on several areas, some of which are industry-standard and simply good room management practices. Others, such as our automated temperature, humidity, and static pressure controls, are proven elements of air-cooled designs developed by Intel.

- **Air management.** Air segregation and automated control sensors enable us to efficiently manage capacity and cooling costs in our data centers.
- **Thermal management.** Excess heat is a major contributor to IT equipment downtime due to internal thermal controls that are programmed to shut down the device. For our high-density data centers, we concentrate on maximizing thermal ride-through.
- **Architectural considerations.** To support and optimize our air and thermal management techniques, we attend to architectural details such as vapor sealing walls, floors, and ceilings; removing raised metal floors (RMF); increasing floor structural loading limits; and using overhead structured cabling paths.
- **Electrical considerations.** Our high-density data centers use flexible redundant busway circuit designs and 415/240 Vac distribution to reduce electrical losses, helping us design an efficient, fault-tolerant room.
- **Stranded capacity.** We proactively identify stranded capacity in all areas of the data center, such as power and cooling capacity, to increase efficiency and avoid having to make expensive changes to the existing infrastructure.

By focusing on these facility design areas, we can increase room infrastructure capacity, which enables higher densities while reducing operating costs and maintaining the levels of redundancy required by our business models.

Contents

IT@INTEL

The IT@Intel program connects IT professionals around the world with their peers inside our organization – sharing lessons learned, methods and strategies. Our goal is simple: Share Intel IT best practices that create business value and make IT a competitive advantage. Visit us today at www.intel.com/IT or contact your local Intel representative if you'd like to learn more.

BACKGROUND

Intel's business strategy includes consolidating our data centers to increase utilization. To reach this goal, we have increased the density of our data centers over the last decade by putting more powerful servers into less space. Each new generation of Intel® Xeon® processor features more compute power and more cores. Therefore, as we refresh our servers on a proactive three- to four-year schedule, a single new server can take the place of several older ones—accomplishing more work in a smaller footprint. In addition, building Intel's private enterprise cloud, by virtualizing more than 60 percent of our Office and Enterprise applications, has put more demands on our data centers, thereby increasing their density. We anticipate that the density of our data centers will continue to increase with current designs including 15-kilowatt (kW), 22.5-kW, and even 30-kW cabinets.

Although high-density data centers can be more efficient in terms of energy consumption and use of space, they also pose a potentially greater business continuity risk. A low-density data center with less than 5 kW per cabinet might have 2,500 servers in 20,000 square feet. An average incident, such as a rack- or row-level event, might result in the loss of only 10 percent of those servers. In higher density data centers, however, that same compute or storage capacity may be located in only 500 square feet. This higher concentration of servers makes it more likely that an incident will result in the failure of at least 50 percent of the servers and network devices, depending on the switch topology.

Mitigating these risks and challenges, while minimizing costs, affects every aspect—from the slab to the roof—of data center design. The savings, risks, and the resulting design must be carefully balanced to meet customer business requirements based on the intended use of each data center.

HIGH-DENSITY FACILITY DESIGN AT INTEL

Recognizing both the benefits and the risks of high-density computing and storage, Intel IT focuses on several design areas, some of which are industry-standard and simply good room management practices. Others, such as our automated temperature, humidity, and static pressure controls, are proven elements of higher density air-cooled designs developed by Intel.

We focus our design efforts on five main areas: air management, thermal management, electrical design, architectural design, and harvesting stranded capacity. Focusing on these areas helps us to reach our goal of consolidating our data centers to reduce operating cost while still maintaining the levels of redundancy required by our business models.

Air Management

Airflow management is the key to increasing the cooling capacity of facilities. Airflow design for new construction can influence the size and shape of rooms or buildings that house data centers and can determine the maximum efficiency and density of an air-cooled data center. Efficient airflow design in existing data centers increases energy efficiency, reduces operations cost, avoids hot spots, and, most importantly, supports high-density computing by enabling us to use larger per-cabinet loads, such as 15–30 kW.

We use computational fluid dynamics (CFD) algorithms, developed for Intel's manufacturing facilities, to model and design air segregation configurations for our data centers. The right configuration enables our servers to operate at optimum temperature without wasting energy.

