The Clear Solution for Data Center Cooling™

Extreme cooling performance with low-cost dielectric fluid submersion cooling

How a top-10 supercomputing facility begins to rethink its cooling efficiency.

Texas Advanced Computing Center, Austin, TX, April 2010

ABSTRACT: When researchers at Texas Advanced Computer Center purchased a CarnotJet™ 42U end-to-end cooling system unit from Green Revolution Cooling in April 2010, they were surprised to find that it offered an 85% reduction in cooling power when compared to their already-efficient air-cooled system. TACC is currently planning a larger install for 2012.

Texas Advanced Computing Center (TACC) is one of the world's leading scientific computing and research centers, in large part because it houses Ranger, the largest open research supercomputer in the world. When TACC switched on Ranger in early 2008, it paved the way for the next generation of scientific research and advancement by bringing online half a petaflop (one thousand-trillion operations per second) of computational capability.

For more than two years, researchers from all over the world have accessed Ranger's petascale resources through TeraGrid, a nationwide scientific computing infrastructure. TACC's Linux-based Sun Constellation cluster has been utilized for projects as diverse as sequencing parts of the human genome and mapping the path of the Gulf Coast oil spill to developing new smart-memory metals and tracking particular inner-ear mutations that can cause deafness.

With 3,936 nodes, 15,744 quad-core processors, 123 terabytes of memory, and 579.4 teraflops of shared storage space, Ranger can support a lot of groundbreaking research. But a supercomputer of this magnitude also requires a lot of energy to keep cool.

Despite the fact that leading edge data centers like TACC have been getting more efficient through the use of clever air-manipulating schemes—in-row cooling, raised floors, hot aisles and cold aisles, sophisticated temperature gauges—air cooling remains inherently inefficient. Even with a committment to saving energy and low Texas energy prices, TACC spends hundreds of thousands of dollars a year to cool its facilities.

> *As server density continues to increase, so too will the need for more efficient cooling solutions.*

So these gains in efficiency have all been relative. Data center energy expenditure now accounts for nearly 3% percent of total U.S. electricity use, up from 1% in 2000. As server density continues to increase, so too will the need for more efficient cooling solutions.

In 2009, we approached TACC with a simple question: How much cooling efficiency could be gained by replacing your air-cooling system with GRC's fluid-submersion technology?

Our mission was simple: Prove to TACC that our CarnotJet™ fluid-submersion solution outstrips even the most efficient air-cooling systems when it comes to data center energy efficiency. The question proved too good to leave unanswered, TACC was eager to see how the system performed.

Installation

On April 15, 2010, we installed a 42U CarnotJet™ cooling system in TACC's facilities. The installation, the first of its kind in the world, focused on independently verifying three key items:

- Cooling power efficiency
- Ability to accept full range of form factors
- Potential for over-clocking servers

The unit was built at GRC's headquarters in Austin and transported to TACC where it was installed on a concrete pad in a loading dock completely independent from the existing data center infrastructure. One week after loading-dock delivery, our system was up and running at full capacity.

The space features no chilled water, CRACs, or climate control of any kind. Instead of air conditioning, the system rejects heat through an evaporative cooling tower outside that uses very little power. Warm coolant exchanges heat with water provided by the tower.

The Rack occupies about thirteen square feet and weighs about 250 pounds per square foot (PSF), a weight/sq. ft. similar to that of a typical air-cooled

Figure 1. CarnotJet Evaluation Unit Installed at TACC

rack. The CarnotJet's™ Rack sits on portable secondary containment that prevents any incidental coolant splashes from causing slippery floors.

