



# WHITE PAPER



## **Oil Immersion Cooling for Today's Data Centers**

**An analysis of the technology and its business**

In response to a growing demand for computing power and technological advancement, server densities have increased rapidly across industry. Oil immersion cooling provides groundbreaking technology in an industry that has struggled to innovate with demand.

# Table of Contents

Introduction	1
Data Center Cooling Technologies	1
How Immersion Cooling Works	2
What This Means for Operators	5
Lower PUE	5
Lower Energy Draw	5
Carbon Reduction	6
Lower OPEX+CAPEX	7
Increased Reliability	7
Granularity Helps Capacity	8
Monitoring and Detection	9
GRC Dashboard	10
Proven Results	11
Conclusion	11
References	12
Glossary	12

## Introduction

The past few years have seen an exponential increase in global data center traffic. It is estimated that this trend will grow threefold (a 25 percent CAGR) from 2012 to 2017, while global cloud traffic will grow 4.5-fold (a 35 percent CAGR) over the same period<sup>1</sup>.

In response to the growing demand, we have seen a rapid increase in computer densities following Moores law. However, one critical area that has not managed to keep pace with this growth is data center cooling technology. This mismatch, clubbed with the growing challenges of energy and environmental sustainability, make data center cooling not just an important concern but an urgent issue to tackle.

## Data Center Cooling Technologies

There are lots of precision cooling solutions out there: aisle containment, rear door cooling, free cooling, liquid to the chip cooling, etc. None of these novel cooling strategies are as simple and industrial as immersion cooling.

## Oil Immersion Cooling

Oil immersion cooling, as the name suggests, involves immersing entire servers in a dielectric coolant. Such methods of cooling have been used for nearly a century in electrical transformers and industrial capacitors. Similar applications of fluid immersion cooling in the world of HPC date back to as early as the 1980s<sup>2</sup>, but concerns regarding the cost and safety of these coolants have led to limited adoption across the industry.

Green Revolution Cooling's revolutionary CarnoJet cooling system addresses these issues by taking advantage of a special non-proprietary mineral oil blend, called ElectroSafe as the coolant.

## The Coolant: ElectroSafe

The ElectroSafe coolant is a non-toxic, clear, odorless, dielectric mineral oil blend. Dielectric oils have a history of being used across a variety of industries, from computing and electrical, to the medical field and domestic applications.

This high performance, low cost oil is readily available and is rated as w0-1-0 substance by the National Fire Prevention Association (NFPA), meaning that it poses no risk to human health, does not ignite readily, and can be treated with any fire suppression substance, including water<sup>3</sup>.

## How It Works

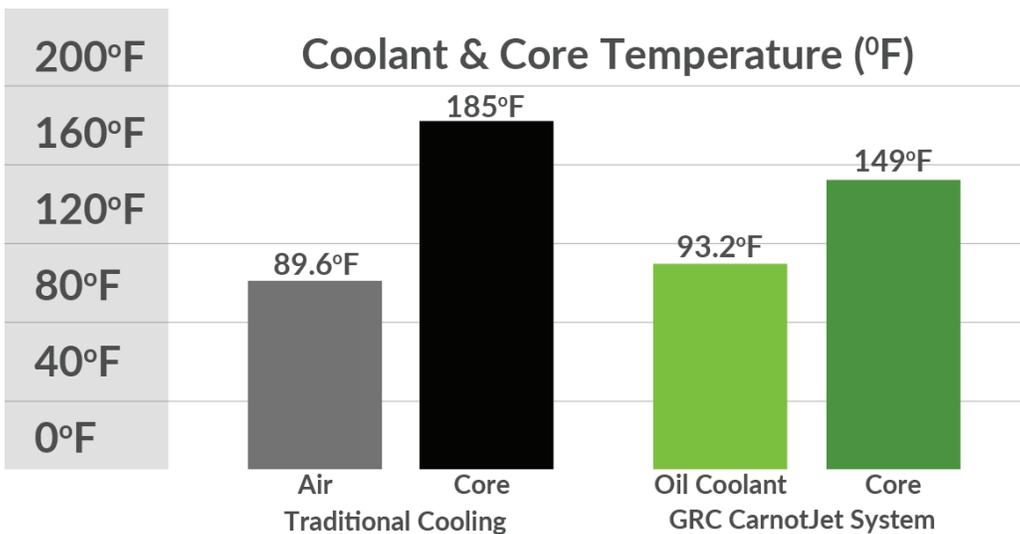
Green Revolution Cooling has brought to the data center market an industrial scale oil immersion cooling solution that works with servers from any manufacturer.