Our air management goals include the following:

- Supply the necessary amount of cool air to each cabinet
- Remove the same amount of hot exhaust air from each cabinet
- Keep the hot exhaust air away from the equipment air intake
- Maximize the temperature difference between cold and hot air (Delta T) to improve cooling efficiency

These goals must be achieved in a redundant and uninterrupted manner, according to business requirements.

AIR SEGREGATION

In our experience, with segregated air paths and flooded room designs we can support densities up to 30 kW per rack server anywhere in the data center. A key design requirement for air segregation is to maintain a constant pressure differential between the supply and return air paths, so that enough conditioned air is provided to meet the actual server demand.

We set the supply air temperature just below the thermal boundary for the first server fan speed increase (80° F (26° C)). This prevents speed-up of the internal fans, which would use 2 to 10 percent additional uninterruptable power supply (UPS) server power, depending on the number of fans and speed steps programmed into the device's BIOS. Although the American Society of Heating, Refrigerating, and Air-Conditioning Engineers guidelines allow a supply air temperature of 80.6° F (27° C), we keep our supply air at 78° F (25.5° C) in the cold aisle, measured between six and seven feet above the finished floor.

Intel's data center designs use both hot aisle enclosures (HAEs) and passive ducted cabinet air segregation, also known as "chimney cabinets." Although our CFD modeling shows that hot aisle and cold aisle enclosures provide the same thermal solution, we have experienced greater benefits with the HAE configuration with return air temperature design range from 95° F (35° C) to 105° F (41° C), such as:

- Greater heat transfer efficiency across the cooling coils due to elevated return air temperatures and a greater Delta T.
- By carefully developing our work procedures, these higher temperatures are practical because data center staff enter this space only for new landings or maintenance, so our staff works in an ambient temperature environment most of the time.
- During a mechanical or electrical failure that causes a loss of cooling to the data center, the walls, floor, and cabinets serve as a "thermal battery," releasing stored energy and contributing to thermal ride-through (discussed in the Thermal Management section later in this paper).

...all of our data centers—even those with 30-kW cabinets—are air cooled, and our design roadmap supports 50 kW per rack for new construction.

Table 1. Rack density and watts consumed per square foot^

Year	Rack Density	Watts Per Square Foot
2004	15 kilowatts (kW)	250
2006	22.5 kW	375
Current	30 kW	500

^ Data is based on a 5,000-square-foot room with cooling units and power distribution units (PDUs) in the same space.

- In our data centers that use a raised metal floor (RMF) for supply air delivery, all of the leaked air—0.5 to 2 percent—is available to cool the IT equipment in areas outside the HAE.
- The common building electrical components, such as panels, breakers, and conductors, outside of the HAE, are not subjected to the elevated temperatures found in the cold aisle designs and therefore do not need to be evaluated for thermal de-rating.

AIR COOLING

The density of server racks—and therefore the power used per square foot—has been steadily increasing, as shown in Table 1.

Although there is no consensus in the industry that high-density racks can be air cooled effectively, all of our data centers—even those with 30-kW cabinets—are air cooled, and our design roadmap supports 50 kW per rack for new construction.

15-kW Cabinets

For low-density cabinets up to 15 kW, we use traditional air cooling with the computer room air conditioning (CRAC) or computer room air handler (CRAH) units in the return airstream without any air segregation. This configuration is illustrated in Figure 1.

22.5-kW Cabinets

For 22.5-kW cabinets, we use HAEs with the CRAC unit in the return airstream, as illustrated in Figure 2.

30-kW Cabinets

For 30-kW cabinets, we use flooded supply air design, passive chimneys—ducted server cabinets without fan assist—as well as HAEs. We have found that passive ducted equipment cabinets are very effective at delivering server exhaust to a return air plenum (RAP). Passive ducting eliminates the need for fans, which are problematic from a monitoring, maintenance, and replacement perspective. Our design for 30-kW cabinets is shown in Figure 3.

