# Case Study: San Diego Supercomputer Center

***SDSC saves energy, money, and critical data with new datacenter cooling design.***

## Keeping Big Data Cool

Where do researchers turn to analyze fluid dynamics during the design and testing of hypersonic aircraft?

How will medicine eventually become personalized treatment and prevention plans based on a patient's genome sequencing?

What will researchers use to map the brain and unlock how it functions, so they can defeat diseases like Parkinson's and cure disorders like epilepsy?

How will they co-analyze data on greenhouse gases and ocean currents to understand global climate changes?

Will researchers eventually be able to predict earthquakes? Or someday discover how stars interact, find hidden galaxies, and even study the evolution of the universe over time?

At the **San Diego Supercomputer Center** (SDSC), these applications and many others are being solved right now, using high-performance computing running on cluster-based supercomputers.

SDSC is an organized research unit of the University of California, San Diego (UC San Diego). When most people think of a supercomputer center, they may think of one large computer, performing a single task. But inside the datacenter at SDSC live several large supercomputer systems, each performing many tasks and being operated by many different researchers at any given time.

SDSC enables international science and engineering discoveries through advances in computational science and data-intensive, high-performance computing.


SDSC is considered a leader in data-intensive computing, providing resources, services, and expertise to the national research community. The supercomputer center's mission is to extend the reach of scientific accomplishments by providing tools such as high-performance hardware technologies, integrative software technologies, and deep interdisciplinary expertise.

SDSC pursues research in the areas of high-performance computing, grid computing, computational biology, geoinformatics, computational physics, computational chemistry, data management, scientific visualization, and computer networking.

### Getting Useful Data from Big Data

According to IBM, "Every day, we create 2.5 quintillion bytes of data—so much that 90% of the data in the world today has been created in the last two years alone. This data comes from everywhere: sensors used to gather climate information, posts to social media sites, digital pictures and videos, purchase transaction records, and cell phone GPS signals, to name a few. This data is big data."

Sorting through and obtaining value from this vast amount of data can be a challenge for several reasons.

- Most of the big data created today doesn't have much inherent value but regardless is logged and warehoused in databases across the world. Because there's so much in them, these large repositories of collected data obscure the valuable data researchers are trying to get to.

- Big data is also generated very quickly, and the pace at which it's being generated is constantly increasing.
- Finally, big data comes from many different sources and is transmitted and stored in a variety of formats.

These are the challenges researchers face today as they try to solve complex problems using high-performance computing.

Sorting through massive amounts of data in research projects and filtering out useful data from inconsequential or erroneous data takes massive amounts of computing power. Ordinary PC architecture simply doesn't have the computing power to process big data as we know it today. A different kind of computing system is required.

## Measuring Computing Power

Computing power is measured in FLOPS, an acronym for floating-point operations per second. In layman's terms, FLOPS show how fast a computer can calculate numbers with decimals.

The number of equations a computing system can execute per cycle (FLOPs) is based on the system's hardware architecture and calculated using the equation below.

$$FLOPS = Sockets \times \left(\frac{Cores}{Socket}\right) \times Clock \times \left(\frac{FLOPS}{Cycle}\right)$$

To make sense of big data and apply it to solve complex problems, SDSC supercomputers use a system architecture drastically different from a standard PC. While the modern PC may have two or four core processors that make up its central processing unit (CPU), high-performance computing requires significantly more processing power and memory.

SDSC's latest supercomputer, *Comet,* has a total of 47,776 Intel Haswell cores and 247 terabytes of memory. *Comet's* peak performance is around 2 petaflops, or about 2 quadrillion floating point calculations per second.

A computing architecture like the one found in *Comet* requires one megawatt of power to operate the system. Using that much electricity generates a tremendous amount of heat in the SDSC datacenter.

All computers generate heat when operating. But the supercomputers at SDSC generate substantially more heat than typical PC systems.

**When you're trying to protect millions of dollars' worth of research, you need a control system you can rely on.**

**- Todor Milkov, Senior Project Engineer, SDSC**

Keeping the SDSC datacenter cool so *Comet* and the other supercomputer systems at SDSC don't overheat is a complex and mission-critical task, says Todor Milkov, Senior Project Engineer at San Diego Supercomputer Center.

## The Challenge

The supercomputers at SDSC are used by a wide variety of research organizations across the country for analysis. Not all of the systems are running at full capacity at any given time. Systems installed at the supercomputer center also change regularly as research projects start up or stop.

These constantly changing factors cause the amount of heat dissipated from the supercomputer systems to fluctuate from minute to minute. The datacenter cooling system has to quickly adjust to accommodate these fluctuations in temperature.