The technology consists of four major subsystems: the tank with a horizontal rack (42U), the pump module including the heat exchanger and strainer, the control system, and a cooling tower to extract heat from the system. The tank functions similarly to a rack turned on its back and immersed in a tub of circulating dielectric liquid. The liquid captures heat produced by IT equipment and dissipates in a heat exchanger. The open system tanks support servers manufactured by Dell, HP, Supermicro, IBM, Cisco, Intel, SGI, Bull, Quanta, and more.

## The Physics Behind It

### Air is a Poor Conductor of Heat

The primary advantage of a liquid immersion cooling system comes from the fact that liquid coolants are much better conductors of heat than air. This increased heat conductivity drastically improves the system's ability to both extract heat from the servers and subsequently expel that heat out of the system. Tests revealed that the CarnotJet system was able to maintain significantly lower CPU temperatures even with higher-than-air coolant temperatures.



*CPU operating close to 100% utilization in 89.6°F (32°C) air yielded an average core temperature of 185°F (85°C).  
When submerged in electrosafe at 93.2°F (34°C) the same CPU yielded an average core temperature of 49°F (65°C).*

The lower delta means that for the same target core temperature, the temperature set point for oil can be significantly higher than for air. Further, maintaining ElectroSafe at 93.2 °F (34°C) requires significantly less energy than chilling water to 55°F (12.8°C) for air conditioning.

### Higher heat capacity means higher ride-through time and higher server density

Fluid coolants typically have a significantly higher heat capacity than air, by volume. The ElectroSafe coolant has a heat capacity that is 1200 times that of air. This means that it would take 1200 times the amount of heat to raise the temperature of ElectroSafe by 1.8°F (1°C) as compared to air. As a result, in the case of a catastrophic failure the system can prevent equipment from reaching critical temperatures for longer. This additional disaster response time can be crucial in a mission critical environment.

### Lower Server Failure: Lower Chip Temperature, No Dust, No Moisture, No Oxygen, No Corrosion

Fluid immersion can significantly improve server life and performance through enhanced thermal management, improved electrical connections, and by blocking off dust, moisture, and oxygen. The CarnotJet system, through its highly efficient cooling mechanism, maintains lower surface temperatures throughout the IT equipment with a difference of less than 5.4°F (3°C) across any two points in the tank. The lower temperatures along with the elimination of hot spots is known to improve server life and performance. Further, the immersion in ElectroSafe blocks off dust, moisture, and oxygen from coming in contact with the equipment surface. The constant flow of the coolant along the surface also helps prevent any particulate accumulation and reseat errors. The ElectroSafe coolant can significantly reduce reseating errors due to its high dielectric strength, which is 6 times that of air (12MV/meter vs. 2 MV/meter). The use of dielectric compounds to increase connector reliability is well documented and common in other industries such as industrial automation and automotive. Another significant improvement in system reliability comes from the removal of the server fans, which are an extremely common source of failure in traditional cooling solutions. To learn more about server reliability, download the hardware reliability [case study](#).

## What This Means for Data Center Operators

The CarnotJet system puts all the advantages of liquid cooling in the simplest package on the market. Low infrastructure requirements and simple open-rack design means that the CarnotJet has extremely low capital requirements. Consistent 1.03 PUE operation means that the CarnotJet operating expenses of the data center are drastically lower than any other technology on the market.

Beyond straightforward value pricing, the CarnotJet system has proven to reduce failure rates in server components, reduce operating temperatures of IT equipment, and provide greater thermal ride through time for IT equipment in the event of cooling power loss.