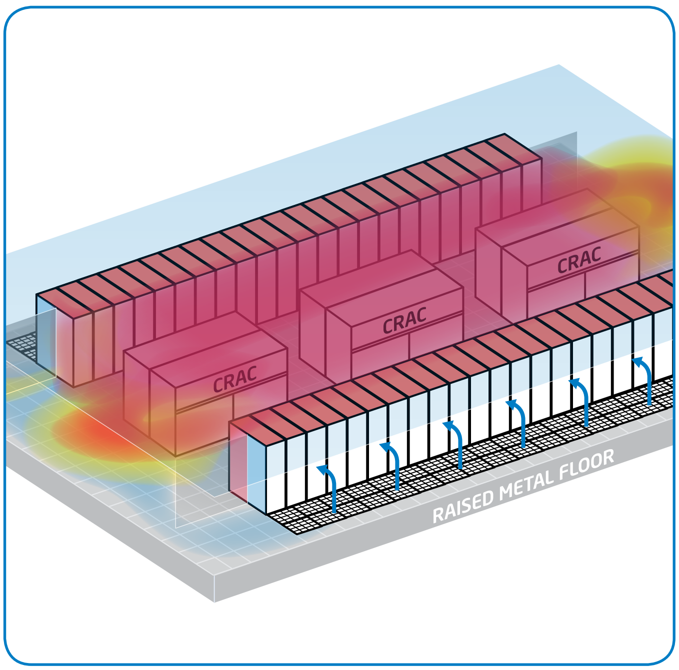

Figure 1. Simply moving the computer room air conditioning, or computer room air handler units, into the return air stream without air segregation provides sufficient cooling for 15-kilowatt cabinets.

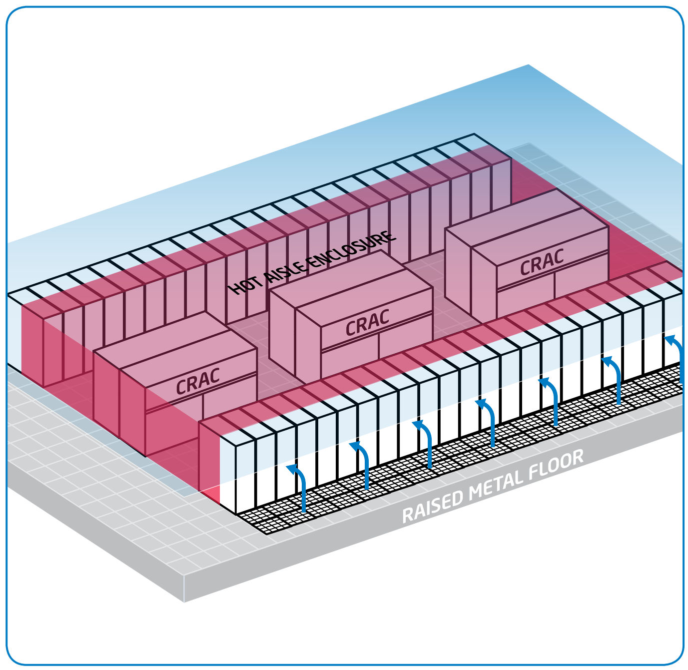

Figure 2. To accommodate the greater cooling needs of 22.5-kilowatt (kW) cabinets, we bolster the standard 15-kW design with walls, forming a hot aisle enclosure.

FLOODED SUPPLY AIR DESIGN

A flooded supply air design, used in conjunction with chimney cabinets or HAEs, eliminates the need for a raised metal floor (RMF) and enables conditioned air to be introduced into the room from any direction at the required volume. An optimum design for a data center without an RMF can use air segregation, an air-side economizer, or a fan-wall configuration.

In this type of design, we keep the supply-air velocity in front of the server and storage devices at or below what we call "terminal velocity," targeted at less than 400 feet per minute (fpm). A supply air velocity greater than 400 fpm can exceed the ability of some IT equipment to draw the air into the device.

Some facilities use an RMF as an access floor for power cooling and network infrastructure. This does not prohibit a flooded air design—we simply consider the RMF elevation as the slab height, and raise the room air handlers' supply path above the RMF, as shown in Figure 4.

AUTOMATED CONTROL SENSORS

In the past, our room air handler designs used uncontrolled direct-drive fans to meet the air moving and cooling requirements, with the return air temperature control sensors. However, this was inefficient because Intel data centers rarely operate at 100-percent design load. Typically, data centers operate at about 60 percent of peak load; however, a particular job might push the workload to 90 percent for a few hours, which then requires maximum design cooling during that period.