Many datacenters use a standard hot aisle/cold aisle design. This design involves lining up server racks in alternating rows, with cold air intakes facing one way and hot air exhausts facing the other:

- The rows composed of rack fronts are called cold aisles. Typically, cold aisles face air conditioner output ducts.
- The rows the heated exhausts pour into are called hot aisles. Typically, hot aisles face air conditioner return ducts.

Containment systems can help isolate hot aisles and cold aisles from each other and prevent hot and cold air from mixing. Containment systems started out as physical barriers that simply separated the hot and cold aisles with vinyl plastic sheeting or Plexiglas covers. Modern containment systems offer plenums and other commercial options that combine containment with variable fan drives (VFDs) to prevent cold air and hot air from mixing.

At SDSC, however, the entire area under the raised floor is used for the supply plenum, and the entire area above the ceiling for the return plenum:

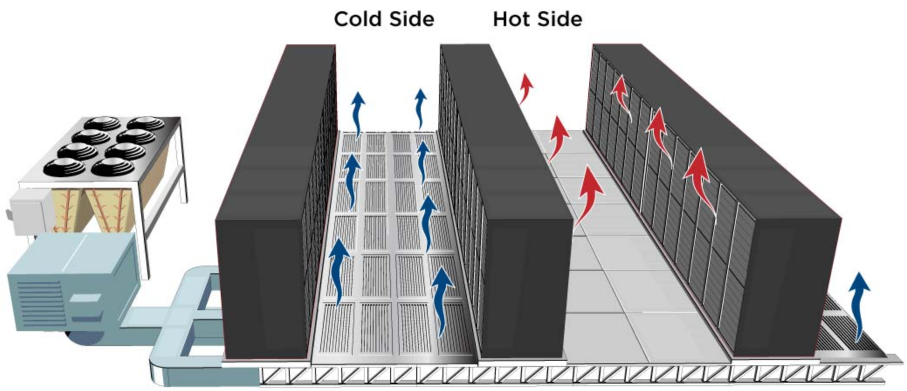


*SDSC's datacenter uses a modified hot aisle/cold aisle design for racks.*

- Cold aisles use perforated floor tiles with designed hole sizes to control the air flow volume from the space below the floor.
- Hot aisles use ceiling grates that allow heated air to enter the space above the ceiling.

Controlling the air flow from all air handlers discharging into one common plenum presents a difficult problem, especially since these spaces also contain obstructions such as pipes and conduit.

### Cooling system research

To design a cooling system that could handle the massive amount of heat SDSC's supercomputers generate, the center developed three cooling system prototypes and conducted research to determine the most efficient system.

Each prototype system was designed using vendor-specific technology controlling five air handlers, as a baseline to evaluate system performance.

One of the prototypes used wireless temperature sensors that read the temperature of the hot and cold aisles every three minutes to increase battery life.

"We learned a lot during the prototype and research phase of the cooling system design," says Milkov. "First we started by collecting a lot of data on how air flowed through the datacenter. We found that three minutes between temperature readings was too long an interval to keep the datacenter within the desired temperature ranges.

"Because of the longer interval, we used more electricity bringing the datacenter back to its temperature setpoints than we needed to if we took temperature readings over shorter intervals and could make changes to the cooling system sooner, before it got too far out of its setpoints."

## The Solution

When the facility began looking for ways to reduce energy and increase capacity through improved control, Todor Milkov was tasked with finding solutions.

His research began to form a picture that's standard in most industries: many companies were using the same basic and repeated datacenter design philosophies and just using different methods and equipment to get there.

Realizing that the research being conducted at SDSC was far from standard and that perhaps a different approach was needed, Milkov put together a vendor evaluation process for an updated datacenter management system, with the objective of reducing energy use and increasing the level of control capability available to the SDSC operations staff.

After extensive research, Milkov selected three companies for prototype installations. At the conclusion of a detailed evaluation, systems integration company Earth Base One (EBO) Corporation and a SNAP PAC-based control system were chosen for providing extensive control capabilities and superior energy savings.

Milkov and Michael Hyde of EBO approached the project with the same vision. "Rather than adapting an off-the-shelf datacenter management system to SDSC, we designed a tailor-built system for SDSC's unique challenges," said Hyde. "Opto 22 was chosen as the primary controls manufacturer because of their quality and history

**After developing three prototype systems, we chose the SNAP PAC System from Opto 22 for its robust flexibility, high quality, and the outstanding technical support SDSC received during our design and prototype phase.**

**- Todor Milkov, Senior Project Engineer, SDSC**

of building what their customer needs, rather than building something and then trying to find a market for it."

"The Opto 22 hardware and software not only won the competition for control and energy savings but was also the least expensive vendor solution." Hyde noted. "The software's excellent historical data collection and trending abilities have allowed Mr. Milkov to continue improving the system based on real data."