### Lower PUE

The CarnotJet system boasts a PUE of 1.03, which is significantly lower than most free-air cooling data centers. The highly efficient system can have been tested to work efficiently with input water temperatures as high as 122°F (50°C) even in extreme heat conditions where ambient temperatures approach 135°F (57°C). Further, the return water from the CarnotJet system can be used for heat recapture to supplement domestic hot water and space heating needs

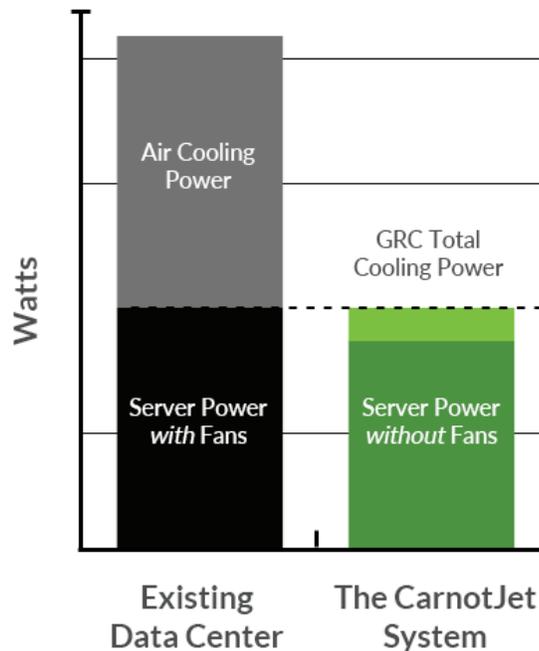
### Lower Server Energy Draw

Beyond the superior cooling performance of the CarnotJet system, the removal of the server fans along with the lower operating temperatures of the IT equipment facilitate lower power draw from the servers themselves. The overall reduction in server power draw has been observed to range from 10% to 20%, with certain systems yielding savings beyond 20%.

*The GRC system as a whole consumes less power than the server fans it turns off*



## Data Center Cooling + Server Power: Air-Cooling vs. Green Revolution Cooling



*CPU Operating close to 100% utilization in 89.6°F (32°C) air yielded an average core temperature of 185°F (85°C). When submerged in electrosafe at 93.2°F (34°C) the same CPU yielded an average core temperature of 49°F (65°C).*

### Carbon Reduction, CDP, and Grants

The CarnotJet system can help organizations not only to save money directly, but to open doors for energy efficiency grants and programs offered by local utilities, state and federal governments. The direct reduction in carbon emissions associated with cooling energy reduction translates into lower Scope-1 emissions, helping organizations reach their carbon reduction goals and achieve better scores on the CDP (Carbon Disclosure Project).

## Not Just Lower OPEX, but Lower CAPEX as Well

The CarnotJet system helps organizations not just reduce their operating expense, it also reduces the capital expenditure on new builds. Due to the higher rack density, the system takes up only 60% of the space required for a conventional air-cooled data center. Further, the lower power consumption means that the requirements for battery backup, generators, and power distribution are also significantly reduced, while larger CRAC and CRAH equipment are completely eliminated. In the case of a retrofit, the CarnotJet system frees up valuable power and water capacity to support more servers within the same data hall.

## Increased Data Center Reliability

The CarnotJet system simplifies the complete data center setup by eliminating a number of components such as server fans, CRAC and CRAH units, etc. This reduction in the number of components means fewer points of failure and higher levels of overall system reliability. In addition, the CarnotJet system comes with complete redundancy (2N). However, even in the case of a catastrophic failure, the ElectroSafe coolant, due to its high heat capacity is able to give extended response time in the form of power-free cooling. Each 275 gallon tank in the CarnotJet system provides 13 kWh of ride-through time, so a 12 kW rack would have over an hour of power-free cooling in the event of a catastrophic failure.

## Operations and Maintenance

The CarnotJet systems horizontal rack design has been optimized for ease of access and serviceability of the submerged equipment. Most traditional servers can easily be pulled out of the coolant and be placed on support bars located on top of the tanks. Heavier servers and blade chassis can be retrieved using assisted lifts. While this marginally increases the time taken per maintenance event, the reduction in the number of such events more than makes up for it.



*Watch the video on server maintenance in the CarnotJet System*

## Granularity that Helps High Capacity Planning

The current data center layout poses considerable challenges for capacity planning. Forecasting demand for computing power has proven to be an uphill task, and with millions of dollars on the line, forecasting errors can turn out to be extremely expensive. Data center Infrastructure planners need to make sure that they are not becoming the bottleneck for growth within the organization, while ensuring efficient use of capital.