To improve the energy efficiency of our data centers, we have adopted manufacturing clean room management technology that uses real-time data to supply the appropriate temperature and volume of air for the data center load at any particular time. We have also replaced single-speed, direct-drive fan motors in the air handling systems with variable frequency drives (VFDs) to allow the throttling of cold air during fluctuations in server load and heat generation.

Sensor case study

To gather real world data, we installed wired and wireless automated temperature, static pressure, humidity, and power monitoring sensors at a 1.2 megawatt production data center that did not have air segregation. These sensors sample the environment every 15 seconds and average the readings over 1 to 6 minutes.

The sensor array includes:

- **CRAC and air handling unit (AHU) flow and temperature sensors.** We placed flow meters in the chilled water supply and return pipes to determine the actual real-time work done by each unit. We placed two additional temperature sensors in the supply air duct and six sensors in the return air path to identify short cycling and the Delta T across the cooling coil.

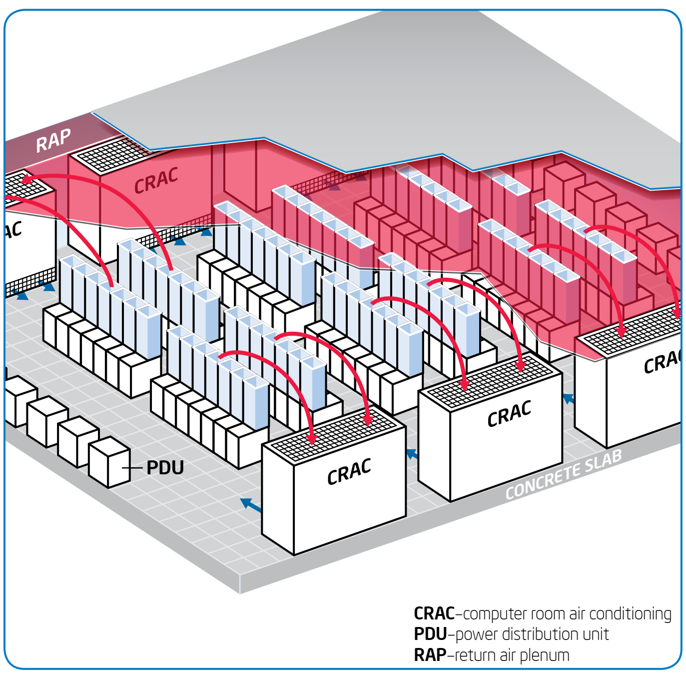

Figure 3. Passive ducted equipment cabinets are very effective at delivering server exhaust to a return air plenum in the ceiling, allowing us to air-cool 30-kW cabinets.

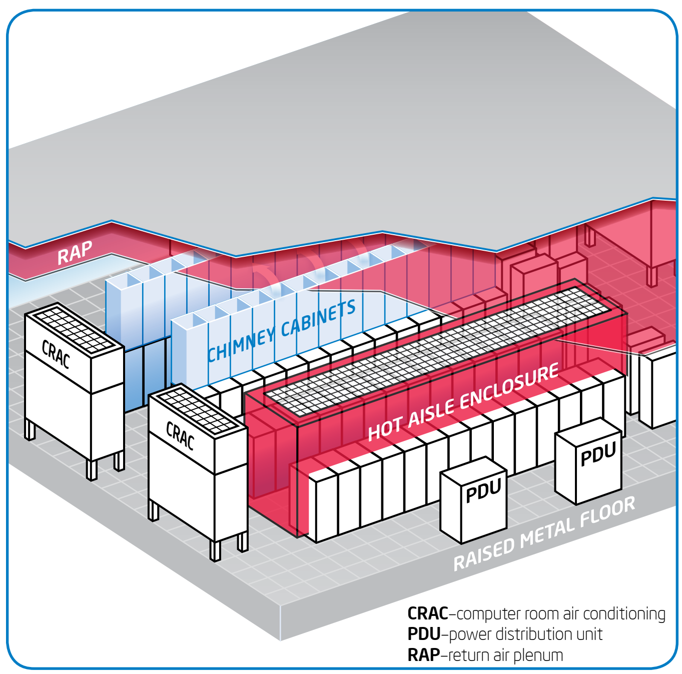

Figure 4. For rooms with a raised metal floor (RMF), we raise the room air handlers' supply path above the RMF to achieve a flooded supply air design used in conjunction with chimney cabinets or hot aisle enclosures.

