AL MARJEIA MODERN DEVELOPED TRADING GROUP.

TABLE OF CONTENTS

About Al-Marjeia

Our History

1-We are operating in the Saudi Market since 2013.
2-We are spreading with the Go green concept in the EMEA, Pakistani & Indian markets.
3-We are providing the latest solutions and innovations that reduce operating expenses and lessen impact on our environment.
We aim to build a trust worthy name in the market, collaborate to save energy and enable businesses to function efficiently.

Our view

1- **Al-Marjeia Modern Developed Trading Group** formed a new division to import the latest innovations in (IT Infrastructure, Energy solutions, non-chemical (physical) water Treatment solutions)
2- We help all sectors (Industrial – commercial - Governmental) to Decrease both capital and operational expenses (CAPEX & OPEX) as well as protecting the environment for the future generations.

- Less opex
- More profit
- Clean environment
- Better future

Our vision

To be a strategic partner to our clients through presenting innovative, cost effective green solutions and excellent services across the EMEA, Pakistan & the Indian subcontinent regions.

Our mission

Promote the growth of businesses by providing cost effective, green and efficient energy solutions. We engage in the sustainability of our environment by introducing excellent services through our dedicated teams

Our future

Al-Marjeia's R&D Department is always searching and seeking reliable, innovative solutions that would be of added value to our clients.
We excel to provide the industrial, commercial and governmental sectors with wide varieties of products, services and customized solutions that fit their needs and exceed their expectations

Al Marjeia Modern Developed Trading Group(LTD)

We offer IT services, products and solutions by alliances with global technology vendors, customers and specialist solutions providers, as well our environmental system consists of many companies in various technology areas. Besides our partnership for specific customer requirements.

Network services we offer includes:

Active Services

IP based network with multi vendors (HP, Huawei, Cisco and Juniper) including supply, configuration and implementation.

Passive Services

Fiber-optic system: designing, supply, implementation, pulling, welding, termination and testing.

Copper systems: designing, supply, implementation, termination and testing (for all systems CAT6, CAT6A & CAT7).

FIBER OPTICS NETWORKING

In recent years it has become apparent that fiber-optics are steadily replacing copper wire as an appropriate means of communication signal transmission. They span the long distances between local phone systems as well as providing the backbone for many network systems, cable television services, university campuses, office buildings, industrial plants, electric utility companies and residential compounds. Fiber optic networks' high speed of operation, bandwidth, resistance to electromagnetic noise, long signal transmission and low cost of maintenance resulted in Fiber optics being the fastest growing global broadband technology with significant deployments occurring across the globe. Al Marjeia LTD offers premium quality fiber network and FTTX services for implementing and maintaining mega projects as well as providing turnkey solutions and material supply

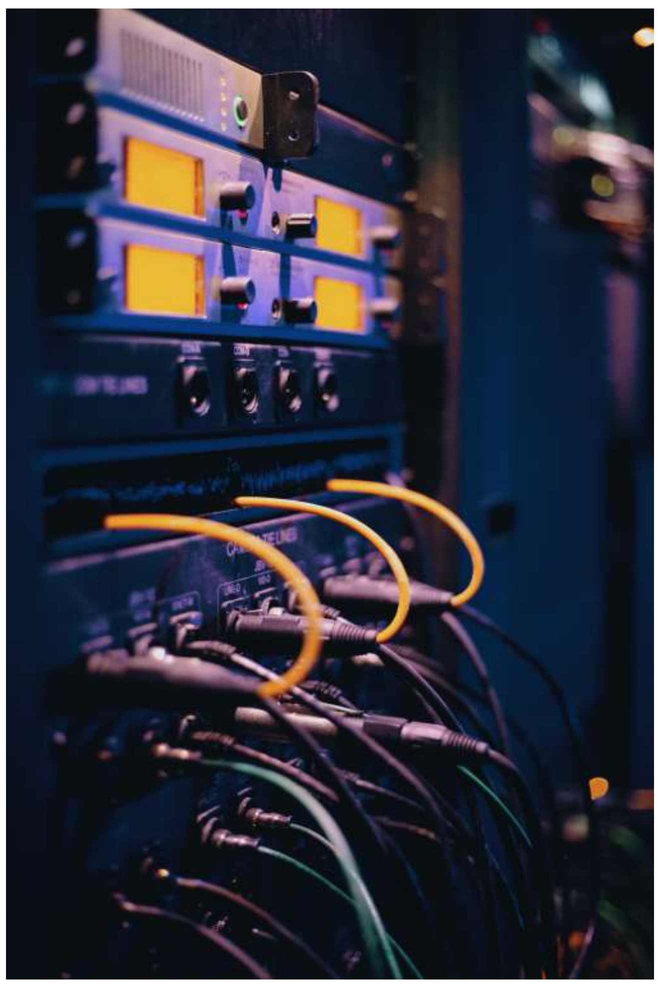

SURVEY & DESIGN

- Verify the best routs for the Fiber optic link.
- Verify the existence of infrastructure.
- End to end design for the Fiber Optic Network.

FIBER INSTALLATION

- Automated and conventional open trenching Cable laying techniques
- Excavation and back filling of all type trenches
- Bridges crossings through attachment or existing bridge ducts, Road Crossing & Railway Crossings
- HDD horizontal directional drilling for rivers and water channels crossing
- Construction and installation of hand holes, manholes and all concrete works
- Laying HDPE/PVC pipes in trenches, jointing of pipes with water/air tight sockets according to the specs
- Cable pulling/blowing in pipes
- Fusion splicing and termination of Fiber Optic cables.
- Testing and commissioning
- Documentation and as built drawings

FIBER MAINTENANCE

- Regular preventive surveys & environmental faults fixing (corrective maintenance) as per SLAs
- Fault management teams, ready on the move to handle faults around the clock (24/7)
- Fault detection teams, allocated along the routes to detect and fix faults within minimal response time

WIRELESS COMMUNICATION

Making faster and better business decisions requires access to vital data in all types of situations. AL Marjeia's suite of Wireless Connectivity Solutions including WiMax, Pre-WiMax, WiFi and Microwave, delivers an unmatched combination of range, capacity, security, and reliability, even in very challenging environments. Best-in-class data rates and spectral efficiencies conserve valuable spectrum and reduce the costs of providing high-performance access and deliver characteristically robust, extreme-distance/high-bandwidth performance. With built-in security encryption, at speeds of up to 90 Mb/sec, our solutions can be used to transmit signals as far as 100 Km. Ideal for a variety of access, backhaul and private network applications, our solutions are meeting the needs of carriers, service providers and enterprises worldwide. Alkan's Wireless Solutions maintain a reliable link in challenging environments, thus minimizing installation costs reducing location rentals, and allowing users to overcome the limitations of fixed wired access technologies. Applications include:

- Wireless Surveillance for city traffic monitoring
- Voice TDM (E1's/T1's) transport for telecom operators
- Internet connectivity and low latency VOIP applications for ISP's and fixed lines operators
- Triple Play and VOD applications for cable TV operators
- Oil rigs connectivity solutions Military and mission critical communications
- Backhauling and disaster recovery solutions for GSM traffic

INTEGRATION & IT INFRASTUCTURE

As much as the business dependence on IT increases the risk for services interruptions both for internal and external users increases , therefore reliable and scalable infrastructure is becoming a fundamental cornerstone for today's business continuity. **AL Marjeia LTD** provides a multiple array of solutions that can efficiently manage new technological developments, application changes, responds to security threats and growing IT capacities. Offerings include Networking Integration, Bandwidth Optimization, IT Asset Management, Hardware Virtualization, Business Service Management and Enterprise Network Management

NETWORKING INTEGRATION

Communication Convergence is the keyword, understanding the importance of Scalable continuity to the success of business for the financial sectors' institutions,**AL Marjeia LTD** provides turnkey converged networking solutions (multi-service networking) that enable integration of data, voice and video solutions onto a single (IP based) network and empower customers to gain a competitive advantage today and tomorrow. Benefits include improving network performance, enabling implementation of complex network initiatives and reducing risks associated with network changes

BANDWIDTH OPTIMIZATION

WAN optimization appliances can make more efficient use of existing connections and may even yield better performance than merely throwing bandwidth at the problem. **AL Marjeia LTD** offers a wide range of Application acceleration and WAN optimization products that are needed to improve the performance of applications and services operating over their WAN. Among the benefits of bandwidth optimization are IT cost reductions, Business continuity, disaster recovery and improved application performance

NETWORK SECURITY

Top notch security solutions that enable organizations of all sizes to effectively minimize threats of intrusion, secure networks and protect data, by deploying a comprehensive and proactive multilayered system that smoothly integrates with existing infrastructure, thus allowing better utilization of resources, enhancing performance and eliminating the danger of security breaches and hacking. Services include risk evaluation, design and implementation of policies, procedures and customized solutions both for Network and Data Security. Services include ethical hacking, Security Risk Assessment, Vulnerability assessment, Penetration testing, PCI Compliance preparation and ISO 27001 certification preparation

DATA CENTERS

IT operations are a crucial aspect of most business operations, where organizations totally rely on their information systems to run their operations. If a system becomes unavailable or an information security breach is detected, transactions may have to be stopped completely. A data center therefore has to offer a secure environment and keep high standards for assuring the integrity and functionality of its hosted computer environment that minimizes the chances of a security breach and provides a reliable infrastructure for IT operations that guarantees business continuity. **AL Marjeia LTD** offers a complete data center solution including Storage, Server consolidation, backup, disaster recovery, Racks, UPS and Physical Security needs

IT BUSINESS PROCESS OUTSOURCING (IT BPO)

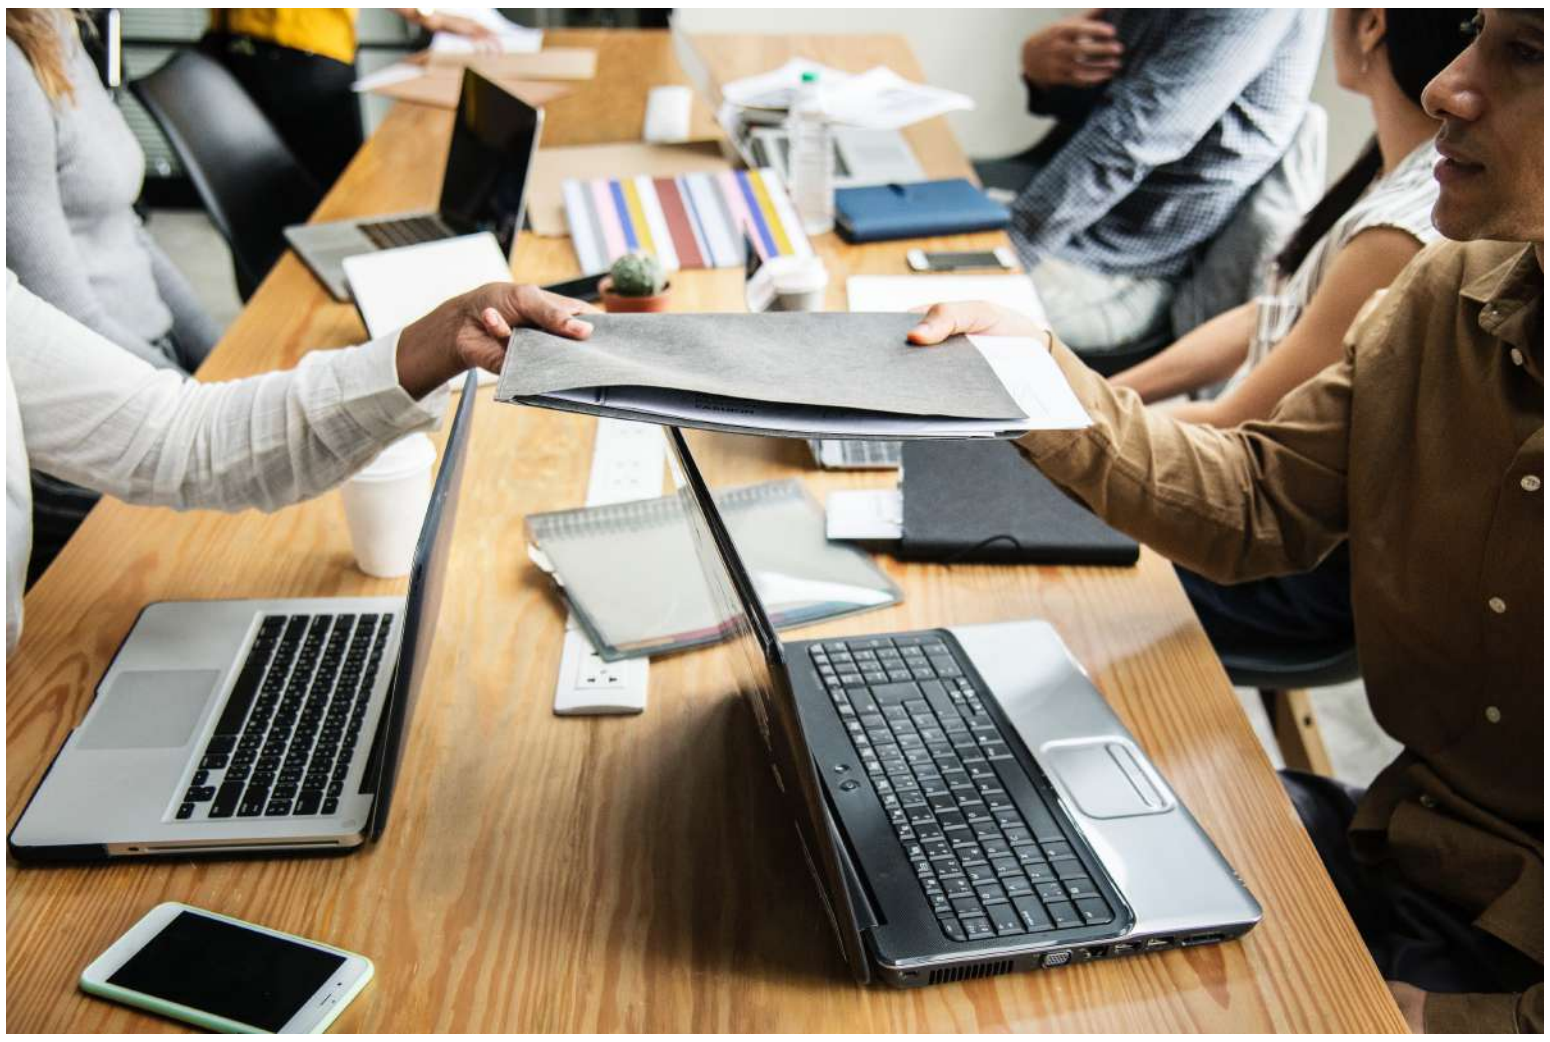

IT BPO is the sought after strategy for companies seeking new ways to achieve high performance by controlling costs, reducing risk, fostering collaboration and increasing transparency. Today's businesses tend to focus on core businesses and reduce costs by outsourcing, especially on bringing in different skills following a major restructuring, acquisition or divestment, also for obtaining maximum value from a non-core business, but reliance on others can increase risk, that is why businesses need safe, secure and cost-effective ways of securing the support they need for day-to-day operations, development and growth. **AL Marjeia LTD** is one of the few companies providing all the elements of professional outsourcing services for communication and Information Technology infrastructure across the different fields. Our professional services cover all aspects of networking and IT infrastructure including desktop, server, network and application, storage and security, taking into account business process updates and regular technology updates. **AL Marjeia LTD**'s services are vital to success of any organization; it also allows clients to focus on their core business, safe in the knowledge that their ICT infrastructure remains secure and continue to play their role in developing their services and meeting critical business