"We appreciated the outstanding technical support SDSC received from Opto 22 during our design and prototype phase," said Milkov. "When you're trying to protect millions of dollars' worth of research, you need a control system you can rely on."

## Control System Details

The SNAP PAC control system at SDSC consists of:

- 6 rack-mounted SNAP-PAC-R1 controllers
- 3 standalone SNAP-PAC-S2 controllers
- 19 SNAP-PAC-EB1 Ethernet I/O brains
- 1 *groov* Box
- 2 computers
- 22 I/O panels

In total the SNAP PAC System operates 1,593 Opto 22 I/O points and approximately 400 Modbus points from third-party power distribution units (PDUs).

Using the SNAP PAC System allowed SDSC to take monitoring and control of the datacenter to a whole new level. Hot and cold aisle temperatures are monitored by the control system through temperature sensors wired to SNAP-AICTD modules.

Using PAC Display HMI software from Opto 22 allowed SDSC to develop easy-to-understand HMI screens that give operators overall situational awareness of everything going on with their datacenter cooling system. SDSC's HMI screens provide a clear overview of cold and hot aisle temperatures as well as air pressure levels within the datacenter.

"Many variables go into controlling the temperature of the datacenter," says Milkov. "We prefer to cycle air very slowly through the datacenter to reduce air turbulence. Air turbulence wastes energy and reduces the ability to control the overall temperature of the datacenter by causing hot spots in specific datacenter locations."

To remain at an optimum operating temperature for the supercomputers to run, the datacenter uses 17 separate air

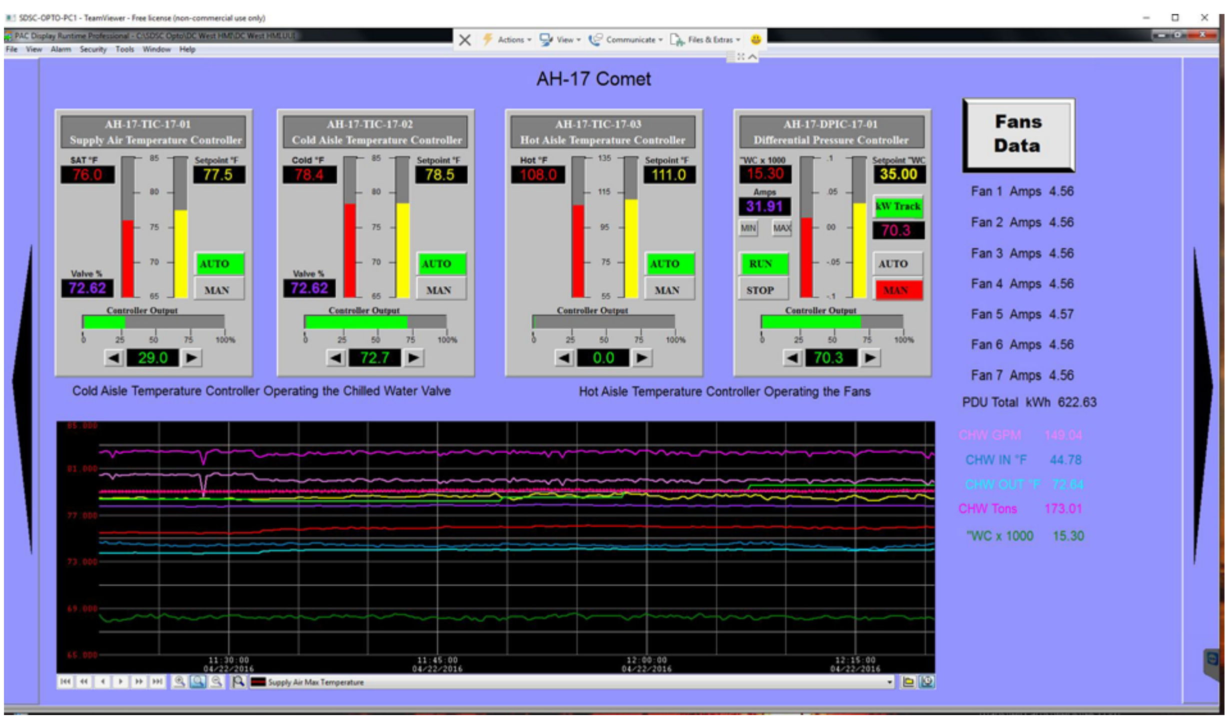


*The PAC Display HMI at San Diego Supercomputer Center indicates air handler status. Yellow is slowing; red is stopped; orange is starting up.*


OPTO 22

handlers. Controlling each air handler and the flow of chilled water used to cool air sent to the datacenter requires constant temperature monitoring of the hot and cold aisles. Accurate placement and sizing of the floor tiles is also a critical part of the system.