Business Value of Automated Control Sensors

In one five-year-old room with traditional hot and cold aisles and no air segregation, we were oversupplying 164,000 cubic feet per minute (CFM) of air and 120 more tons of cooling than necessary to maintain the room environmentals within the upper operating thermal boundaries of 78° F (25° C). With a USD 205,000 investment in HAEs, VFDs, and cascade supply air temperature and pressure differential sensors, we were able to power off six of sixteen 30-ton CRAC units and operate the remaining units at 60-percent fan speed. The AHUs currently provide supply air at 75° F (24° C). This generated an annual power savings of USD 88,000 and captured a stranded cooling capacity of 490 kW, providing a return on investment of less than 2.5 years. Based on the success of this project, we modified seven additional rooms globally in a similar manner.

- Rack-level temperature sensors. We established rack-level temperature monitoring through multi-elevation supply air sensor readings and from on-board server sensors. The latter provide a more accurate picture of the data center's cooling air distribution. We placed an additional sensor at the top rear of the rack to identify air recirculation from the front or rear of the rack. We used this information to identify uneven distribution, recirculation, or mixing of supply and return airstreams.
- Electrical power sensors. For monitoring and efficiency metrics we used substation, UPS, and busway meters to identify data center loads and losses. We deployed electrical power sensors throughout the chilled water plant electrical distribution system and individually on the CRAC or AHUs. We also set up a wireless weather station in the chiller equipment yard to record the local conditions and correlate these to chiller efficiency.
- Rack-level power monitoring. These sensors identified what influence the rack density heat load had on cooling in a specific area of the room.

As a result of the case study, we made some design changes. We installed cascade temperature and static pressure controls. We programmed these controls to reset supply air temperature at the IT device to the highest possible design temperature (78° F (26° C)) and to deliver the needed air volume to the room. The controls used direct supply air temperature feedback at the top of the rack between six and seven feet from the finished floor and averaged pressure differential feedback from multiple sensors in the supply and return air paths.

Using data from the sensors, the sensor controls adjust the chilled water supply valve and the supply air fans on the room cooling units as needed to meet demand.

Thermal Management

After power loss to IT equipment, excess heat is a major contributor to IT equipment downtime due to internal thermal controls that are programmed to shut down the device. For our high-density data centers, we concentrate on maximizing thermal ride-through and controlling room environmental factors.

THERMAL RIDE-THROUGH

Thermal ride-through is the amount of time available to absorb the heat generated in the data center environment and maintain acceptable equipment temperatures after loss of cooling. Thermal ride-through is important because as long as devices are operating, they are generating heat without an active means to neutralize it.

It is possible that a data center will continue running for a period of time—using a UPS, batteries, fly-wheel, or other power-storage system—even if power to the cooling system fails due to a failure of utility power, a generator, or an automatic transfer switch. In order to prevent extreme heat buildup, we use various approaches to supply residual cooling. Our thermal ride-through designs extend the cooling capacity beyond the power storage capacity time by five to seven minutes. This removes any residual heat from critical devices before temperatures exceed the physical, logical, or pre-determined thermal boundary of any data center equipment.

The amount of time an environment can reach a certain temperature is not an exact formula. However, we consider the amount of air changes through the equipment, the inlet air temperature to the equipment, and the Delta T across the equipment in determining how fast a room will overheat.

Intel's chilled water thermal ride-through solutions use the existing thermal storage capacity in the form of existing pipe fluid volume and the addition of chilled water storage tanks, if needed, to meet the minimum cooling demand.[1]

[1] See "Thermal Storage System Provides Emergency Data Center Cooling," Intel Corp., September 2007.

In these data centers, we place room air handler fans and a chilled water pump on a separate mechanical load UPS to move existing air and water through the room air handler coils. The amount of total heat removed is proportional to the temperature and volume of chilled water available. For direct expansion systems, we install the mechanical load UPS to support the minimum cooling requirements before the UPS batteries are drained or before the start of the generator and compressor cycle times.