VIDEO TECHNOLOGIES

A picture is worth a thousand words, but when it comes to video technologies "they tell the whole story". Video Surveillance systems are becoming a necessity for leveraging the effectiveness of organizational security systems, whether internally or externally and literary enable organizations monitor their activities anywhere around the globe. As for video conferencing, organizations are empowered to communicate, collaborate and cut down travel expenses on timely manner

SURVEILLANCE

The span of video surveillance services has recently widened to comprise both remote and urban operations where safety is directly tied to cost efficiency. **AL Marjeia LTD** utilizes state of the art technologies and latest innovations in digital technology and IP surveillance to offer scalable turnkey solutions that effectively enable preventing robberies, detecting suspicious behavior and enabling Perimeter security. Offerings are complemented with a portfolio of high-value services, including consultation services for assessing current systems, solution design and configuration including camera locations, networking platform, system installation and integration in addition to ongoing maintenance and future expansion.

VIDEO CONFERENCING

In today's business, organizations increasingly need to maintain communications and confidentiality across their geographically dispersed offices. **AL Marjeia LTD** advanced professional services and video conferencing solutions enable organizations to hold secure, distance meetings and collaborate with remote colleagues, share important information and reach critical decisions quickly. **AL Marjeia LTD** features the latest technologies, including HD video conferencing and audiovisual integration, which make video calls as vibrant, collaborative and productive as face- to-face meetings.Offerings include Telepresence,Conference RoomSolutionsandDesktopVideoConferencing