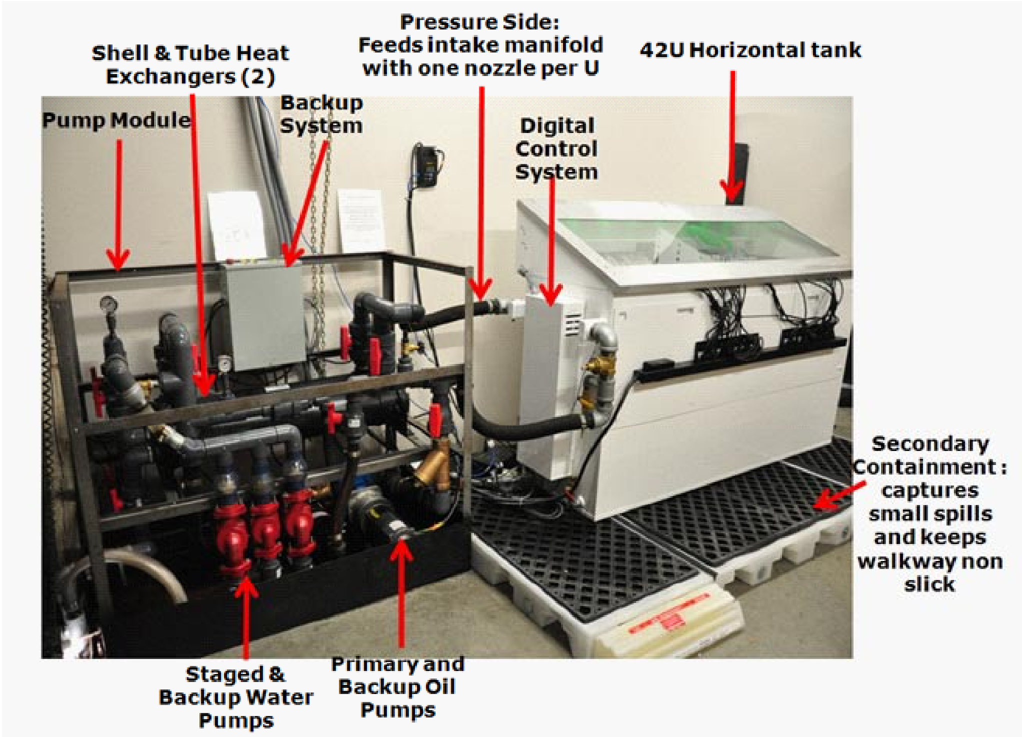

Figure 2. Details of the Rack and Pump Module

There are very few requirements in general for installation of the system: a level surface, space for the Rack and the Pump Module, and access to a heat-exchanging agent such as outside air or a water loop with water less than 30° C. There is no need for raised floors, chillers, or air conditioners and there are no concerns about room temperature or humidity.

Table 1. Installation Specifications

Install Location	Loading dock
Installation Time	One week (after delivery)
Room Climate Control	Unconditioned
Rack Size	42U
Rack Dimensions	13 square feet
Load Capacity (Pump)	40 kW
Secondary Containment	Modular
Heat Exchange Method	Evaporative cooling tower

System Design

The CarnotJet™ system consists of three components: the Rack, the Pump Module, and the Control Module. Typically 42U, each Rack has the same capacity as an industry standard vertical cabinet.* Each Pump Module supports up to four Racks simultaneously, so expansion typically occurs in increments of four.

**Racks are also available in 13U and 60U models.*

The Rack typically holds about 250 gallons of our GreenDEF™ fluid coolant. The coolant is a dielectric (non-conductive) formulation of mineral oil, meaning it has no negative interaction with electronics. In fact, it's quite the opposite: GreenDEF™ has 6x the dielectric strength of air! The true utility of GreenDEF™ is that it retains 1,200x more heat by volume and is many times better at convective cooling than air. Additionally, GreenDEF™ is non-toxic.

Before servers can be installed in the Rack, there are three simple modifications that are required: remove server fans, encapsulate hard drives, and replace thermal grease on heat sinks with an insoluble material.

Once this process is completed, servers may be installed and uninstalled from the Rack very easily. The horizontal rack system that the CarnotJet™ employs is useful for quick access to servers—when the need arises to replace parts or upgrade components, it's as simple as lifting a server from the liquid and setting it on the service rails above the Rack. Replacing a stick of internal DIMM, for example, takes less than sixty seconds on average. Coolant drips back into the Rack and the server is available for immediate modification.

The Rack is capable of accommodating virtually any 1U-4U server, blade, or 19" server, as well as all necessary switches. The design also provides ample space for all data and power supply cables, including cables with large bend radii such as Infiniband. Separate routing for power and Ethernet minimizes electromagnetic interference and "cross talk."

> *"Since the installation, the system has consistently shown energy savings approaching 40% of total power usage."*

The Pump Module [Figure 2], creates liquid circulation in the rack by natural and forced convection. Hot coolant is pumped out of the Rack, circulated through the Pump Module, cooled, and returned to the Rack. The Pump Module also features 2N redundancy: the primary system controls a set of variable speed pumps with a digital controller. An electrically isolated analog secondary system with a separate set of pumps serves as a backup system in the event the primary system goes down, and each unit is capable of running the entire system independently.