In some designs, we use the outside air environment as a supply or exhaust solution during emergency thermal events to increase the total usable room cooling capacity. We exhaust air out of the room, either to a hallway or outside the building through a wall or roof vents.

At times, we mix data center and facility motor loads on the UPS systems because we have found this approach can reduce harmonic current content and improve the overall output power factor (PF) of the UPS units. We use VFDs to feed larger motors because they provide in-rush isolation from the other UPS loads. Some motors are served from UPS units without VFDs but are usually smaller in horsepower compared to the UPS capacity. Depending on the size of motors to be used and UPS manufacturer, we may consult with the UPS manufacturer to have the UPS properly sized for the motor in-rush current.

ROOM ENVIRONMENTALS

Whenever possible, we use room-level global integrated control logic instead of standalone device control logic. This allows all the mechanical devices to work together to efficiently maintain the environmentals within the zone of influence for a particular cooling or airflow device. We carefully control the humidity, static pressure, and temperature in our data centers, as shown in Table 2.

For example, typically we keep the relative humidity between 20 to 80 percent for server environments. However, tape drives and tape storage may have different environmental specifications and may require a separate conditioned space.

Architectural Considerations

To support and optimize our air and thermal management efforts, our data centers feature several architectural modifications that may not be typical in many data centers.

- **Sealed walls, ceiling, and floors.** Sealing these surfaces with a vapor barrier enables better humidity control. Concrete and drywall are porous and allow moisture to enter the data center environment.
- **Removal of RMFs.** We began removing RMFs from our data centers 10 years ago. In addition to enabling better air management, removing RMFs offer the following benefits:
  - More available floor space from no ramps or stairs
  - Floor installation and maintenance cost savings
  - Reduced risk from seismic events and simplified seismic bracing
  - Reduced dead weight floor loading for above-grade rooms
  - Reduced overall room slab-to-deck height requirements
  - Opportunity for ceiling RAP
  - No need for under-floor leak detection and fire suppression
  - No floor safety issues
  - No electrical bonding of flooring components to the earth ground

  Because of the benefits of a no-RMF design, we made this part of our standard data center design three years ago.

Table 2. Room environmental settings used at Intel

Environmental Factor	Setting
Humidity	20% to 80% for servers
Supply Air Temperature	78° F (25° C) Top of equipment, cold aisle
Pressure Differential between Supply Air and RAP	0.015 WC Positive Supply Air (equivalent to 3.736 Pa)
CFM per kW	80 to 120 CFM

CFM – cubic feet per minute; **kW** – kilowatt; **Pa** – Pascal; **RAP** – return air plenum; **WC** – water column

- Increased floor loading specifications. Our older data centers are limited to 125 pounds per square foot; the current design standard is now 350 pounds per square foot. The higher floor loading limit enables the floors to support the additional weight of more fully loaded blade server racks and storage arrays—each weighing between 2,000 to 2,500 pounds.
- Integrated performance "shakedown" testing. We test all systems in a way that resembles actual data center events, including using controlled overload conditions, to determine if controls and procedures are established and can be used to prevent or mitigate loss of data center services. All systems are integrated and tested to the maximum fault capacity and tolerances, per engineered or factory specifications.

STRUCTURED CABLING

Departing from industry-standard data center design, for the past 10 years we have used custom-length power and network cables and have eliminated the use of cable management arms on server racks and patch panels in our row-end network access distribution between the switch and the server. As a standard practice, we replace the layer-one networking cabling with each server refresh cycle, allowing for a cable plant upgrade every three to five years. While patch panels and patch cables offer flexibility, they also add four more possible points of failure, along with an approximate doubling of the cost of installation. We have found that network paths rarely change, so the benefits of lower cost and having only two points of failure instead of six far outweigh the flexibility benefit.