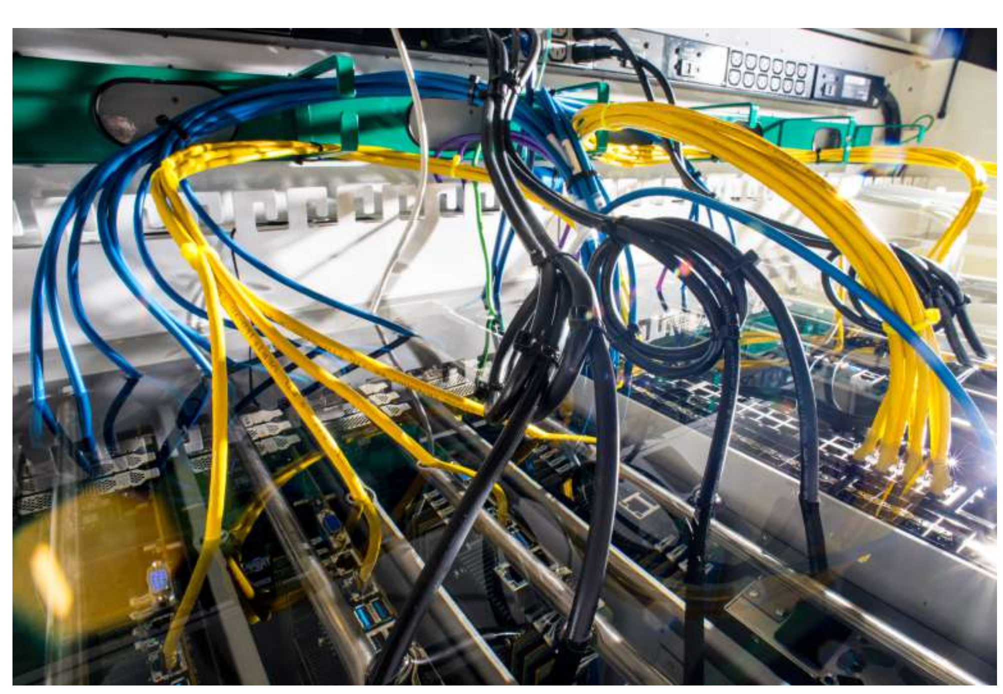

The Experience

STC DC Project

Active and Passive project in STC Tabouk and Shaqra

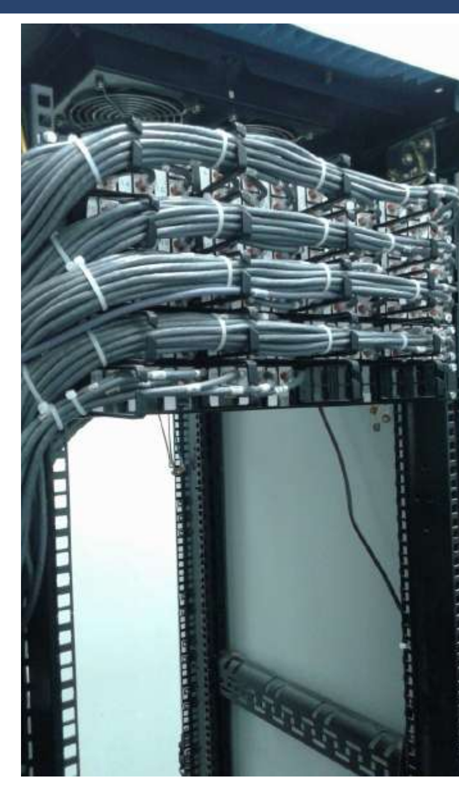

Fiber Project Jeddah-AL Madinah

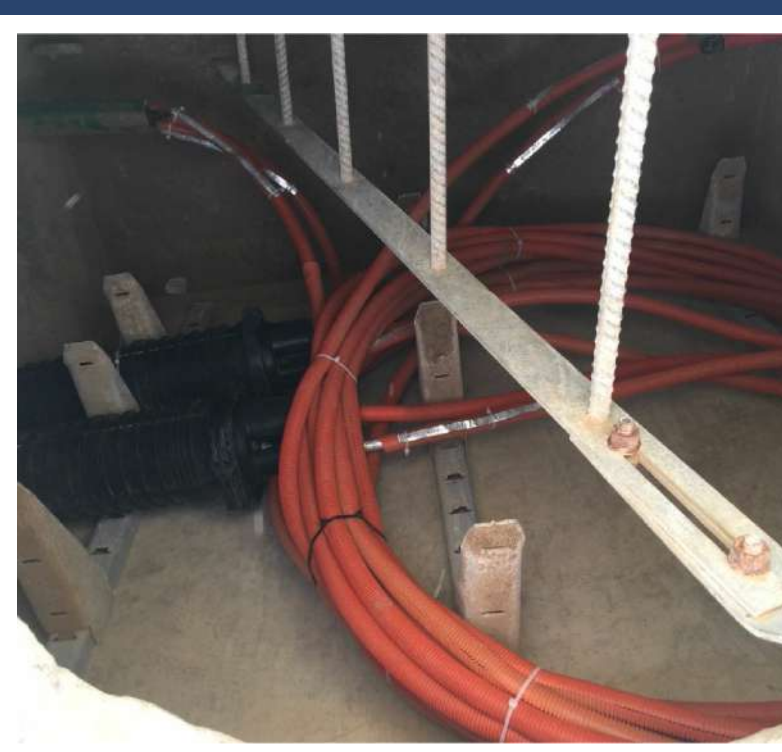

CONFIDENTIAL

Riyadh

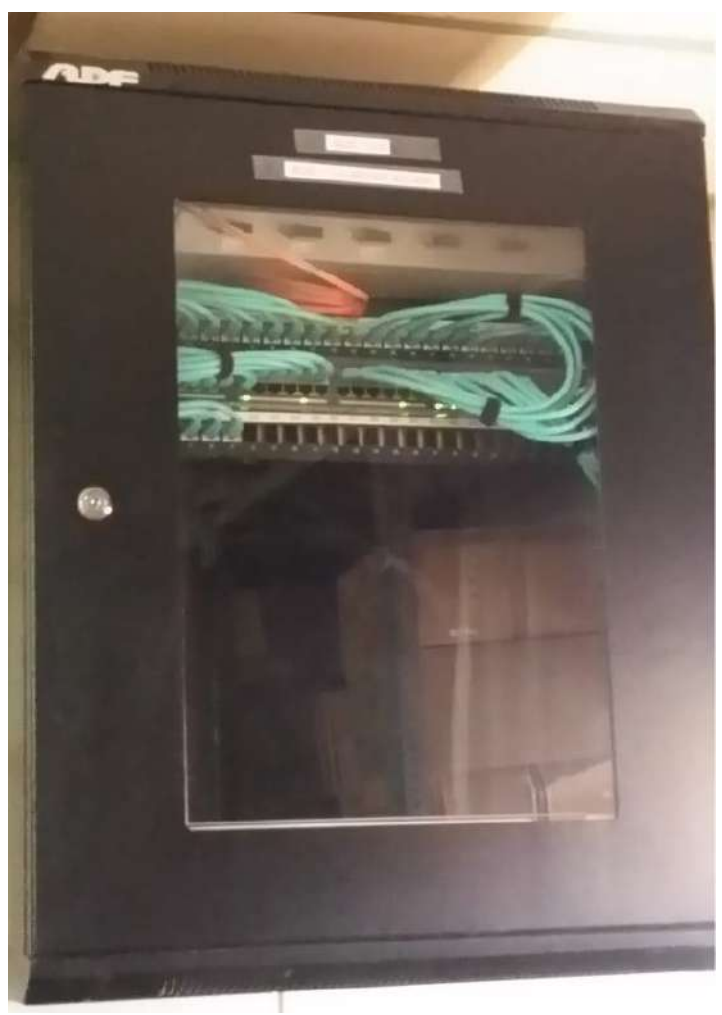

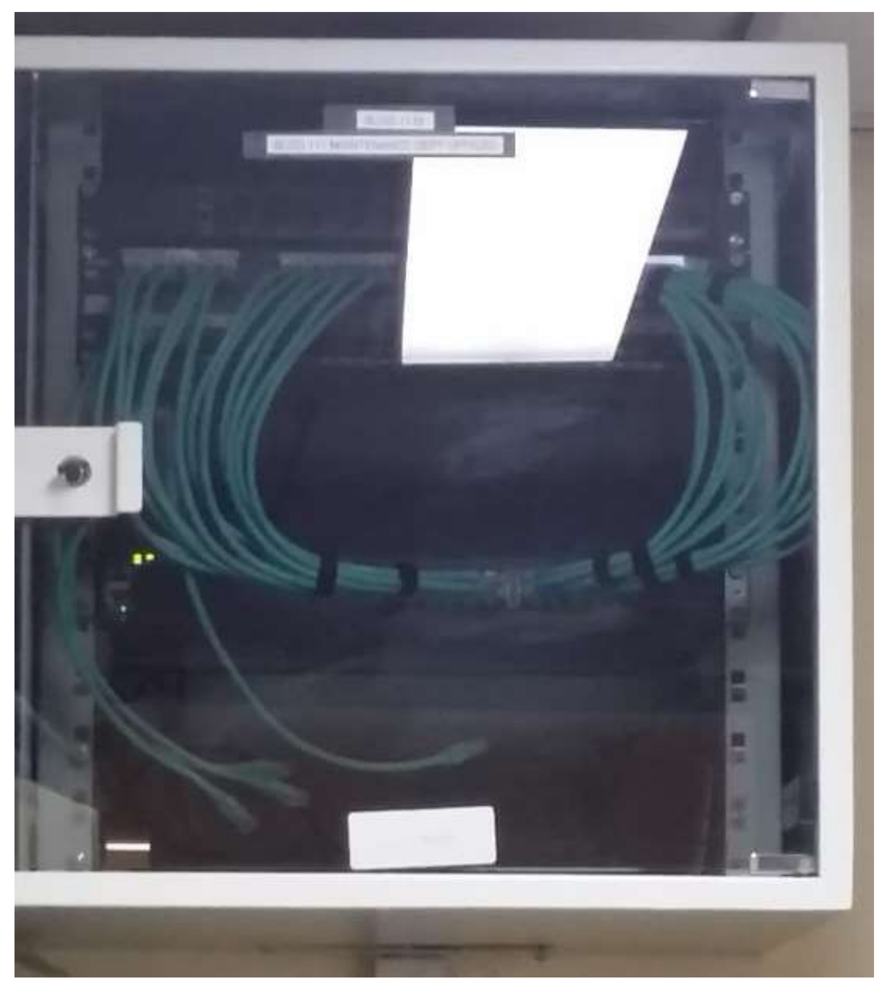

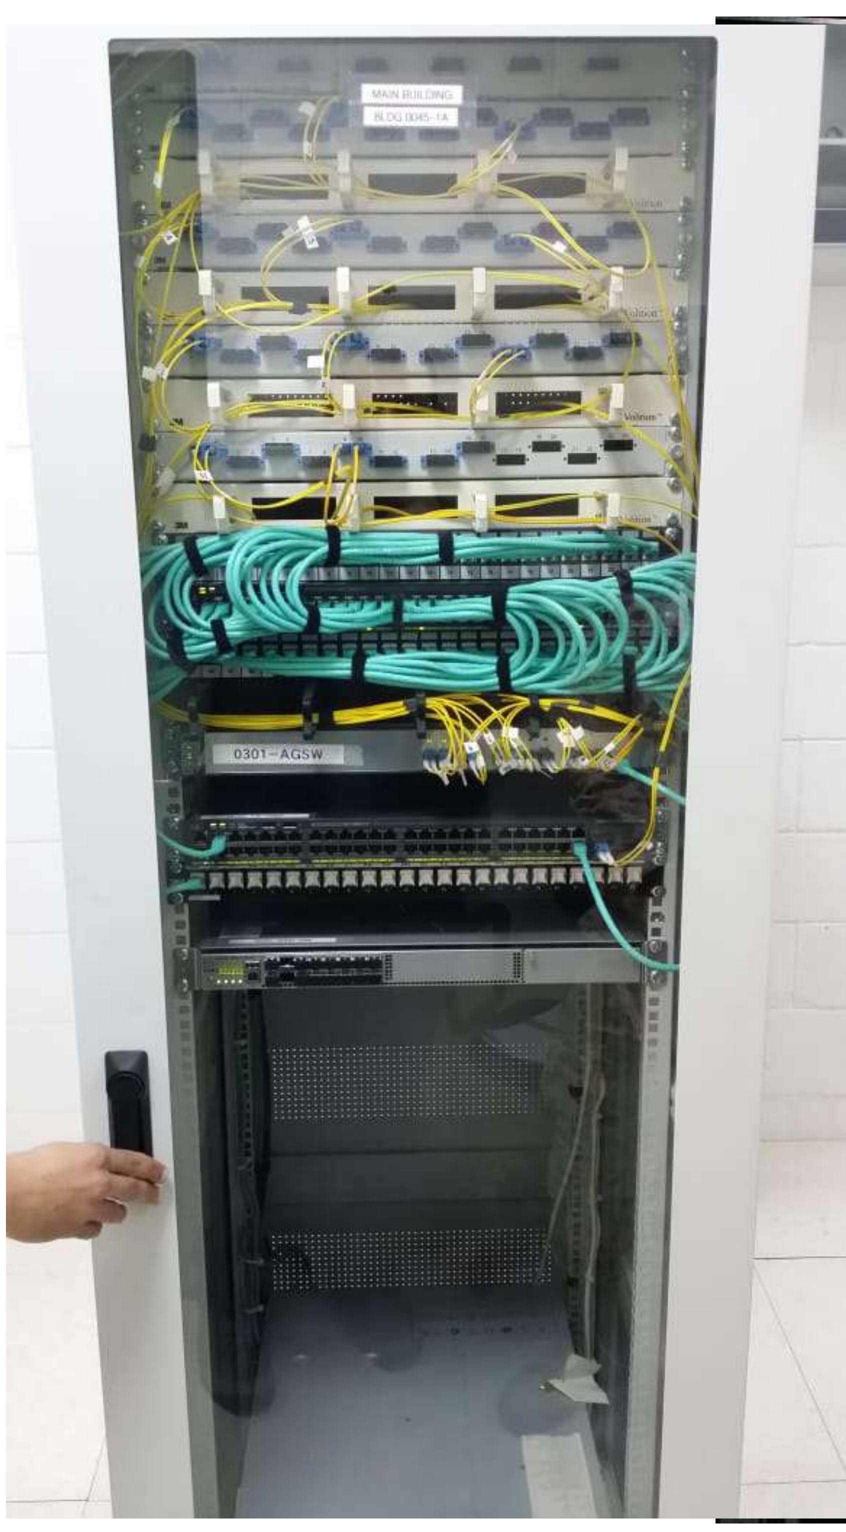

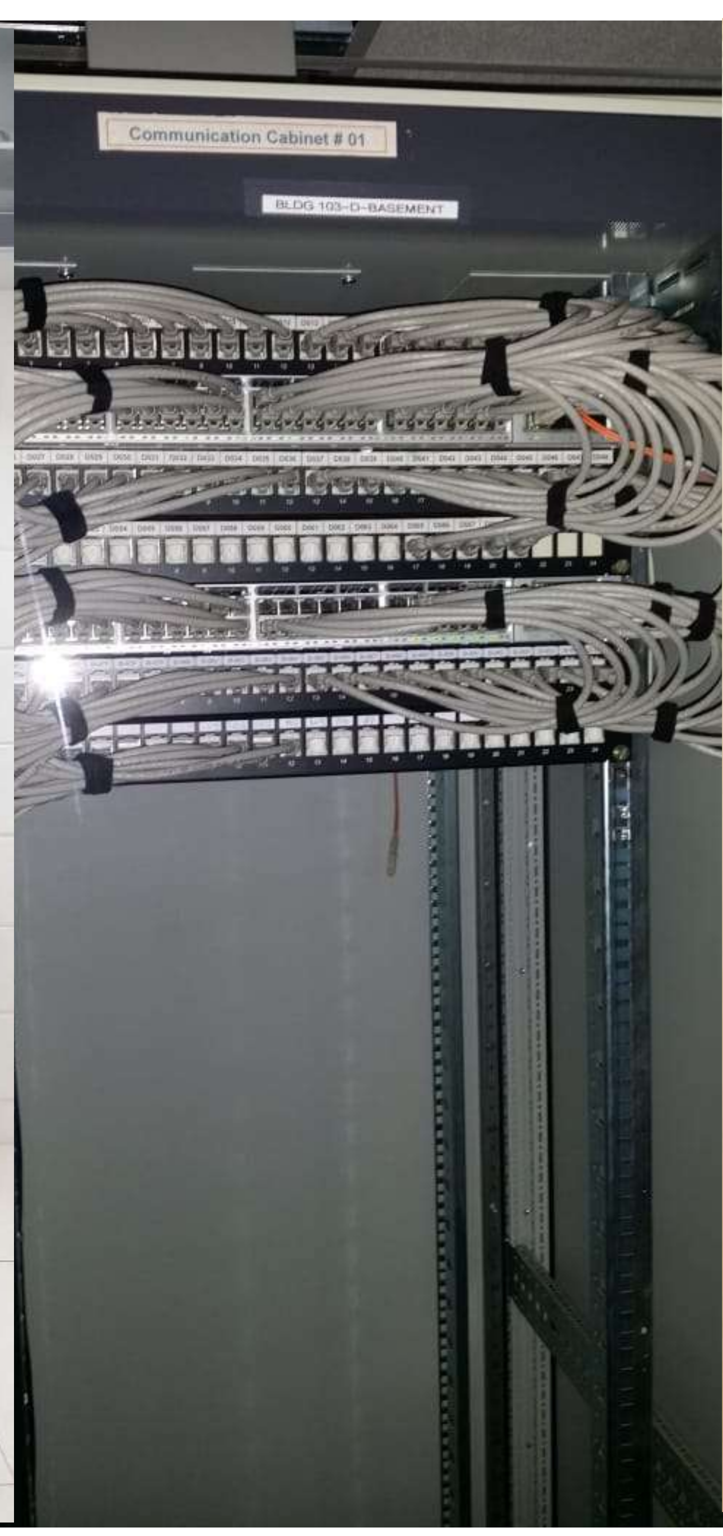

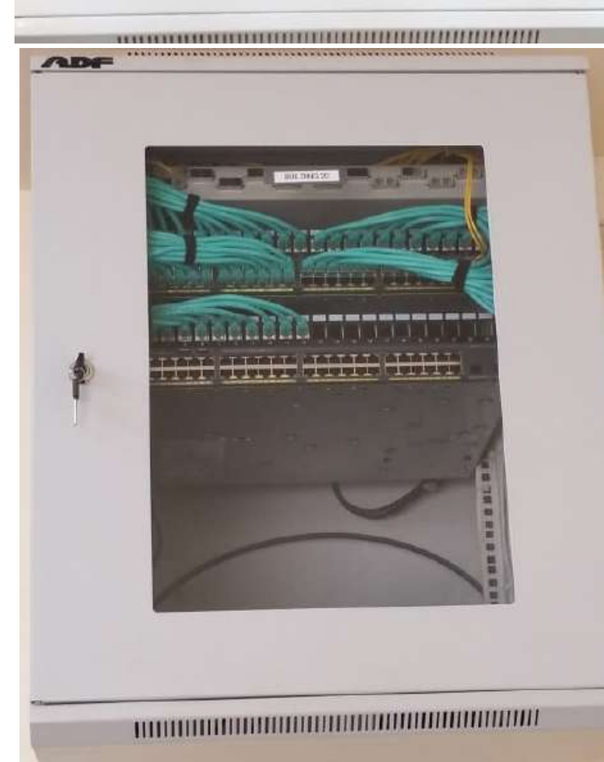

GRC
GREEN REVOLUTION COOLING

The Immersion Cooling Authority

The Immersion Cooling Authority

Future-Proof Your Data Center with Next-Generation Immersion Cooling Technology

Traditional Data Center Cooling Can't Take the Heat

Quick Facts:

- Founded in 2009
- Pioneer of single-phase immersion cooling
- Holder of 10 patents—with 8 pending
- Strong vetting by IT icons such as the NSA and Intel®
- World wide installations in 13countries

GRC Delivers Game-Changing Cooling

Business success today is increasingly tied to IT. Yet traditional data centers are very costly to build. What's more, amid a soaring demand for greater computing power, server heat dissipation has become a huge barrier.

We're the World's Immersion Cooling Authority

Green Revolution Cooling (GRC) was founded with a vision to change the way data centers are designed, built and operated.
With our CarnotJet system we pioneered single-phase immersion cooling technology to help companies grow computing power easily, cost-effectively and limitlessly through a dramatic leap in data center cooling capabilities.
Perfected over the course of a decade, our innovative liquid-cooled server racks have been vetted by IT giants like Intel® and the NSA, and are at work within some of the world's largest cloud,enterprise, education, government and telecom organizations.Buoyed by CarnotJet's widespread acclaim, we have since developed two new solutions — ICE and Hash.

Break the Heat Barrier — Take Your Data Center Beyond Limits

From the beginning, data center demands have increased year-over-year in lockstep with business growth. To keep pace, each new computing generation has brought cooling closer to the heat source. Yet today's high-performance servers are consuming more power than ever.
Traditional methods can't cool new and emerging systems.

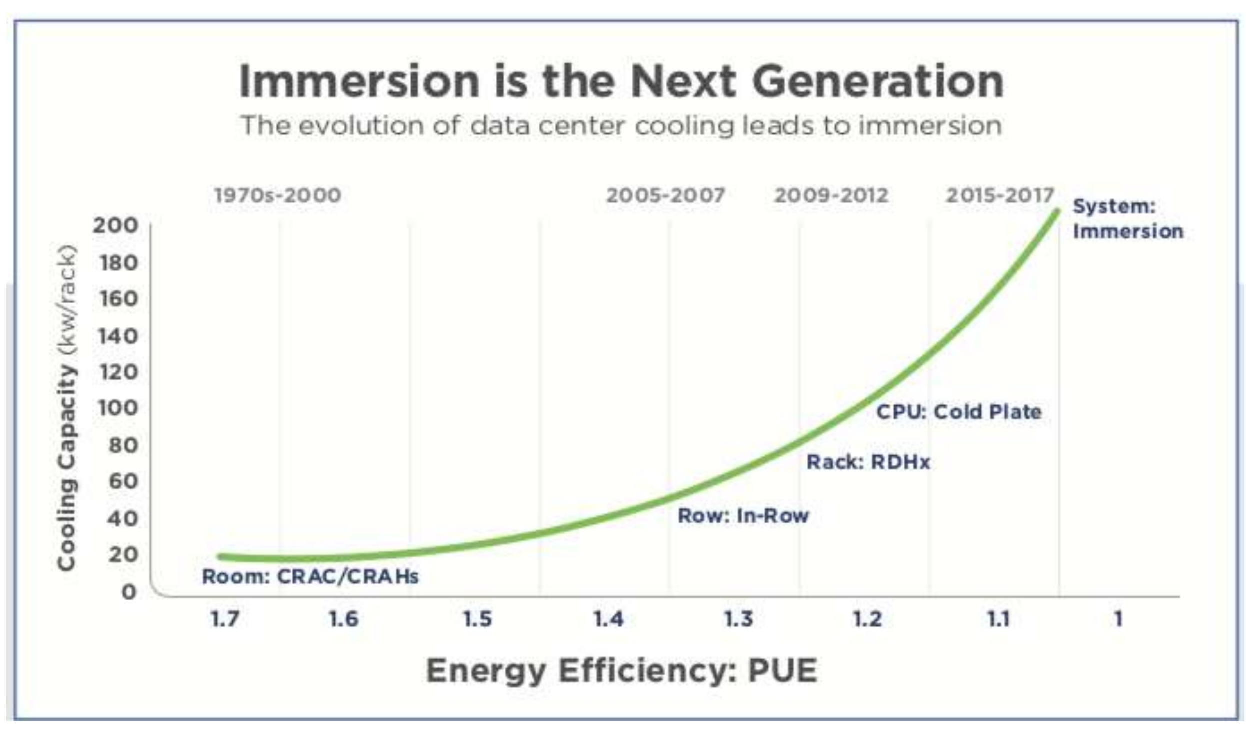

The Next Evolution Is Here

Radically simplifying the design, build and operation of data center cooling infrastructure, GRC's patented immersion cooling solutions provide the breakthrough growth and cost-saving potential you're looking for.
Immersing servers in liquid affords a major across-the-board improvement in rack density, cooling capacity, data center layout and location options, enabling you to:

- Scale easily, cost-effectively — and limitlessly
- Slash capex up to **50%** — and cooling energy costs up to **95%**
- Increase cooling capacity up to **100kW/rack**

"We saturated the power envelope by putting twice as many systems as we would normally have, if it had a normal way of cooling." — CGG

How GRC Liquid Immersion Cooling Works:

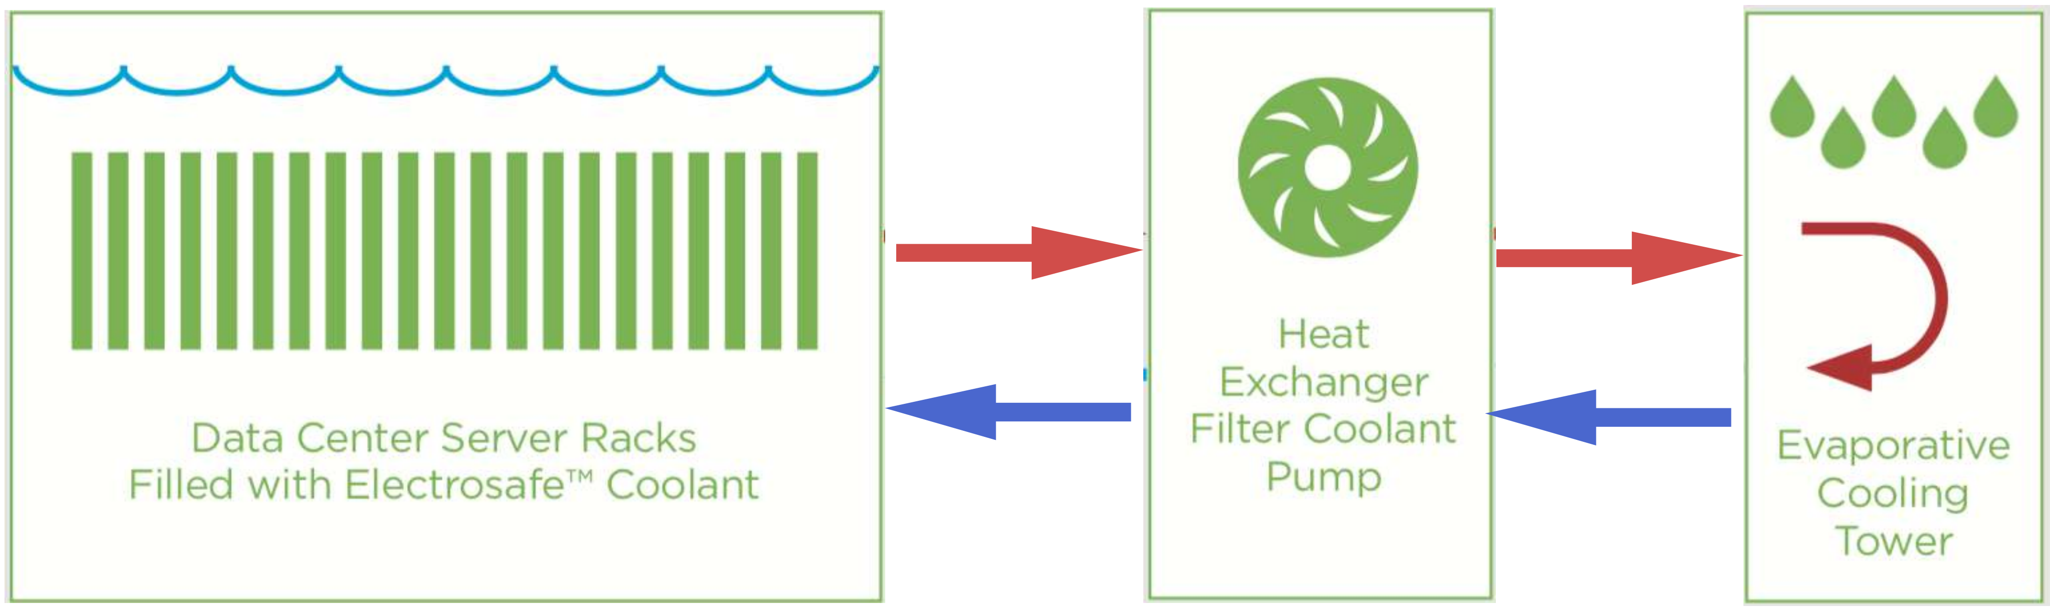

ElectroSafe coolant is a clean, odorless, nontoxic liquid with 1,200X the heat retention capacity of air by volume. Immersing servers in this liquid enables cooling up to 100kW/rack!

GRC Has Your Data Center Performance Down Cold

Choose the solution that's right for your operation — and your vision.

ICE Line of Products

*ICE*raQ™ Rack-Based System

Cut the cost of building, running and expanding your data center — to achieve unprecedented cost, performance and space efficiencies.

Ideal for data centers supporting:

- Artificial Intelligence (AI)
- High-Performance Computing (HPC)
- Cloud/On-Premise Data Centers
- High-Frequency Trading (HFT)
- Edge Computing

Includes:

- Immersion cooling server racks
- Cooling unit (CDU)
- ElectroSafe coolant
- Accessories

*ICE*tank™ Containerized System

Eliminate concerns about location or environment with our turnkey data center in a box. You provide the power, water and networking. We do the rest.

Available in two Models:

- IT20: 20' container with four ICEraQ's
- IT40: 40' container with eight ICEraQs

IT20

IT40

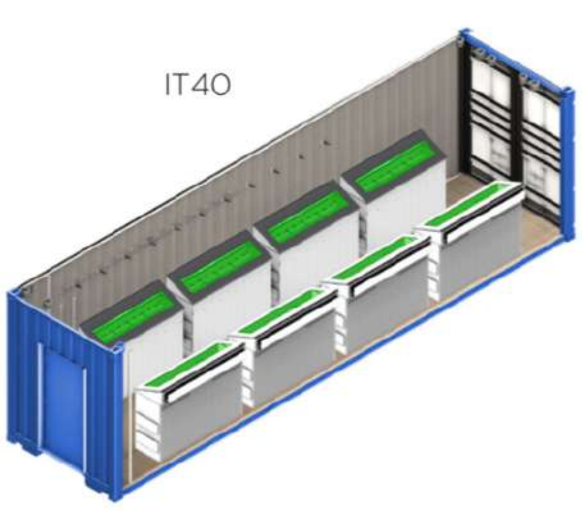

Hash Line of Products

*Hash*raQ™ Rack-Based System

Make crypto mining more profitable, slash data center capex and opex, mine different currencies based on relative profitability, and respond quickly to market fluctuations.

Ideal for crypto mining & blockchain applications

Includes:

- Immersion cooling server racks
- Cooling unit (CDU)
- ElectroSafe coolant
- Accessories

*Hash*tank™ Containerized System

Space or environment a problem? Drop your crypto mining or blockchain data center in a parking lot, the desert — virtually anywhere.

Available in two Models:

- HT20: 20' container with two HashraQ's
- HT40: 40' container with six HashraQs

H20

H40

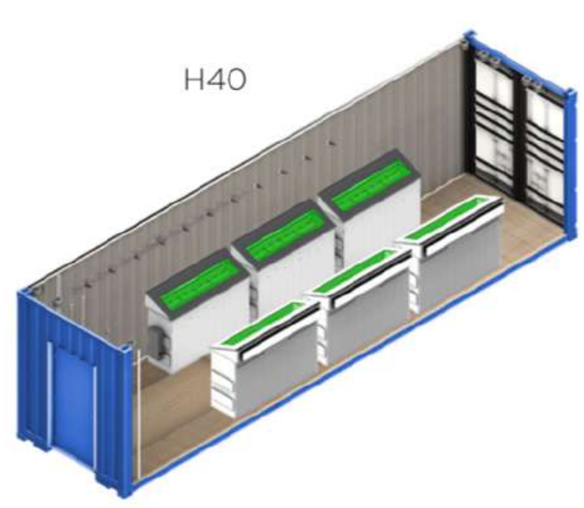

GRC's Solutions Are :

Cost-Effective
Cut data center construction costs by up to 60% by eliminating chillers, CRACs, CRAHs and raised floors.

Scalable
Minimize the pains of capacity planning and inaccurate forecasts. Build as you go and grow your operation.

Resilient
Eliminate fan vibrations, dust and moisture contamination, oxidation risks and hot spots.

Fast + Flexible
Locate your data center in virtually any environment — and be up and running within weeks, not months.

Power-Efficient
Experience a <1.05 PUE, 10-20% server load reduction plus an average power reduction of ~50%. Get more from your power envelope, too.

Future-Proof
Cool the most powerful servers—up to **100kW/rack** — and capitalize on emerging applications.

GRC cooling solutions support servers manufactured

Compute Cooler. Run Smarter. Grow Faster

ICEraQ ™

Quad/Duo/One Models

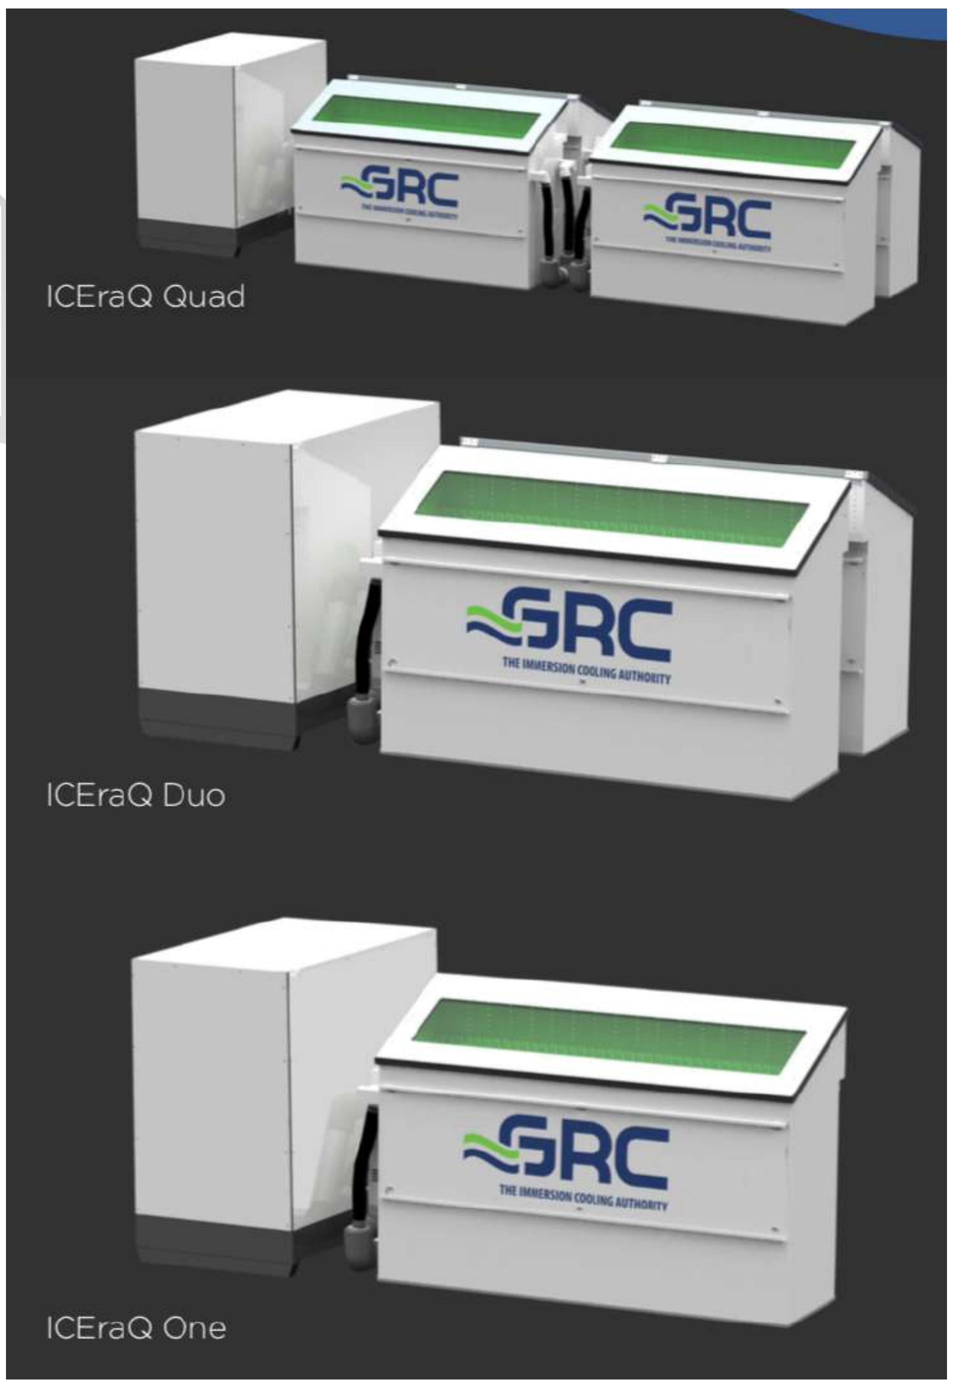

Low-Cost, High-Efficiency Immersion Cooling for Data Centers

Product Specifications	Quadt	Duo	One
Number of Immersion-Cooled Rackst	4 x 52U	2 x 52U	1 x 52U
Number of Coolant Distribution Units(CDU)	1	1	1
Chillerfree, Warm Water Operation (95° F) Cooling Capacity/Per Rack Density	100kW/25kW	100kW/50kW	100kW100kW
Chilled Water Operation (55° F) Cooling Capacity/Per Rack Density	200kW/50kW	200kW/100kW	200 kW/200kW
Mechanical PUE	1.03		1.03
Redundancy[1]	Coolant Pumps: 2N, Control System: 2N		
Rack Dimensions (lxbxh)	100"x27"x52" (2.54mx0.7mx1.3m)		
Rack Floor Loading	154 lbs/sq ft. (752 kg/m2)		
CDU Dimensions (lxbxh)[2]	63"x36"x65" (1.6mx0.9mx1.7m)		
CDU Floor loading	57 lbs/sq ft. (278 kg/m2)		

Delivery & Installation

Shipping Lead Time	• 10-12 weeks after receipt of purchase order
Shipping Terms	• Ex Works
Installation	• Three days for the first unit, plus two days for every subsequent unit. GRC service engineer/s will be onsite to perform installation and training.