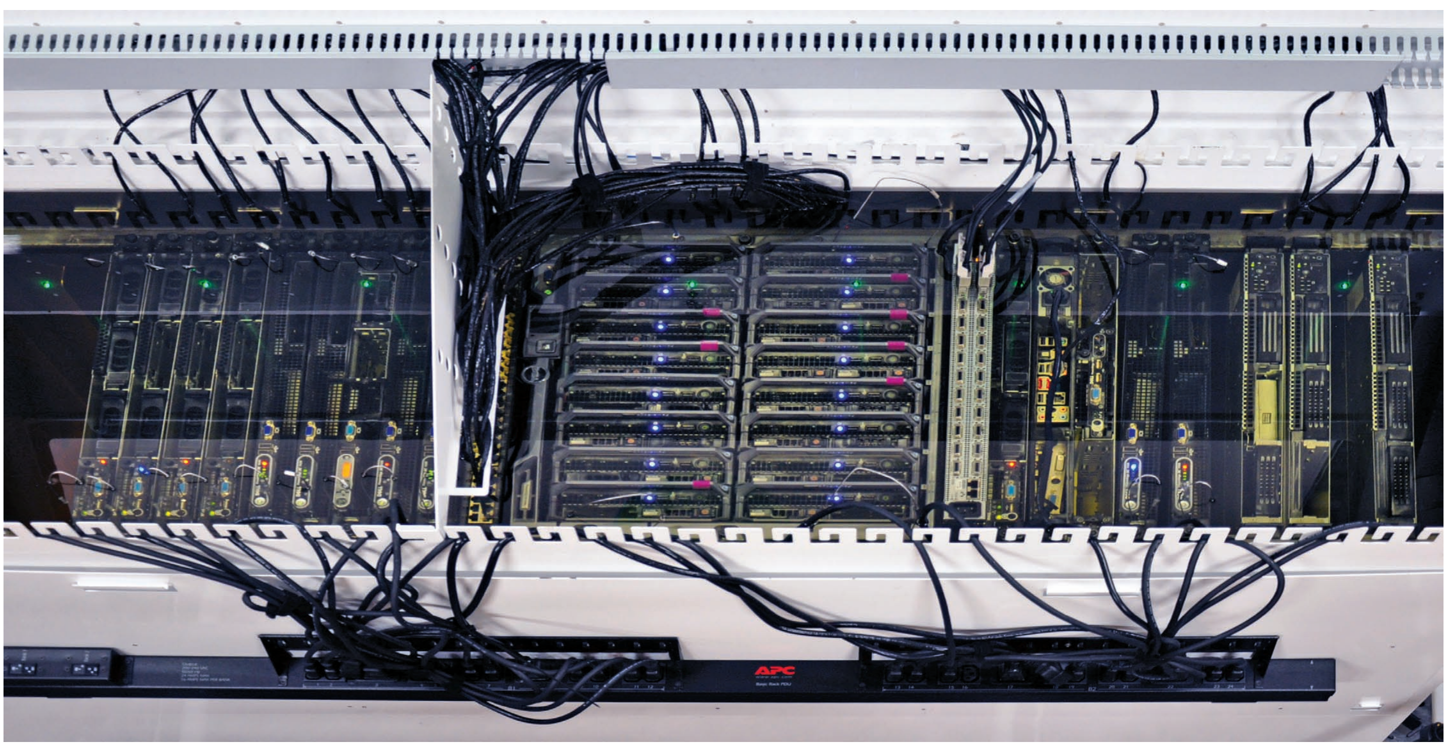

Figure 3. Servers installed in TACC's CarnotJet system

An extensive Software Control System takes measurements of the system, issues commands to the cooling system components, and stores diagnostic data in an easily accessible log file. In addition, hourly status updates as well as diagnostic alerts are sent automatically via email.

System Details:

Figure 3 shows the servers installed at TACC, excluding additional servers that were added more recently:

- » Dell M1000e system chassis with 16 dual-socket blades with 12 GB per CPU
- » Dual-socket 1U servers
- » Two switches

The Rack will accommodate any OEM (Original Equipment Manufacturer) server and can support in excess of 100kW per unit.

Results:

Timing of Test and Uptime:

The CarnotJet™ system has run continuously since April 15, 2010 with 100% uptime.

Efficiency:

Since the installation, the system has consistently shown reductions in tototal power usage (i.e. servers and cooling) of ~35% as compared to the same hardware cooled by air. Power consumption savings are as follows:

- » Cooling System Power: Our system at TACC has consistently used 6 watts or less of cooling power for every 100 Watts of IT since April 2010.

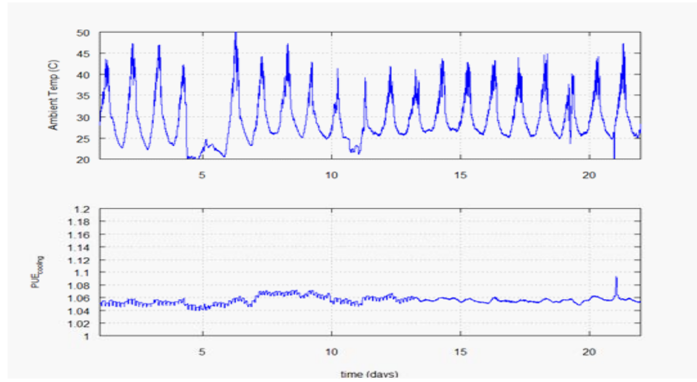

Figure 4. Measured system efficiency compared with ambient outdoor temperature

- » Server power consumption decreased by a larger percentage than expected. When compared to servers that were tested at room temperature with internal fans, the servers that were submerged in fluid coolant regularly reduced energy usage by 10-20%. Some liquid-cooled servers reduced energy use by more than 20%. To test actual fan power consumption, we ran the server at 75° F ambient temperature using Linpack for 100% CPU utilization, and then we removed the fans and ran the same program and server in our system.