Other layer-one network practices we use to minimize risk and total loss of connectivity include the following:

- Network switch physical separation. We add this to the logistical network design to prevent loss of redundant switches or routers caused by a local catastrophic event, including accidental fire sprinkler discharge, overhead water from plumbing or building leaks, damage from falling objects or local rack electrical failure. We use a preferred distance of at least 40 lineal feet (12 meters) between redundant devices.
- Overhead cabling trays. These provide reasonable access and flexibility to run copper and fiber throughout the data center without restricting airflow in older data centers or creating a need for an RMF.
- Armor-covered fiber cable for multiple strand conductors. We use this arrangement in paths going into or out of the data center room or building to prevent damage from activities that are not controlled by data center operators.
- Cable management. We are careful to dress cables away from server chassis fan outlets.

Electrical Considerations

The data center electrical distribution is critical for an efficient, fault-tolerant room. Our high-density data centers use flexible branch circuit designs, which reduce electrical losses. Some of the ways we make our data centers' electrical systems reliable and cost-efficient include the following:

- Overhead busway power distribution. In 2000, our first installed overhead bus was only 100 amps and our rack power demand was 3 kW to 5 kW supporting eight server racks in a redundant configuration. In 2012 designs, the amperage has increased to 800 amps per bus-duct to support higher density server racks.
- UPS derating. The large-scale deployment of refreshed servers that use switching power supply units (PSUs) can trend toward a leading PF. This in turn can cause higher temperatures that negatively affect the UPS output transformer and, depending on the design of the UPS electronics, can create other internal UPS problems. Because of this, we have had to derate some older UPS modules and specific UPS manufacturer's designs. We have seen a leading PF greater than one with new servers in our data centers, accounting for a greater percentage of the attached UPS load. Intel engineers have consulted with UPS manufacturers, determined the correct rating for specific UPS equipment, and updated Intel's design construction standards to reflect these changes. Newer UPS units are designed to handle leading PFs and have resolved the issues generated by the newer PSUs.
- Variable-frequency drives (VFDs). All new fan and pump equipment installations use VFDs, and we retrofit them to existing systems in the field. The building management systems control the VFDs on cooling fans, with inputs from pressure differential sensors.
- Power diversity. Depending on the power configuration, some UPSs can be fully loaded to the manufacturers' design rating as power source one. If the UPS fails, a three-way "break before make" static transfer switch (STS) transfers the power source to utility power as one of our redundant sources of equipment power. Then, when the generator starts, the STS transfers the power source to the generator. When the UPS is restored, the STS transfers back to utility power, then back to the UPS without having to drop the load to the data center.
- High voltage. We use the highest voltage alternating current (Vac) the equipment accepts. This allows for higher efficiency and greater kWs per circuit.
- 415/240 Vac power distribution. In 60-Hz countries, we can provide twice as many kWs per electrical path—one-half the installation cost per kW compared to conventional 208/120 Vac installations—at

the same current. This configuration is compatible with major suppliers of PSUs, which are rated 90 to 250 Vac or 200 to 250 Vac. For existing electrical equipment, we place the 480/415 Vac transformer downstream of the UPS module, generator emergency panel, and UPS bypass breaker. This keeps the different power sources at the same voltage. We achieve an additional benefit by using common worldwide power distribution solutions for the racks in the rooms, allowing 400/230 Vac to be used in Europe and 380/220 Vac in Asia.

- Fully rated electrical components and conductor sizing. For continuous duty, we size and rate our branch circuit breakers, busways, distribution breakers, disconnects, panels, conductors, and rack power strips at the stated value on the component. This provides a significant capacity capability within the distribution design, in addition to providing a safety margin for peak loads. We follow all applicable local electrical codes and rules when specifying and installing electrical components.

Stranded Capacity

Stranded capacity is the available design capacity of the derated installed space, power, or cooling infrastructure minus the actual use or the unavailable capacity of one of these data center elements due to the limitations of another interdependent element. We proactively identify stranded capacity, such as power and cooling capacities, in all areas of the data center. Harvesting stranded capacity can help us avoid having to make expensive changes to the existing infrastructure.

HARVESTING STRANDED IN-ROOM COOLING CAPACITY

As server density goes up, the rack power and heat density also increase, creating a need for greater airflow or cooling in a specific area of the room. The total demand on the cooling system is the same but the distribution is wrong. We have found that by simply moving a CRAC unit to another location, we can deliver the once-stranded CRAC capacity to a specific area of the room where needed, thus eliminating the stranding.