Power & Water Specifications

Final Heat Rejection Options[3]	Evaporative Cooling Tower or Facility Water Loop / Chiller Water Loop
Water Requirements	Water input temperature 37° F to 122° F (3° C to 50° C) - Recirculating water flow rate 50-150 gpm (11 - 34 m3/hr) - Connections 2.5" FNPT or hose barb
Power Requirements	Two electrical feeds (primary & secondary) each with the following characerteristics: - 3 Phase 200 - 240 VAC, OR 380 - 480 VAC, 50-60 Hz - Max Power Consumption: 2.3 kW

Monitoring and Reporting

Platform	Foresight DCIM: Proprietary, web-based remote/local monitoring and reporting system
Alerts	Configurable email alerts
DCIM / BMS Integration Protocols	BACNET & SNMP
Data Measurements	- Heat load - Coolant pressure - Pump speed - Operating temperatures - water and coolant - Operating pressures - water and coolant - Rack temperature - multiple locations - Power consumption - Liquid level - multiple locations - System health, diagnostics, and early fault detection

Infrastructure

Site Requirements	- Access to power and water - Level installation surface (raised floor or concrete slab)
Operating Guide- lines	- Ambient temperature: 0° F to 122° F (-18° C to 50° C) - Secondary containment at least 110% of the single largest container - Fire Suppression: Standard data center fire suppression

Warranty[4]

• One-year limited warranty including:
- 24/7 on-call GRC support staff
- 24/7 remote monitoring
- Domestic (US): Three business-day onsite response with parts and labor - International: Parts by mail
- Regular maintenance as required
• Additional SLAs, and extended warranty & maintenance contracts are available for an additional fee

[1] Additional redundancy options available.
[2] Underfloor and low profile CDUs available for space constrained sites.
[3] Low water use options available in climates where the design dry bulb temperature does not exceed 90°F / 35°C. Additional costs apply.
[4] Warranty is Void if the ICEraQ units are run outside of their operating parameters defined in the installation specification.

HashRaQ™

Model HR without Miners

Immersion-Cooled Racks for Blockchain Computing Simple. Efficient. Limitless.

Product Specifications

Number of Antminer S9 Units Supported1 (not included)	72
Number of Immersion Cooled Racks	1 x 52U
Power Consumption-Miners	108 kW (Max)
Power Consumption-Cooling	2.3 kW (Max)
Rack Dimensions-Approximate	100”x27”x52” (2.54mx0.7mx1.3m)

HashraQ HR without Miners

Power & Water Specifications

Final Heat Rejection Type (not included)[2]	Evaporative Cooling Tower
Water Usage-Estimated	1-1.5 GPM (standard water hose)
Input Voltage	CDU: 3 Phase 200 - 240 VAC, OR 380 - 480 VAC, 50-60 Hz Custom PDU (optional): 3 Phase 380 - 415 VAC, 50-60 Hz
Nominal Input Amperage	200 A

Delivery & Installation

Shipping Lead Time	• 12-14 weeks after receipt of purchase order, plus two weeks after receipt of Bitmain Antminer S9’s
Shipping Terms	• Ex Works
Installation	• Three days for the first unit, plus two days for each subsequent unit. GRC service engineer(s) will be onsite to perform installation and training.

Data Sheet

Site Preparation

Customer to Provide

• Mining hardware
• Installation preparation, including electrical & plumbing design and construction • Cooling towers, including water treatment and plumbing to CDUs
• Electrical distribution with appropriate breakers and panels
• Applicable permits
• Firewall protection / routers / switches
• Internet connection for miners (estimated at 10kB bandwidth per miner)
• Bitcoin wallet and mining pool account

Warranty & Maintenance[3]

• One-year limited warranty on infrastructure only, including:
- 24/7 on-call GRC support staff
- 24/7 remote monitoring
- Domestic (US): Three business-day onsite response with parts and labor - International: Parts by mail
- Regular maintenance as required
• Extended warranty & maintenance contracts available for an additional fee

[1]Antminer units not included. Other miner hardware supported. Contact GRC for more info.
[2] Low water use options available in climates where the design dry bulb temperature does not exceed 90°F / 35°C. Additional costs apply.
[3]Warranty is Void if the HashRaQ units are run outside of their operating parameters defined in the installation specification. Bitmain products are tested in air prior to immersion to verify no DOA units are accepted. Warranty on Bitmain products does not extend past acceptance testing.
GRC believes the information in this Data Sheet to be accurate; however, GRC does not make any representation or warranty, express or implied, as to the accuracy or completeness of any such information and shall have no liability for the consequences of the use of such information.
This Data Sheet and its contents do not constitute an order by GRC to sell any product. An order is made only by a quotation provided by GRC. The terms of sale in such quotation may vary from those set forth in this Data Sheet. GRC’s acceptance of any order shall be in GRC’s sole discretion, and all quotations and sales are subject to GRC’s Terms and Conditions of Commercial Sale.

ICEtank™
Models IT20/IT40

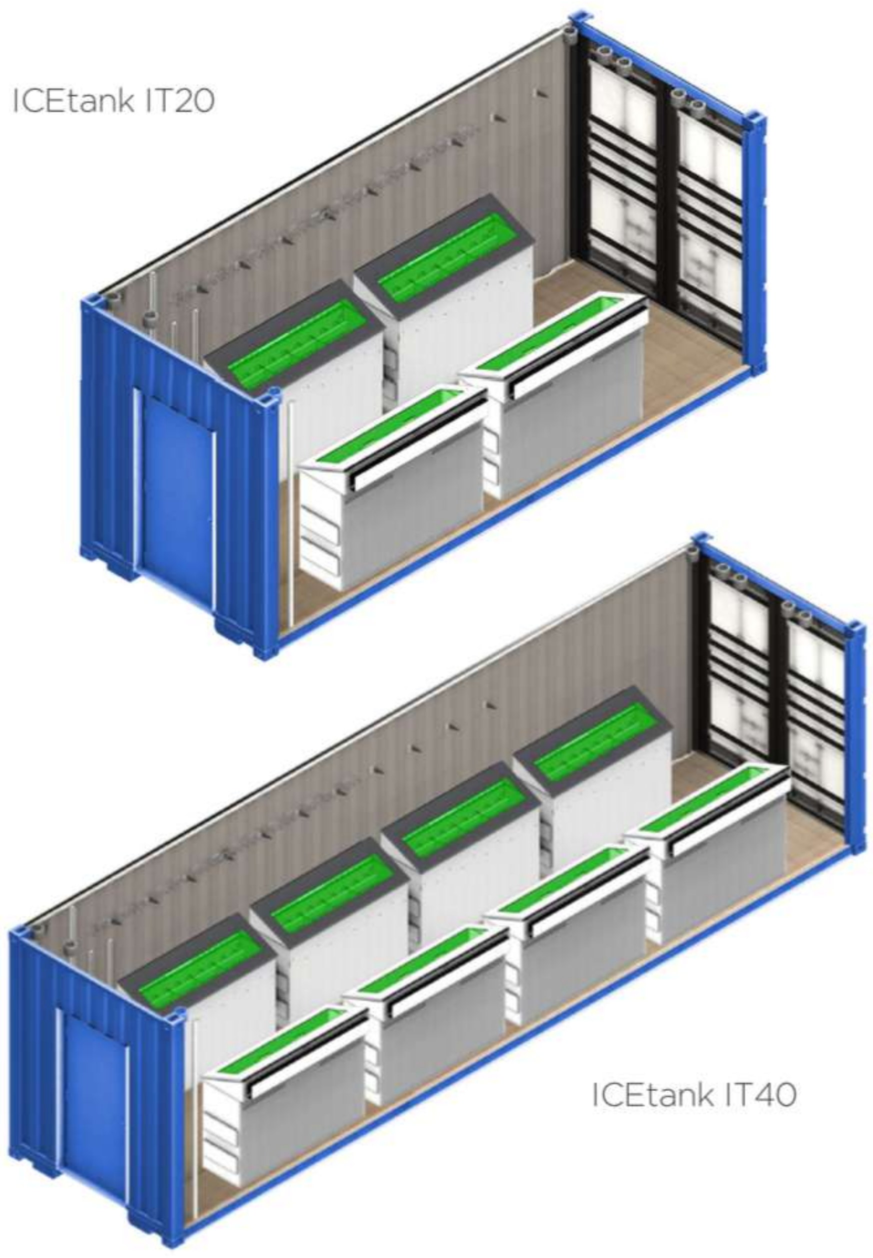

Containerized, Immersion-Cooled Data Center
Quick Deployment in Any Environment

Product Specifications	IT20	IT40
Number of Immersion-Cooled Racks	4 x 42U	8 x 42U
Number of Coolant Distribution Units (CDU)	1	2
Chillerfree, Warm Water Operation (95° F) Cooling Capacity/Per Rack Density	100 kW/25kW[1]	100 kW/25kW[1]
PUE	1.03-1.09	1.03-1.09
Redundancy[2]	Coolant Pumps: 2N, Control System: 2N	
Dimensions (lxbxh)	20' ISO shipping container: 20' x 8' x 9'6"	40' ISO shipping container: 40' x 8' x 9'6"
Approximate weight (Not including IT)	25,000lbs	42,000lbs

Delivery & Installation

Shipping Lead Time	• 12-14 weeks after receipt of purchase order
Shipping Terms	• Ex Works
Installation	• seven to ten days for the first unit, plus five to eight days for every subsequent unit. GRC service engineer/s will be onsite to perform installation and training.

Monitoring and Reporting

Platform	Foresight DCIM: Proprietary, web-based remote/local monitoring and reporting system
Alerts	Configurable email alerts
DCIM / BMS Integration Protocols	BACNET & SNMP
Data Measurements	- Heat load - Coolant pressure - Pump speed - Operating temperatures - water and coolant - Operating pressures - water and coolant - Rack temperature - multiple locations - Power consumption - Liquid level - multiple locations - System health, diagnostics, and early fault detection

Power & Water Specifications

Final Heat Rejection[3]	Cooling towers / Hybrid Coolers
Water Usage-Estimated	1-2 GPM (standard garden hose)
Power Requirements	3 Phase 380 - 480 VAC, 50-60 Hz

Warranty[4]

• One-year limited warranty including:
- 24/7 on-call GRC support staff
- 24/7 remote monitoring
- Domestic (US): Three business-day onsite response with parts and labor
- International: Parts by mail
- Regular maintenance as required

• Additional SLAs, and extended warranty & maintenance contracts are available for an additional fee

[1] High density/custom configurations available. 2 Additional redundancy options available.
[3]Zero water use options available in climates where the design dry bulb temperature does not exceed 90°F / 35°C. Additional costs apply.
[4]Warranty is Void if the IceTank units are run outside of their operating parameters defined in the installation specification.
GRC believes the information in this Data Sheet to be accurate; however, GRC does not make any representation or warranty, express or implied, as to the accuracy or completeness of any such information and shall have no liability for the consequences of the use of such information.

This Data Sheet and its contents do not constitute an order by GRC to sell any product. An order is made only by a quotation provided by GRC. The terms of sale in such quotation may vary from those set forth in this Data Sheet. GRC's acceptance of any order shall be in GRC's sole discretion, and all quotations and sales are subject to GRC's Terms and Conditions of Commercial Sale.

HashTank™

Models H20/H40 without Miners

HashTank H20

HashTank H40

Containerized, Immersion-Cooled Data Center for Blockchain Computing

Product Specifications	HashTank H20	HashTank H40
Number of Antminer S9 Supported [1] (not included)	144	432
Number of Immersion Cooled Racks	2 x 52U	6 x 52U
Power Consumption-Miners	216 kW (Max)	648 kW (Max)v
Power Consumption-Cooling	11 kW (Max)	28 kW (Max)
Container Size	20' ISO Shipping Container	40' ISO Shipping Container

Power & Water Specifications

Final Heat Rejection Type–Included[2]	Evaporative Cooling Tower	Evaporative Cooling Tower
Water Usage-Estimated	2-3 GPM(standard garden hose)	5 - 8 GPM(standard garden hose)
Input Voltage	3 Phase 380 - 415 VAC, 50-60 Hz	3 Phase 380 - 415 VAC, 50-60 Hz
Nominal Input Amperage	400 A	1,200 A

Delivery & Installation

Shipping Lead Time	• 12-14 weeks after receipt of purchase order
Shipping Terms	• Ex Works
Installation	• seven to ten days for the first unit, plus five to eight days for every subsequent unit. GRC service engineer/s will be onsite to perform installation and training.

Site Preparation

Customer to Provide

- Installation site preparation, including electric & water utilities
- Grade level concrete pad (12' x 30' for H20c or 12' x 50' for H40c)
- Applicable permits
- Cooling tower water treatment chemicals, water connection for cooling towers
- Firewall / routers / switches
- Internet connection for miners (estimated at 10kB bandwidth per miner)
- Bitcoin wallet and mining pool account

Warranty & Maintenance[3]

- One-year limited warranty on infrastructure only, including:

- 24/7 on-call GRC support staff
- 24/7 remote monitoring
- Domestic (US): Three business-day onsite response with parts and labor - International: Parts by mail
- Regular maintenance as required

- Extended warranty & maintenance contracts available for an additional fee

Antminer units not included. Other miner hardware supported. Contact GRC for more info.

[1]Low water use options available in climates where the design dry bulb temperature does not exceed 90°F / 35°C.

[2]Additional costs apply.

[3]Warranty is Void if the HashTank units are run outside of their operating parameters defined in the installation specification.

ElectroSafe™ Dielectric Liquid Coolant

“Not all liquids are made equal. ElectroSafe conducts heat, not electricity!“

Frequntly Asked Questions

What is ElectroSafe?

ElectroSafe is an odorless, non-toxic, single phase coolant that is both electrically and chemically inert. The proprietary blend of high-performance fluids is the result of years of development, testing, and deployments across the globe. ElectroSafe ensures maximum performance and material compatibility, enabling your servers to run efficiently and reliably.

How long does ElectroSafe last?

Unlike 2-phase coolants, ElectroSafe does not evaporate under normal operating conditions and does not need to be replaced through the life of the data center (15+ years). Commercial deployments dating back to as early as 2010 have been running efficiently and reliably, with no degradation in the coolant’s performance.

How is ElectroSafe better than 2-phase coolants?