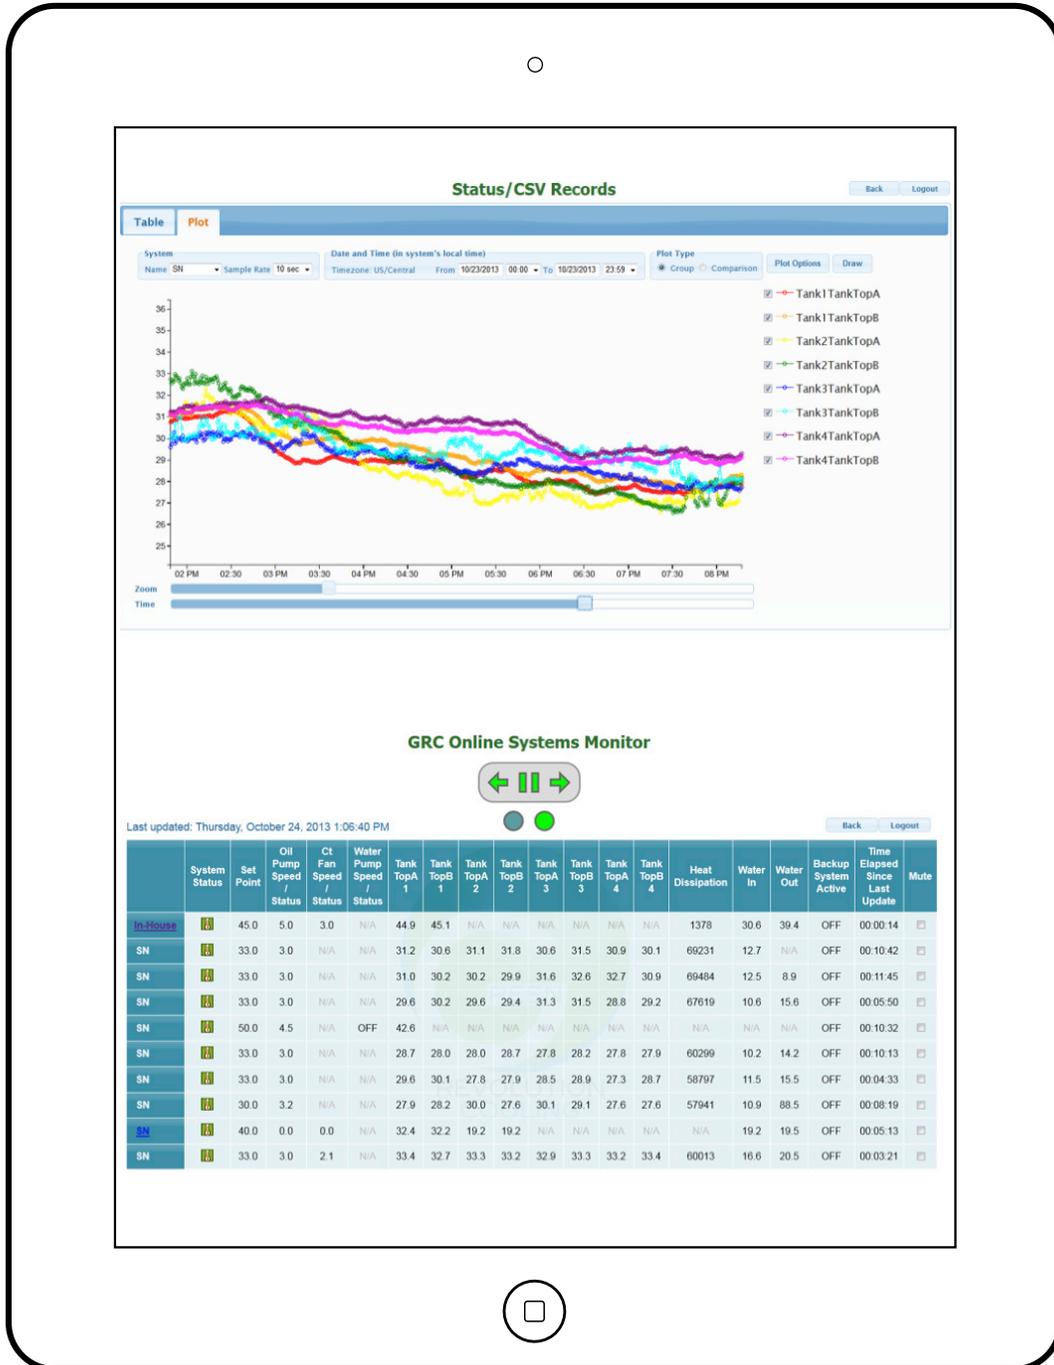
Investing in large CRAC and CRAH systems can only be justified with economies of scale. In fact, even free air cooling only makes sense with large scale installations. In contrast, the CarnotJet system gives data center operators the flexibility to build out computing power and supporting infrastructure almost linearly. With up to four racks sharing a pump module, ramping up the computing power can be done in much smaller increments.