A congested or undersized RMF can strand the available cooling capacity of the room, measured by total CFM delivered at a practical static pressure. To avoid the loss of capacity caused by under-floor congestion or undersized RMF for new densities, we raise the room cooling units above the surface of the RMF—essentially equivalent to a flooded supply air design on grade—or relocate under-floor cabling into overhead trays.

HARVESTING STRANDED CHILLED WATER CAPACITY

The chiller plant may have the required capacity for the data center's cooling requirements but limiting factors, such as an undersized pipe or pump, may contribute to stranded chilled water capacity, or the installed CRAC capacity may be less than the capacity of the chiller or line size. If the supply and return water temperature Delta T is less than the chiller equipment design, the chillers true capacity is not available. We look at these types of factors to determine whether any chilled water capacity is being wasted and then modify installations accordingly.

Data Center Heat Recycling Systems Cut Business Cost

At an Intel data center in Russia, Intel developed and installed a heat recycling system (HRS) that reuses the heat generated by servers in a data center to warm office space during cold months, as well as to heat water for restrooms and cafeterias. In recognition of this project's innovative approach to energy savings, the team was awarded the bronze Intel Environmental Award. A similar HRS system was installed at an Intel data center in Israel.

The Russian project required the IT data center team and Intel Corporate Services to work closely together. Although the system does not currently use 100 percent of the heat generated by the data center, the proof of concept was successful. In its first phase, the HRS in Russia is helping Intel save about USD 50,000 annually. The IT data center team is now working on the second phase of the HRS project, which will use the full amount of heat generated by the data center servers. We anticipate this will save Intel about USD 300,000 per year, or up to USD 1 million over a three-year period. In a larger data center, such a system has the potential to generate even more savings.

HARVESTING STRANDED UPS CAPACITY

Electrical load phase imbalance at several components can limit or strand total conditioned electrical capacity available to data center devices. These components include UPS, distribution breakers, room PDU, and rack-level circuit breakers. We attempt to maintain a phase value difference of 5 to 15 percent between phases, as our experience has shown that this range provides an efficient use of conditioned power and prevents circuit breaker trips or UPS by-pass events.

These same components, if undersized for the load, can also limit or strand downstream electrical capacity to the data center devices. We carefully size these components to provide all the load the system was intended to deliver, with additional capacity for temporary downstream load changes.

We also carefully construct our redundancy topology so that it is in balance with the UPS, distribution board, room PDUs, branch circuits, and rack power strips. The goal is that no single component or circuit limits the capacity of the overall electrical design during a failover event. This is important because failover events typically put additional load on the system during a power transfer.

CONCLUSION

As we continue to consolidate our data centers, proactively refresh servers and storage with newer systems based on the latest generation of Intel Xeon processors, and further increase the virtualization of our Office and Enterprise applications for our enterprise private cloud, the density of our data centers is steadily increasing.

We focus on several areas to optimize the efficiency and reliability of Intel's high-density data centers. These areas include air and thermal management, architectural and electrical considerations, and the harvesting of stranded cooling and power capacity.

Focusing on these design areas enables us to increase the room infrastructure capacity—thereby enabling higher densities and reducing operating cost while still maintaining the levels of redundancy required by our business models.

For more information on Intel IT best practices, visit www.intel.com/it.

CONTRIBUTORS

Ajay Chandramouly
Industry Engagement Manager, Intel IT

ST Koo
Data Center Architect, Intel IT

Ravindranath D. Madras
Data Center Enterprise Architect, Intel IT

Marissa Mena
Data Center Architect, Intel IT

Isaac Priestly
Strategic Financial Analyst, Intel IT

ACRONYMS

AHU	air handling unit
CFD	computational fluid dynamics
CFM	cubic feet per minute
CRAC	computer room air conditioning
CRAH	computer room air handler
Delta T	change in temperature
fpm	feet per minute
HAE	hot aisle enclosure
HRS	heat recycling system
kW	kilowatt
PDU	power distribution unit
PF	power factor
PSU	power supply unit
RAP	return air plenum
RMF	raised metal floor
STS	static transfer switch
UPS	uninterruptible power supply
Vac	voltage alternating current
VFD	variable frequency drive