2-phase coolants, as the name suggests, change phase from liquid to gaseous, i.e. evaporate at low temperatures. ElectroSafe, being a single-phase coolant, does not evaporate and stays in the liquid form under normal operating conditions. This enables GRC’s patented open rack design that gives you easy access to your servers, making hardware maintenance a breeze.Further, the lack of evaporation eliminates the risk of inhalation related health risks that 2-phase coolants pose.ElectroSafe is non-toxic and unlike 2-phase coolants, does not pose significant health risks due to inhalation or ingestion.Finally, in terms of material compatibility and hardware reliability, ElectroSafe is compatible with virtually every electronic and IT component and system out there. GRC has been commercially deploying ElectroSafe based immersion cooling systems since 2010 and hardware reliability data shows a significant increase in the Mean Time Between Failures (MTBF). Initial tests with newer 2-phase coolants have shown some potentially catastrophic issues with material compatibility. A recent study by Lawrence Berkley National Lab (LBNL) showed dramatically high failure rates and vapor leaks that cost 368% of the total cost of IT equipment energy consumption.

How is ElectroSafeTM different from mineral oil or other single-phase coolants?

Back in 2009, when GRC was founded, we started with simple white mineral oil, but we noticed some material compatibility challenges that it posed with certain types of materials. Hence, we set out to find the ideal fluid that would offer better material compatibility without compromising the ease of material handling and safety that mineral oil offers. ElectroSafe is the result of that research and development process that has now proven its efficacy in commercial deployments across the globe, since 2010.

ElectroSafe Coolant Characteristics

Evaporation Rate	Nil
Percent Volatile	Nil
Auto-Ignition Temperature	~605° F (343° C)
Density	7.10-7.13 lbs/gal

ElectroSafe NFPA 704
Fire Diamond

What about fire safety and building code?

ElectroSafe™ coolant has a very low flammability rate and does not readily ignite. The National Fire Prevention Association (NFPA) 704 diamond rates ElectroSafe as a 0-1-0 substance.
This means that it poses no health hazard, has a high flashpoint, and is stable even under fire exposure conditions.

Foresight: Data Center Monitoring Software from Green Revolution Cooling (GRC)

"A watched pot never boils"

Managing data center infrastructure can be a daunting task, especially given the mission critical nature and, often, remote location of these facilities. Whether you're running an AI or high-frequency trading operation, computing on the edge or on premise, ensuring that your infrastructure is operating reliably and efficiently is critical to managing costs and uptime.GRC understands how important your data centers are to your business and are committed to doing everything it takes to help you not just achieve but go beyond 'five nines' of uptime. Hence, we developed the Foresight monitoring system to help you get there.

All of GRC's products come prepackaged with a plethora of sensors, control systems, and the Foresight monitoring platform. While the onboard sensors and control systems optimize the performance, efficiency, and reliability of the system, the Foresight web-app allows you to remotely and securely monitor the real-time and historical performance of your data center infrastructure.

Key Features:

- **Web-app:** allowsforsimple,secure,remote monitoring of your data center's performance,health, and efficiency.
- **Early Fault Detection:** regular system health checks look for early indications of potential issues. The diagnostic system looks into everything from filter life, to pump performance, heat exchanger efficiency, and more. Helping you be proactive in avoiding potential service events.
- **Configurable Email Alerts:** The Foresight web-app allows you to configure email alerts based on trigger events, giving you the peace of mind that you will never miss a thing.
- **Backed by 24/7 Remote Monitoring From GRC:** all GRC products are backed by a limite done- year warranty that come with 24/7 remote monitoring by GRC support staff and technicians, giving you an additional set of eyes, and an added layer of protection. Extended warranty and maintenance contracts also include remote monitoring and diagnostics, beyond the one-year warranty period.
- **Integration with BMS & DCIM Systems:** Already have a BMSorDCIM system that you like?No problem, Foresight can easily integrate with almost any DCIM or BMS out there, using industry

Informative Web-Based Dash Board & Reporting

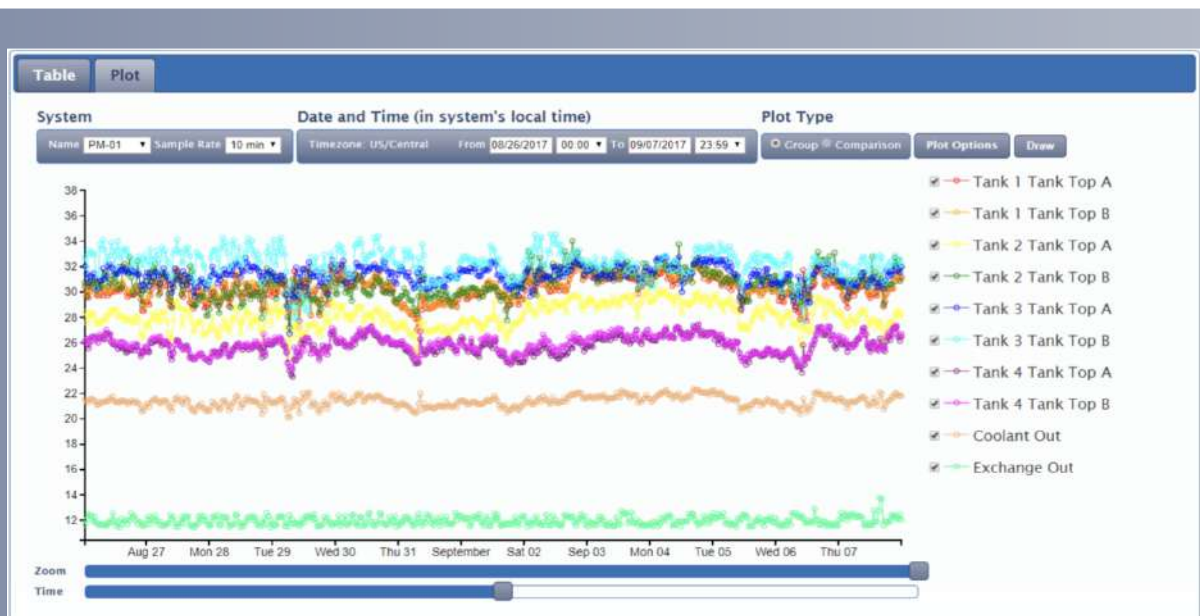

Fig 1. Visualize trends and historical data with Foresight's plotting feature.

	System Status	Set Point	Coolant Pump Speed / Status	CT Fan Speed / Status	Water Pump Speed / Status	Tank TopA 1	Tank TopB 1	Tank TopA 2	Tank TopB 2	Tank TopA 3	Tank TopB 3	Tank TopA 4	Tank TopB 4	Heat Dissipation	Water In	Water Out	Primary Current	Backup System Active	Time Elapsed Since Last Update
In-House		35.0	3.0	N/A	N/A	21.0	21.6	N/A	N/A	N/A	N/A	N/A	N/A	N/A	24.8	24.5	0.0	OFF	0:00:07
PM-01		33.0	3.0	N/A	N/A	30.2	29.8	29.8	31.6	31.6	33.0	25.6	25.6	N/A	N/A	11.3	6.3	OFF	0:00:10
PM-02		30.0	3.0	N/A	N/A	24.5	24.0	25.1	24.8	23.6	24.3	25.0	24.1	61251	13.9	17.0	6.9	OFF	0:00:05
PM-03		33.0	4.2	N/A	N/A	31.6	32.4	30.4	29.7	33.0	32.9	30.2	30.7	68655	10.6	15.6	6.4	OFF	0:00:09
PM-04		33.0	3.0	N/A	N/A	30.6	32.0	27.4	28.8	28.6	28.8	27.4	28.1	57269	14.0	14.0	3.8	OFF	0:00:07
PM-05		33.0	3.0	N/A	N/A	30.5	31.1	27.0	26.6	28.2	28.6	26.5	26.5	48572	12.2	15.1	4.8	OFF	0:00:11
PM-06		33.0	3.0	N/A	N/A	24.0	23.9	24.1	24.0	24.5	24.6	23.2	23.7	37161	11.2	13.5	4.9	OFF	0:00:08
PM-07		30.0	4.0	2.3	N/A	25.2	25.2	29.8	30.7	27.7	28.5	25.0	25.1	55954	12.8	16.5	5.4	OFF	0:00:05
PM-08		30.0	3.4	N/A	N/A	29.8	30.3	29.5	29.6	28.1	27.9	28.4	28.4	62152	37.4	14.6	5.9	OFF	0:00:10
PM-09		30.0	4.0	N/A	N/A	30.9	29.0	27.8	28.6	29.0	27.0	27.8	28.1	42783	11.0	17.8	7.3	OFF	0:00:09

Fig 2. See real-time status and data from each one of your racks in a simple table view.

Fig 3. Drill down to individual CDUs to track historical trends in a tabular or graphical form.

All in all, Foresight is a virtual, relentless employee on your data center team, that's obsessed with performance and uptime. Foresight allows you to focus on bigger better things, with the peace of mind of knowing that everything is running smoothly and efficiently.

Guide to Floor Space Optimization with GRC

Learn How We Are Improving Data Center Space Utilization

An analysis of how Green Revolution Cooling's (GRC) technology can help improve space utilization in server rooms, allowing for less infrastructure, denser layout, and denser racks.

Introduction

Floor space usage in the data center is an important consideration, whether planning a retrofit or a green-field build. There are a number of ways floor space utilization is measured for data centers: Watts per square foot, servers per square foot, square feet per rack, etc. This guide will explore how the ICEraQ™ is able to save valuable floor space in the data center as compared to traditional architectures.

More Compute per Square Foot

While it is true that a horizontal rack has a larger footprint than a conventional vertical rack by itself, this does not necessarily translate into higher space requirements for the complete data center. In fact, when looking at the fully burdened footprint (including aisle space, and supporting infrastructure like power, backup, and cooling systems) the ICEraQ can improve the space utilization of a traditional data center in terms of watts per square foot, servers per square foot, and even square feet per rack.

Further, because the ICEraQ also reduces the power drawn by the IT equipment itself, it is able to support more computing power for the same amount of watts per square foot. For example, assuming that a standard server draws 350 watts of power, a traditional data center with a rack density of 6 kW per rack would support 17 of these servers. The ICEraQ on the other hand, can reduce the same server's power rating by up to 20%. So if the server were to draw only 280 watts, the same 6 kW per rack would support 21 servers, that's an increase of 25% in computing power without any increase in rack density or power consumption.

In addition, the ICEraQ not only supports higher rack density but also a tighter layout and simplified infrastructure, all of which help further improve the space utilization in the data center.

Less Infrastructure

The ICEraQ, completely eliminates the need for large CRAC or CRAH units, while significantly downsizing the needs for backup generators, UPS and batteries, etc.

GRC's technology reduces the power drawn by servers and that needed for cooling. This reduces the overall electrical power rating of the complete data center, for the same amount of computing power. The elimination of bulky CRAC and CRAH units along with the downsizing of power and backup infrastructure helps save valuable real estate both inside and outside the data hall.

CarnotJet™ System

- ✔ Liquid submersion rack
- ✔ Pump modules (QTY 1 per 4-8 racks)

Air Cooling

- ✔ Air-cooled rack
- ✔ Air handlers
- ✔ CRACs
- ✔ Chillers
- ✔ Additional Generators/UPS/Batteries

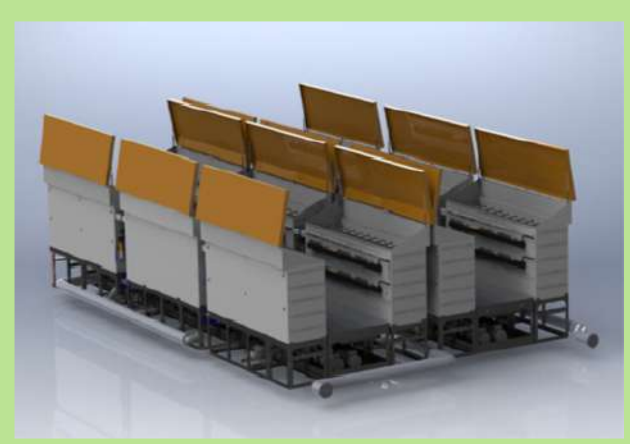

Low pro le pump module covered with tiles to double as a walkway

Back to back racks eliminate every alternate aisle

Further space savings can be found through the use of the low profile pump modules offered by GRC. These pump modules can be installed under existing raised floors or directly on a concrete floor. If installed on a concrete floor, the modules can be covered with tiles to double as walkways, creating a raised floor without the additional labor, costs, or complexity.

Denser Layout

The ICEraQ are installed back to back and end to end next to each other. There is no space wasted in creating hot and cold aisles, so every alternate aisle in a traditional layout can be eliminated completely, packing in the racks closer to each other without concerns of serviceability or thermal management.

Denser Racks

The ElectroSafe™ coolant used in the ICEraQ has 1,200 times the heat capacity of air, i.e. it takes 1,200 times the heat to raise its temperature by 1˚C. This higher heat capacity means that servers can be packed in closer to each other without concerns of heat transfer between them.
Further, since the bulky cabinets and chassis fans are not needed, denser server layouts are possible. Densities of 60 kW/rack are not uncommon with the ICEraQ, and can be increased to as high as 100 kW/rack. This is significantly higher than the average 5.8 kW/rack seen in traditional data centers

Sample Layout

The following drawing is an example of a typical layout of the ICEraQ racks, in a 9,000 square foot data center. The 90' x 100' data hall packs in ten rows of seven quads (collection of four racks and a pump module). That's QTY 280, 42 U racks with 70 low profile pump modules functioning as walkways.