Further, the CarnotJet system (being an integrated data center solution) brings the advantage of prefabrication to data centers, thereby reducing build time down to a few weeks as compared to the industry average of over 9 months. All of that, at a price point that is significantly lower than that of conventional technologies, gives data center operators economies without the constraint of scale.

## Detection Dashboard

The CarnotJet system comes integrated with a number of sensors as well as a control system that provides critical system data and diagnostics tools. The control system continuously measures and monitors system performance, sending out email alerts as and when required. The ongoing monitoring is backed by GRC's 24/7 on call engineers, who can be provided access to system data to help with remote diagnosis. The Foresight system also runs a weekly system health diagnosis, which can be monitored on a web based dashboard or integrated with existing DCIM systems. Historical data can be downloaded in the form of tables, graphical representations, or even CSV log files.

# Dashboard



The Foresight system supports multiple DCIM system integration protocols such as BACnet, SNMP, RESTful API.

## Proven Results

From the most powerful HPC clusters to the cheapest, most basic data center setups (<\$2/Watt), and everything in between, the CarnotJet system has proven its mettle in a range of applications.

The Tsubame-KFC supercomputer, running on GRCs CarnotJet system, was ranked #1 on the Green5004 list of most efficient supercomputers. It was, in fact, the first supercomputer to have breached the 4 GigaFLOPS/ watt (Four Billion Floating Point Operations per Second per Watt), beating its closest competitor by over 30%. The system is also used by 5 of the top 100 supercomputing sites in the world. Global leaders in industries ranging from telecom to oil & gas exploration have adopted the CarnotJet system, due to its capital savings and frugal operating costs. All major server OEMs including Dell, HP, Intel, Supermicro, Cisco, HGST, IBM and many others have had their servers tested and approved for use with ElectroSafe in the CarnotJet enclosures. Intel recently completed a one-year test with servers submerged in ElectroSafe, and reported that they were able to run servers in the CarnotJet system with a PUE between 1.02 and 1.035.

## Conclusion

Liquid immersion cooling is no longer a thing of the future. Green Revolution Cooling is supplying industrial-grade liquid cooling solutions to forward-thinking data centers around the world, taking megawatts of power off the grid, increasing server performance and reliability, while reducing the near and long term costs of owning and operating a data center.



## References

1. Cisco's The Network, (2014). Cisco Global Cloud Index Projects Cloud Traffic to Dominate Data Centers - The Network: Cisco's Technology News Site. [online] Available at: <http://newsroom.cisco.com/release/1274405> [Accessed 11 Jul. 2014].
2. Pfanner, E. (2014). Liquid-Cooled Supercomputers, to Trim the Power Bill. International New York times. [online] Available at: <http://www.nytimes.com/2014/02/12/business/international/improving-energy-efficiency-in-supercomputers.html> [Accessed 11 Jul.2014].
3. ElectroSafe Dielectricliquid coolant.(n.d.).1sted.[ebook]Availableat: <http://www.grcooling.com/wp-content/uploads/2015/03/ElectroSafe-Coolant-Fact-Sheet.pdf> [Accessed 11 Jul.2014].
4. Green500.org,(2014).TheGreen500List-June2014|-TheGreen500.[online] Available at: <http://www.green500.org/lists/green201406> [Accessed 11 Jul. 2014].
5. Miller, R. (2012). Intel Embraces Submerging Servers in Oil | Data Center Knowledge. [online] Data Center Knowledge. Available at:<http://www.datacenterknowledge.com/archives/2012/09/04/intel-explores-mineral-oil-cooling/> [Accessed 11 Jul. 2014].

## Glossary

CAGR: compounded annual growth rate CAPEX: capital expenditure

CPU: central processing unit

CRAC: computer room air conditioner

CRAH: computer room air handler

DCIM: data center infrastructure management GRC: Green Revolution Cooling

ElectroSafe: Trade name for dielectric fluid coolant HPC: high performance computing

OEM: original equipment manufacturer

OPEX: operating expense

PUE: power usage effectiveness (total data center power / IT equipment power)

Written by Dhruv Varma

Layout by Dmitri Mirakyan