SDSC keeps air handlers running at a minimum level to keep the air pressure in the datacenter slightly above atmospheric pressure to keep dust out of the facility.

"With the Opto 22 system we were able to quickly and accurately measure the changes in air pressures and temperatures and supply only the amount of air flow and cooling required for the current server loads," says Hyde.

SDSC also monitors and color codes the amperage each air handler in the cooling system is pulling to determine if the air handler is slowing down, speeding up, or has stopped.

Changes in any of these values can help identify parts of the cooling system that may be wearing out, so operators can replace system components before a failure occurs. All of the air handlers are belt driven, so an air handler slowing down may indicate a belt is about to fail.

The datacenter cooling system is controlled at an independent air handler level but also has a built-in single overall supervisory system. Separate zones are set up within the datacenter, and the air handlers work together using the SNAP PAC System to keep all zones within configured setpoints.

The control system can be operated manually or act autonomously. The SNAP PAC System uses multiple PID loops for monitoring and controlling temperatures in the datacenter. Using preconfigured alarming conditions, if temperatures get above configured setpoints, the SNAP system controls VFDs, individual air handlers, and chilled water valves with cascade and high/low selections of inputs and outputs, as well as overlapping zone control.

The control system has intelligence built into it as well. Using SNAP-AIPM-3V modules, the control system monitors how much power the *Comet* supercomputer is pulling and automatically ramps up cold air production by adjusting chilled water flow and speeding up VFDs to preemptively compensate for temperature changes before they occur.

The SNAP PAC System monitors other supercomputer power usage by interfacing with over 400 power distribution units (PDUs) using the Modbus/TCP protocol.

With built-in PAC Display alarming features, when an alarm occurs, operators receive instructions on what to check in the cooling system to pinpoint the source of problems faster and keep systems running.

The SNAP PAC System and PAC Project software suite allowed SDSC to get a detailed view of exactly how the datacenter cooling system was operating. As a result, SDSC was able to optimize its datacenter cooling system, reducing chilled water use from 2,000 GPM to 330 GPM and reducing power use from 300 KW to 79 KW.

The new control system saved UC San Diego $168,000 per year in energy costs.

## Looking Ahead

Several system components at an adjacent datacenter currently use a Johnson Control building automation system. Milkov plans to retrofit that datacenter with an Opto 22 SNAP PAC System for increased system performance and faster monitoring and control response times.

## About SDSC

The San Diego Supercomputer Center (SDSC) is considered a leader in data-intensive computing and cyberinfrastructure, providing resources, services, and expertise to the national research community, including industry and academia. Cyberinfrastructure refers to an accessible, integrated network of computer-based resources and expertise, focused on accelerating scientific inquiry and discovery. SDSC supports hundreds of multidisciplinary programs spanning a wide variety of domains, from earth sciences and biology to astrophysics, bioinformatics, and health IT. SDSC is a partner in XSEDE (Extreme Science and Engineering Discovery Environment), the most advanced collection of integrated digital resources and services in the world.

## About Earth Base One

Earth Base One (EBO) Corporation was formed in February 1991 to provide technical services in the commercial and industrial market. EBO's mission is to provide their clients with the technical expertise required to enhance their business through energy management and waste reduction. They accomplish this by augmenting clients' staff with EBO technical personnel and expertise. When clients have a technical problem waiting to be converted to an opportunity, and their staff has yet to gain experience in this particular technical area, then EBO provides assistance as needed. EBO has been successful with this approach on projects ranging from water softeners and air handlers to turbine co-generation plants.

## About Opto 22

Opto 22 develops and manufactures hardware and software products for applications in industrial automation, remote monitoring, and data acquisition. Using standard, commercially available Internet, networking, and computer technologies, Opto 22's input/output and control systems allow customers to monitor, control, and acquire data from all of the mechanical, electrical, and electronic assets that are key to their business operations. Opto 22's products and services support automation end users, OEMs, and information technology and operations personnel. Founded in 1974 and with over 85 million Opto 22-connected devices deployed worldwide, the company has an established reputation for innovation, quality, and reliability.

Opto 22 products are sold through a worldwide network of distributors, partners, and system integrators. For more information, contact Opto 22 headquarters at 800-321-6786 (or 951-695-3000) or visit the website at www.opto22.com.


Form 2183-160621 • Opto 22 • 43044 Business Park Drive • Temecula, CA 92590-3614 • **www.opto22.com**
**SALES** 800-321-6786 • 951-695-3000 • FAX 951-695-3095 • sales@opto22.com • **SUPPORT** 800-835-6786 • 951-695-3080 • FAX 951-695-3017 • support@opto22.com