Sample layout with ICEraQ

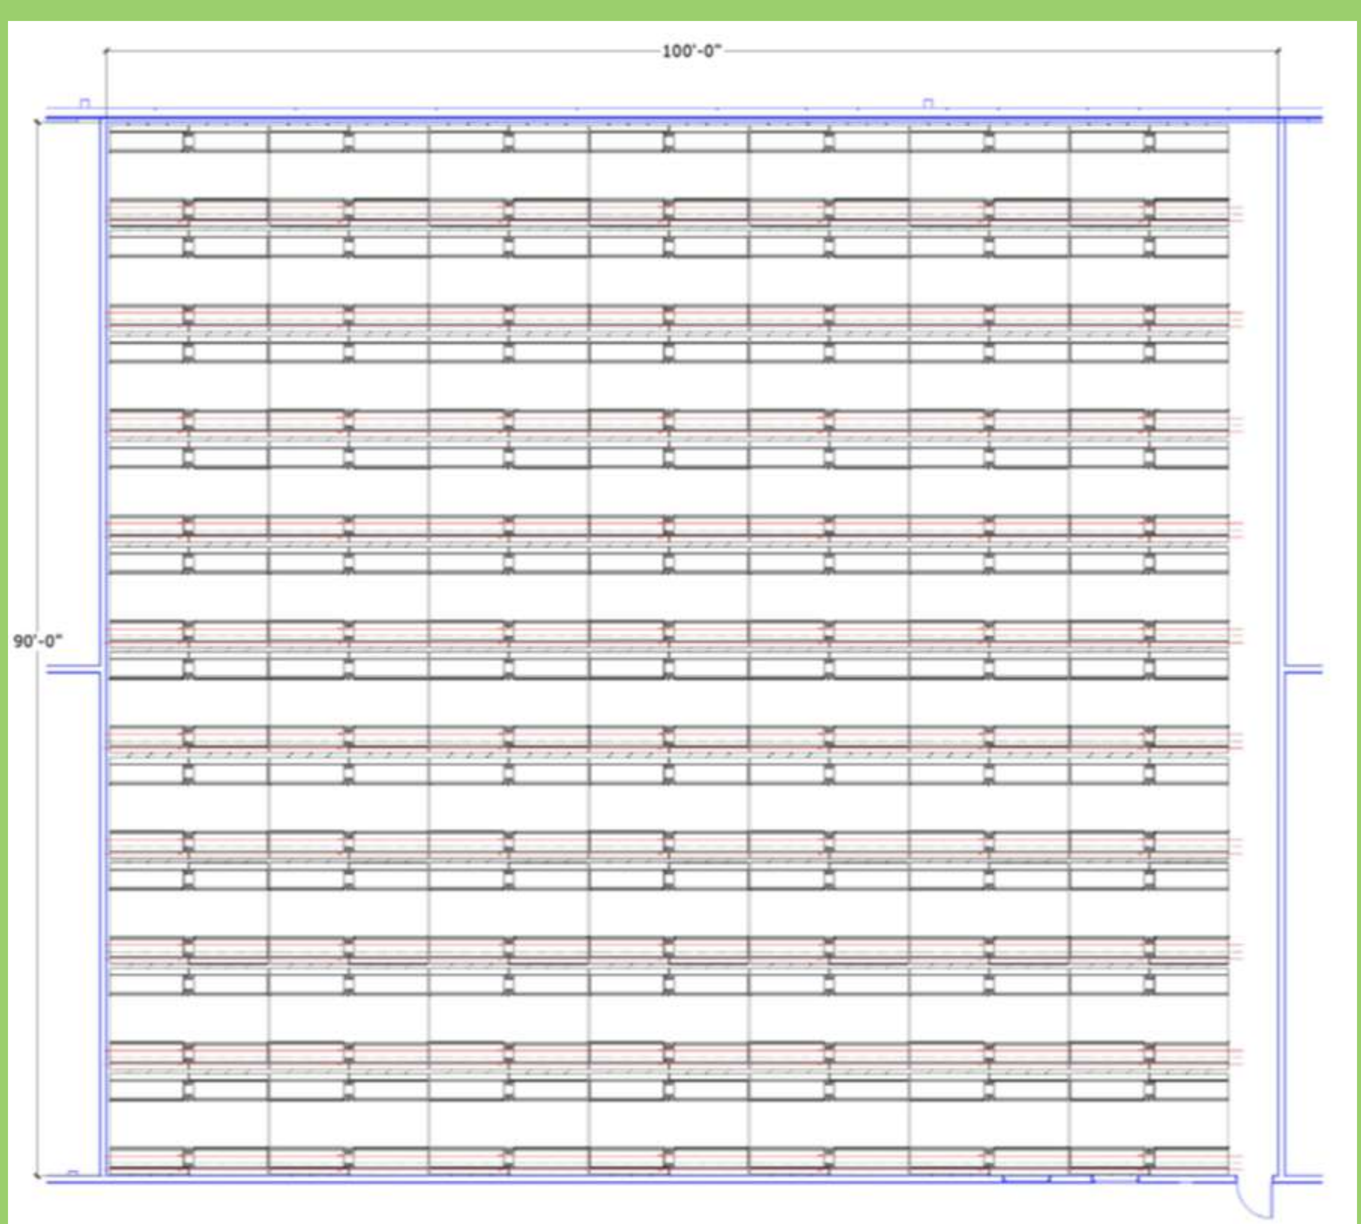

Total Area	9000 sq ft.
Number of Racks	280 racks
Effective Rack Footprint	32.14 sq ft/rack
Server Capacity	1.14 U/sq ft.

The 9,000 square feet, in this example can support over 14 MW of computing capacity, depending on the rack density. The table below shows the possible IT load and the power density of the data center at various rack densities.

Rack Density	10 kW / rack	20 kW / rack	50 kW / rack
IT Power	2.8 Mw	5.6 MW	14 MW
Power Density	311 W / ft.	622 W/sq ft.	1556 W/sq ft.

Conclusion

GRC's technology can help data centers get more from less. With the ever growing demands for computing power, more and more data centers are beginning to become capacity constrained. Our ICEraQ can not only optimize resource utilization today but, support virtually unlimited rack densities, that can prepare the data center infrastructure for the needs of tomorrow as well.

Guide to Data Center Sustainability Metrics with GRC's ICEraQ™

Going Beyond the Traditional Data Center

An analysis of how Green Revolution Cooling's (GRC) technology can help improve the Power, Carbon, and Water Usage Effectiveness (PUE, CUE, and WUE) even at some of the most efficient data centers.

Introduction

A few years ago, the industry average PUE was well above 2.0, which meant that organizations spent more on keeping systems cool than they spent on powering the servers. But over the years, this PUE number has steadily dropped. This drop has primarily been driven by hyperscale operators who have taken prudent steps towards minimizing their environmental impact. The popularity of free air cooling and the increased flexibility of IT equipment to operate at higher temperatures, has facilitated PUEs of as low as 1.15. While it may seem like the low hanging fruit is gone, technologies like liquid immersion cooling promise to bring in the next wave of efficiency improvement. Given the massive scale and exponential growth of data centers, a further reduction to PUEs below 1.05 can have a huge impact on not just the operating costs but on the carbon footprint and on the water requirements for the data centers as well. This guide will look at how GRC's liquid immersion cooling technology can help your data center reduce energy and water consumption, cost effectively.

Intro to Liquid Immersion Cooling

Immersion cooling, as the name suggests is where entire servers are immersed 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. But concerns regarding the cost, safety, and environmental impact of the coolants used led to limited adoption.

Here Comes ICEraQ

GRC has brought to the data center market an industrial scale liquid immersion cooling solution called The ICEraQ. The ICEraQ uses a proprietary dielectric coolant called ElectroSafe, which is a non-toxic, eco-friendly fluid that is a good conductor of heat but not electricity, making it ideal for the cooling of IT equipment. Electrosafe works with servers from any manufacturer, such as Dell, HP, Supermicro, IBM, Intel, SGI, Quanta, and more. The ICEraQ consists of four major components; a horizontal rack, the Coolant Distribution Unit (CDU), the control system, and a cooling tower*. Our ICEraQ is a horizontal rack which comes in 42U or 52U, fiiled with our Electrosafe coolant. Electrosafe captures heat produced by IT equipment and transfers it to a warm water loop via heat exchangers in the CDU.

*Low water use options available in climates where the design dry bulb temperature does not exceed 90°F / 35°C.

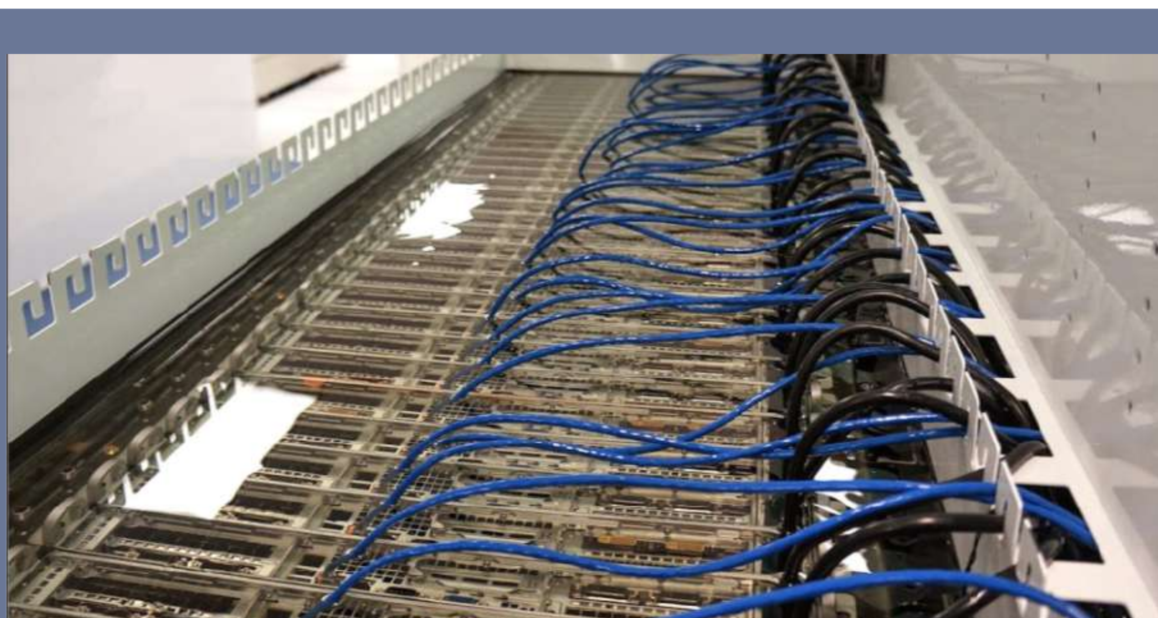

GRC's ICEraQ cooling system in action

Energy Savings

The ICEraQ yields a mechanical PUE of 1.02-1.03. This extremely efficient cooling is enabled by the use of the ElectroSafe coolant that has superior heat conductivity and 1,200 times the heat capacity of air (by volume). The superior heat conductivity means that ElectroSafe is much more efficient in both extract- ing heat from the equipment and subsequently expelling it out of the data hall. For example, ElectroSafe at 100˚F is as effective at maintaining optimal core temperatures as is air at 70˚F. Further, maintaining ElectroSafe at 100˚F requires significantly less energy than cooling water to 45˚F for air cooling.
Furthermore, the ICEraQ allows reduction in the power consumed by servers as well. This server power reduction is driven by the removal of server fans, and through superior thermal management allowing servers to run cooler. This improvement in energy efficiency directly translates into lower carbon foot print and reduction in the water required for cooling the data center.

Emissions Reduction

On average, one kilowatt-hour of electricity purchased in the United States generates about 1 pound of carbon dioxide. Comparing the carbon footprint of an extremely efficient traditional data center (avg. PUE = 1.21), and an ICEraQ enabled data center (avg. PUE = 1.09) of similar computing capacity reveals an annual Carbon emissions reduction of over 10,000 metric tons per annum, a 21% improvement over what is considered an extremely efficient facility.

The reduction in power directly relates to lower Scope-2 and Scope-1 (from diesel generators) emissions on the CDP (Carbon Disclosure Project), helping organizations achieve better scores on the CDP and reach their carbon reduction goals.

Further, the improved energy efficiency and the resulting decrease in carbon emissions, can give organizations access to energy efficiency grants and programs offered by local utilities, or state and federal governments.

50% Less Total Cooling Power

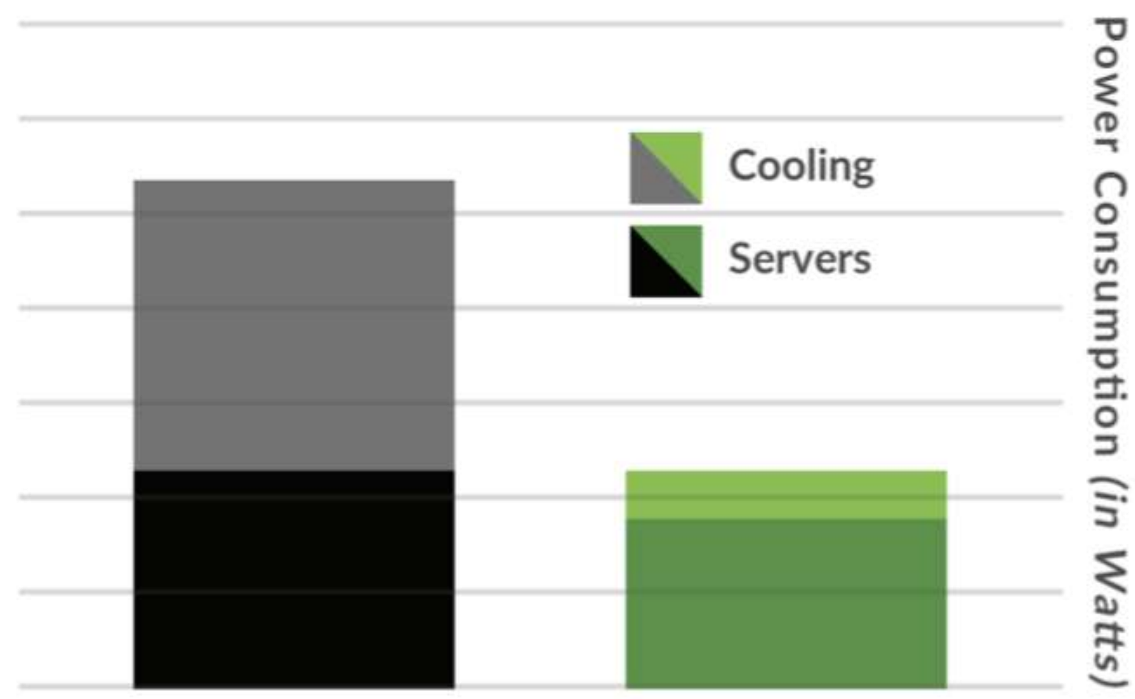

	AIR	GRC
Servers	24,000	24,000
Ave Server Power	350 W	280 W
IT Power	8.4 MW	7.392 MW
Cooling Overhead	15%	3%
Electrical Overhead	6%	6%
Effective PUE	1.21	1.09
Total Facility Power	10.16 MW	8.06 MW
Energy Consumption (Yr)	89 Million kWh	70 Million kWh
Carbon Emissions (Yr)	50 Million kgCO2eq	39 Million kgCO2eq
Reduction w/ CarnotJet		21%
Effective CUE	0.68	0.61 kgCO2eq/kWh

CUE or Carbon Usage Effectiveness is a metric defined by Green Grid. CUE is defined as follows:

$$CUE = \frac{\text{Total CO2 Emissions caused by the total Data Center Energy}}{\text{IT Equipment Energy}}$$

Water Savings

Apart from the indirect water savings related with the generation of electricity, the ICEraQ also helps directly reduce the water use and requirements at the data center.