Table 2, below, displays some of the specific power reductions we were able to achieve.

Table 2: Measured Server Power Reduction

Server Model	*Power Reduction*
IBM x3550 1U server	10.3% at 100% CPU
Dell 1950 1U server	27% at idle

Savings became even more significant when we compared air-cooled servers running in elevated ambient temperatures with servers installed in the CarnotJet™.

As shown in Table 3, below, air-cooled systems are heavily taxed when ambient temperature increases.

Table 3: Increase in server fan power use as a function of increase in ambient temperature, as compared to the CarnotJet™ system. Baseline Server: New Dell 1Userver (R410), one 4 core CPU (E5520), 4 GB RAM, 1 HDD.

	Idle		100% CPU	
	W	% Δ	W	% Δ
CarnotJet	67.9		126.6	
air @ 27 C	68.9	1%	137.7	9%
air @ 29 C	71.5	5%	139.4	10%
air @ 33 C	75.8	12%	144	14%
air @ 35 C	80.5	19%	148.4	17%

On a total energy consumption basis, we use Table 4, below, to display actual savings in power consumption recorded at TACC. Typically, data center efficiency is expressed as total power/server power (PUE), but since the servers in our system also use ~15% less power (due to fan removal), the ratio doesn't fully reflect our energy savings.

Table 4: TACC Total Data Center Efficiency

Units = $\frac{\text{Power}}{\text{Server Power in an Air-Cooled System}}$

System Element	*Current System*	*GR Cooling*
Server	1.0	.85
Power Distribution Losses	.02	.02
Cooling	.4	.06
Total	1.42	.93

Server Reliability

Since installation, server reliability has been carefully monitored. So far, no server failures have been observed except for hard drive malfunctions resultant from the installation of old and used drives. Generically, a few factors improve the reliability of servers submerged in dielectric fluid coolant:

» Decreased number of heat-related failures through removal of dust (a thermal insulator), decreased CPU temperatures of 10° to 20° C, and stable coolant temperature (controlled to +/- 1° C).

» Reduced localized heating due to superior thermal capacity of coolant (GreenDEF™ has 1,200x higher capacity for heat by volume than air).

» Decreased connector oxidation and corrosion and a higher dielectric strength (12 vs. 2 MV/Meter) prevents micro-arcing. This should reduce or eliminate "reseat errors" (caused by oxidized connections) which require the component to be removed and reinstalled in order to function properly.

Additional Benefits

The CarnotJet™ system's technology should help with the following sources of server component failure, here listed by frequency of occurrence for a typical data center:

» Hard Drives: Hard drives are susceptible to vibration from fans and other sources found in a data center. Several studies document the deterioration of performance due to fan vibration, especially with random read-writes and with consumer grade hard drives.

» Power supplies: Engineers that design OEM power supplies report that improved component cooling should increase internal component reliability, specifically with FETs and other high-temperature components.

» Memory: We expect the lack of corrosion issues at the connectors to significantly improve reliability.

» Fans: Completely eliminated.

Total Cost of Ownership

The total savings is nearly $100,000 for a single 22kW server rack. (See Table 5.)

» 85% less cooling energy compared to a very efficient air-cooled facility.

» ~35% less total energy than the same efficient air-cooled facility.

» Lower infrastructure cost.

Conclusion

Before the installation, we knew that a successful installation would demonstrate the stability and reliability of our system while exhibiting the superior cooling capability of the design. We feel that we accomplished this and more. TACC continues to report that the CarnotJet™ system is living up to and exceeding expectations. The Deputy Director at TACC, Dan Stanzione, Ph.D., has responded to GRC with high praise for the CarnotJet™:

> *"I'd like to express our delight with the progress thus far in our project with GRC. The results have been extremely impressive….The total power used to cool the system was reduced more than 85%. Over the course of the last year, we have tried numerous configurations of hardware, including state of the art high density blade systems. In every case, the efficiency and reliability data have been impressive. We consider this a key component of a future Exascale strategy at 100 cabinets and beyond."*

As we make serious progress with our CarnotJet system, we look forward to our continued partnership with TACC and to working on a larger install planned for early 2012.

At Green Revolution Cooling, we hope to bring this technology to the forefront of the data center efficiency conversation. We believe that we can implement world class liquid-cooling solutions that are far more efficient and cost-effective than traditional air-cooling techniques. As we continue to grow, we're eager to partner with other innovative institutions and data centers to help shift the cooling paradigm.