"Water is tomorrow's big problem. The water consumption (in data centers) is super embarrassing. It just doesn't feel responsible. We need designs that stop using water."
-Amazon's James Hamilton

The ICEraQ enables the use of a hybrid cooling tower or an adiabatic dry cooler, both of which only require additional water when ambient temperature crosses a set point (85˚ – 95˚F). Taking the same example as above, in a data center supporting 24,000 servers the ICEraQ helps save close to 300 million Liters or 80 million Gallons of water annually, that's close to 90% reduction in water consumption. Certain climates allow for the use of dry coolers in place of cooling towers, thereby reducing water consumption to zero. Examples of such installations include VSC in Vienna, PIC in Barcelona, and the US Air Force.

	Traditional	GRC
Total Facility Power	10.16 MW	8.06 MW
Daily Site Water Usage*	908,400 L	101,270 L
Energy Source Water (Yr)**	160 Million L	127 Million L
Site Water Usage (Yr)	332 Million L	37 Million L
Site Water Savings (Yr)		89%
WUE (Site)	4.51 L/kWh	0.57 L/kWh
WUE (Source)	6.68 L/kWh	2.53 L/kWh

*Based on James Hamilton's estimate

**US average of 1.88 L/kWh

WUE or Water Usage Effectiveness is a metric defined by the Green Grid. It can be defined as follows:

$$WUE = \frac{\text{Annual Energy Source Water Usage + Annual Site Water Usage}}{\text{IT Equipment Energy}}$$

Conclusion

Liquid immersion cooling can drastically improve your data center's Power, Carbon, and Water Usage Effectiveness, giving you a head start on your sustainability goals. GRC is supplying industrial-grade liquid cooling solutions to 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.

GREEN REVOLUTION COOLING

The Immersion Cooling Authority

Case Studies

United States Air Force: Containerized Data Center

U.S. AIR FORCE

Test Plan Developed by: Air Force Testing Authority

Background of the Problem

The United States Air Force has a need for a containerized data center which:

- Is capable of being deployed in a wide range of environments and temperatures with minimal modifictions.
- Can be deployed rapidly (under24hours) with minimal personnel.
- Reduces the overall amount of support equipmen trequired.
- Reduces overall cost o foperation and maintain optimum operational environment.
- Increases reliability with reduced maintenance cost.
- Can be remotely monitored and controlled as needed.
- Reduces and minimalizes harsh environmental effects on equipment found in various deployment areas.
- Is viable as a stand alone data center or to augment existing data centers.

The suggested technology replaces air conditioning and air-handling infrastructure with a semisealed liquid cooling solution.

Data centers and the information processing capability they deliver are indispensable to warfighting missions. Much of the current generation of Air Force data centers do not meet strict specifications for redundancy (Core Data Center Reference Architecture) and efficiency (AFI 90- 1701, Executive Order 13693). The cost to retrofit existing facilities to meet these requirements is exorbitant, both in terms of dollars and time.

Green Revolution Cooling (GRC), an Austin-based developer and manufacturer of immersion cooling systems which provide high-efficiency and high-performance data centers, aims to produce a commercially viable, fully-integrated containerized data center and cooling solution, called the ICEtankTM, for the US Air Force. The ICEtank uses a proprietary coolant called ElectroSafeTM. ElectroSafe is an odorless, non-toxic, single phase coolant that is both electrically and chemically inert. The proprietary blend of high-performance fluids is the result of years of development, testing, and deployments across the globe. ElectroSafe ensures maximum performance and material compatibility, enabling your servers to run efficiently and reliably.

ElectroSafe NFPA 704
Fire Diamond

The benefit to the Air Force will be the capability to deploy cost-effective and expedited data center build-out and expansion. Featuring total immersion of computer equipment in ElectroSafe, the technology replaces air conditioning and air-handling infrastructure with liquid cooling.

Summary of Research

Two modular data centers were designed with feedback from the Air Force, allowing several design optimizations and significant improvements based on field tests.

The Air Force Test Authority created a test plan to verify all stated functionality and energy savings and found the performance was even greater than claimed.

- 100% uptime for both containers since installation (now approaching nearly 3 years of cumulative testing.)
- Containerized data centers functioned as expected in both extreme cold and warm environments.
- 18.1% less power required to run servers using off-the-shelf servers commonly used by the Air Force.
- Servers Designed for Immersion (SDI) work reliably in ElectroSafe and have increased density and efficiency over air cooled counterparts.
- 93.1% PUE reduction in cooling power from 1.45 to 1.037.
- 41.4% overall reduction in total power usage (computers and required cooling.)
- Power per rack unit increase of 314% due to ability to operate more compute in denser space.
- Allowing smaller data center footprint requirements for the Air Force. One 40ft ISO container with four immersion cooling racks can provide the same compute as 153 air cooled racks.

Task 1: Definition of Specifications

Task 1.1 Kick-Off Meeting (complete)
Throughout the grant, Air Force and GRC employees have collaborated in taking suggestions and feedback to address issues or concerns and make the most relevant product to fulfill the mission for the USAF.

Task 1.2 Detailed Design Specification (complete)
Cooling performance was specified with a maximum heat rejection; each specification was based upon a certain environmental condition. The testing at Hill proved even greater capacity than specified due to conservative modeling.
Environmental Resilience: The ICEtank (“C3”) is expected to remain operational in extreme environmental conditions. GRC will quantify operating and storage temperatures, as shown in later testing; performance was verified.

GRC designed two different versions of specifications. C3 One (Hill AFB) and C3 Two (Tinker AFB) The specifications vary in the following ways:

- Numberofrackspercontainer(2-4)
- FinalHeatRejectionMethod-coolingtower(C3One)vsdrycooler(C3Two)
- Heatexchangerredundancy

The second unit (C3 Two) with a dry cooler has the advantage of greater range of extreme cold

environmental operational temperatures: with the ability to run between – 40 F and 95 F with minimal changes.

- Lower maintenance,higher reliability: with fewer components,the systems are less likely to fail.
- Faster installation time:deployed at any location and setup in less than one day.The net result is a design which will enable the Air Force even greater flexibility and speed during deployment.
- The lowest installation cost of any containerized data center.
- Highest efficiency of any product made by GRC.
- Advanced remote monitoring and controls
- The secondun it delivered will also be the first data center to achieve high efficiency
- (PUE<1.25) with zero water usage and without any requirement for a chiller.

Task 1.3 Code Compliance (Complete)

While code compliance does not apply to non-permanent buildings, GRC engaged code compliance engineers to design the container to meet code: including electrical system design, OSHA guidelines for ingress and egress, and mechanical ergonomics.

Task 2 Design

Task 2.2 CAD Work and Layout (Complete)

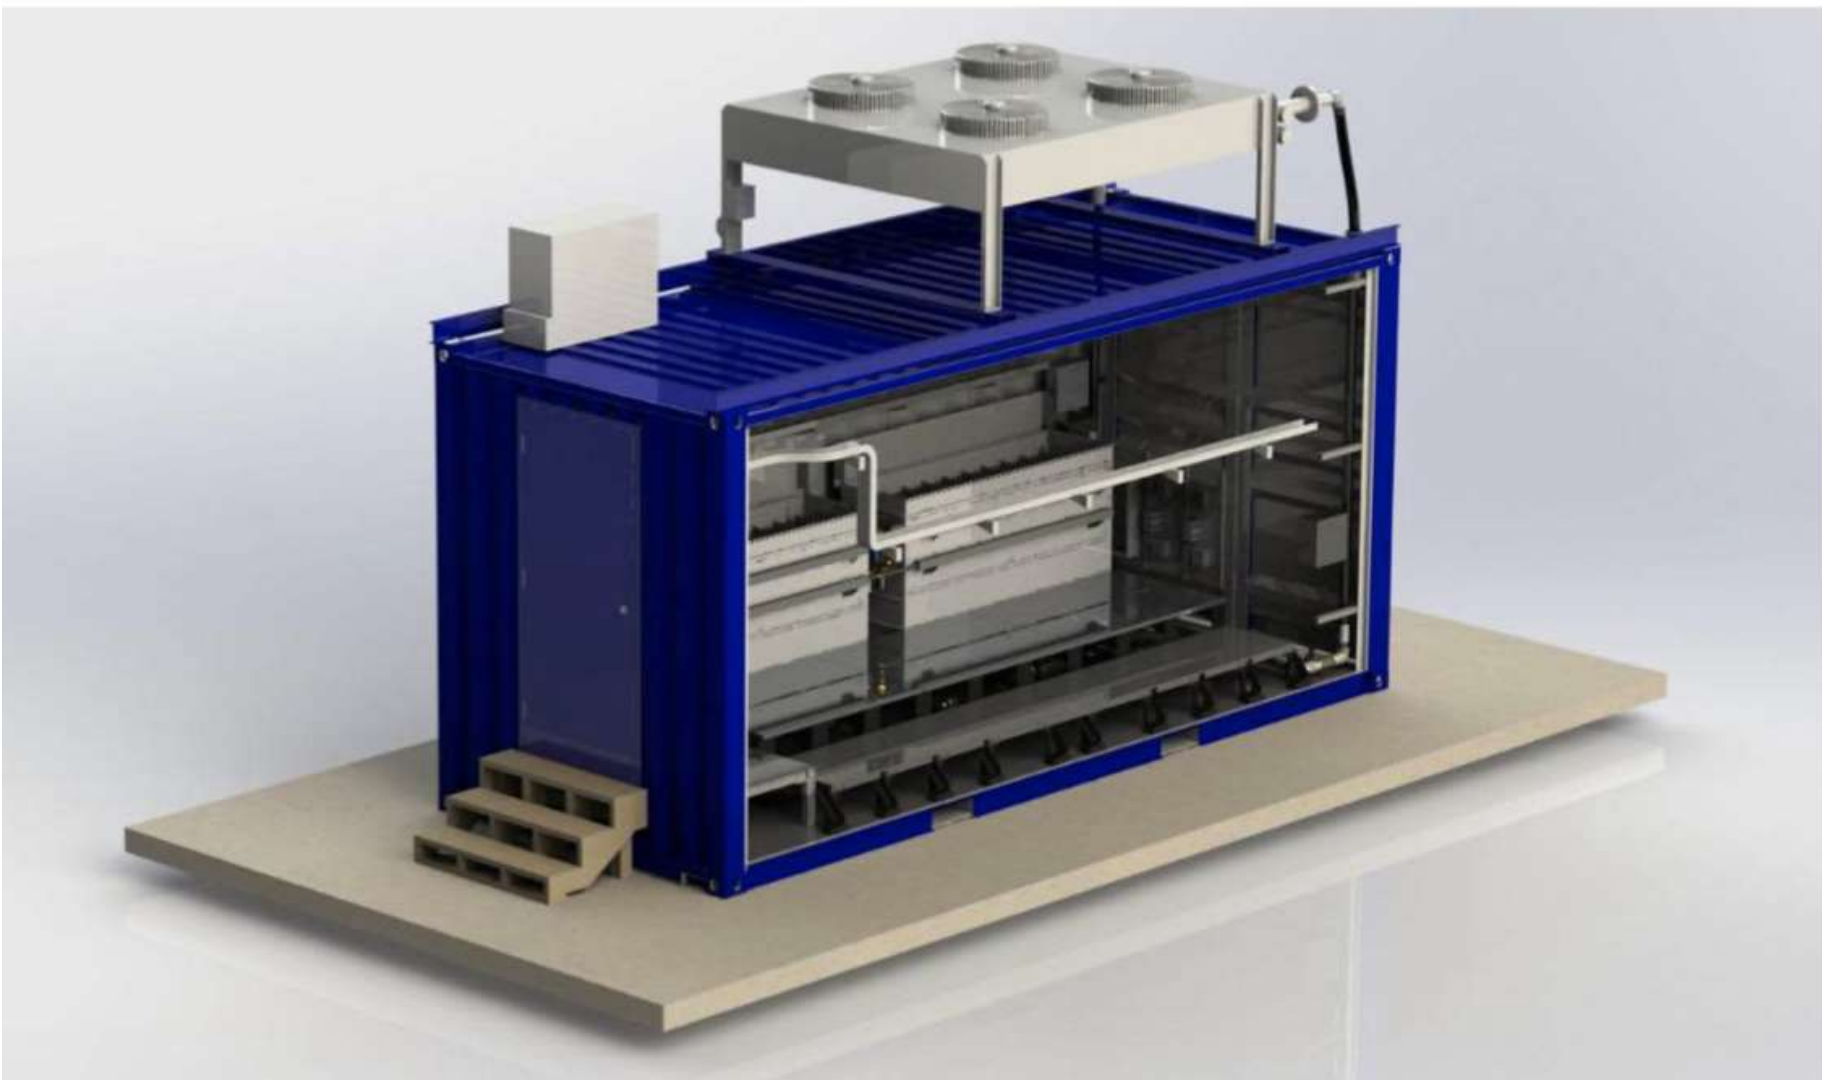

Image 1: (C3 Two – Tinker)

Task 2.3 ICEtank Integration (Complete)

GRC previously manufactured low-profile CDUs (Coolant Distribution Units) with a custom floor for use in raised floor environments, using the same concept, most of the supporting infrastructure was placed underneath the walking surface (See Image 1). The design focused on making the structure possible to securely ship, while being a quick and repeatable manufacturing process.

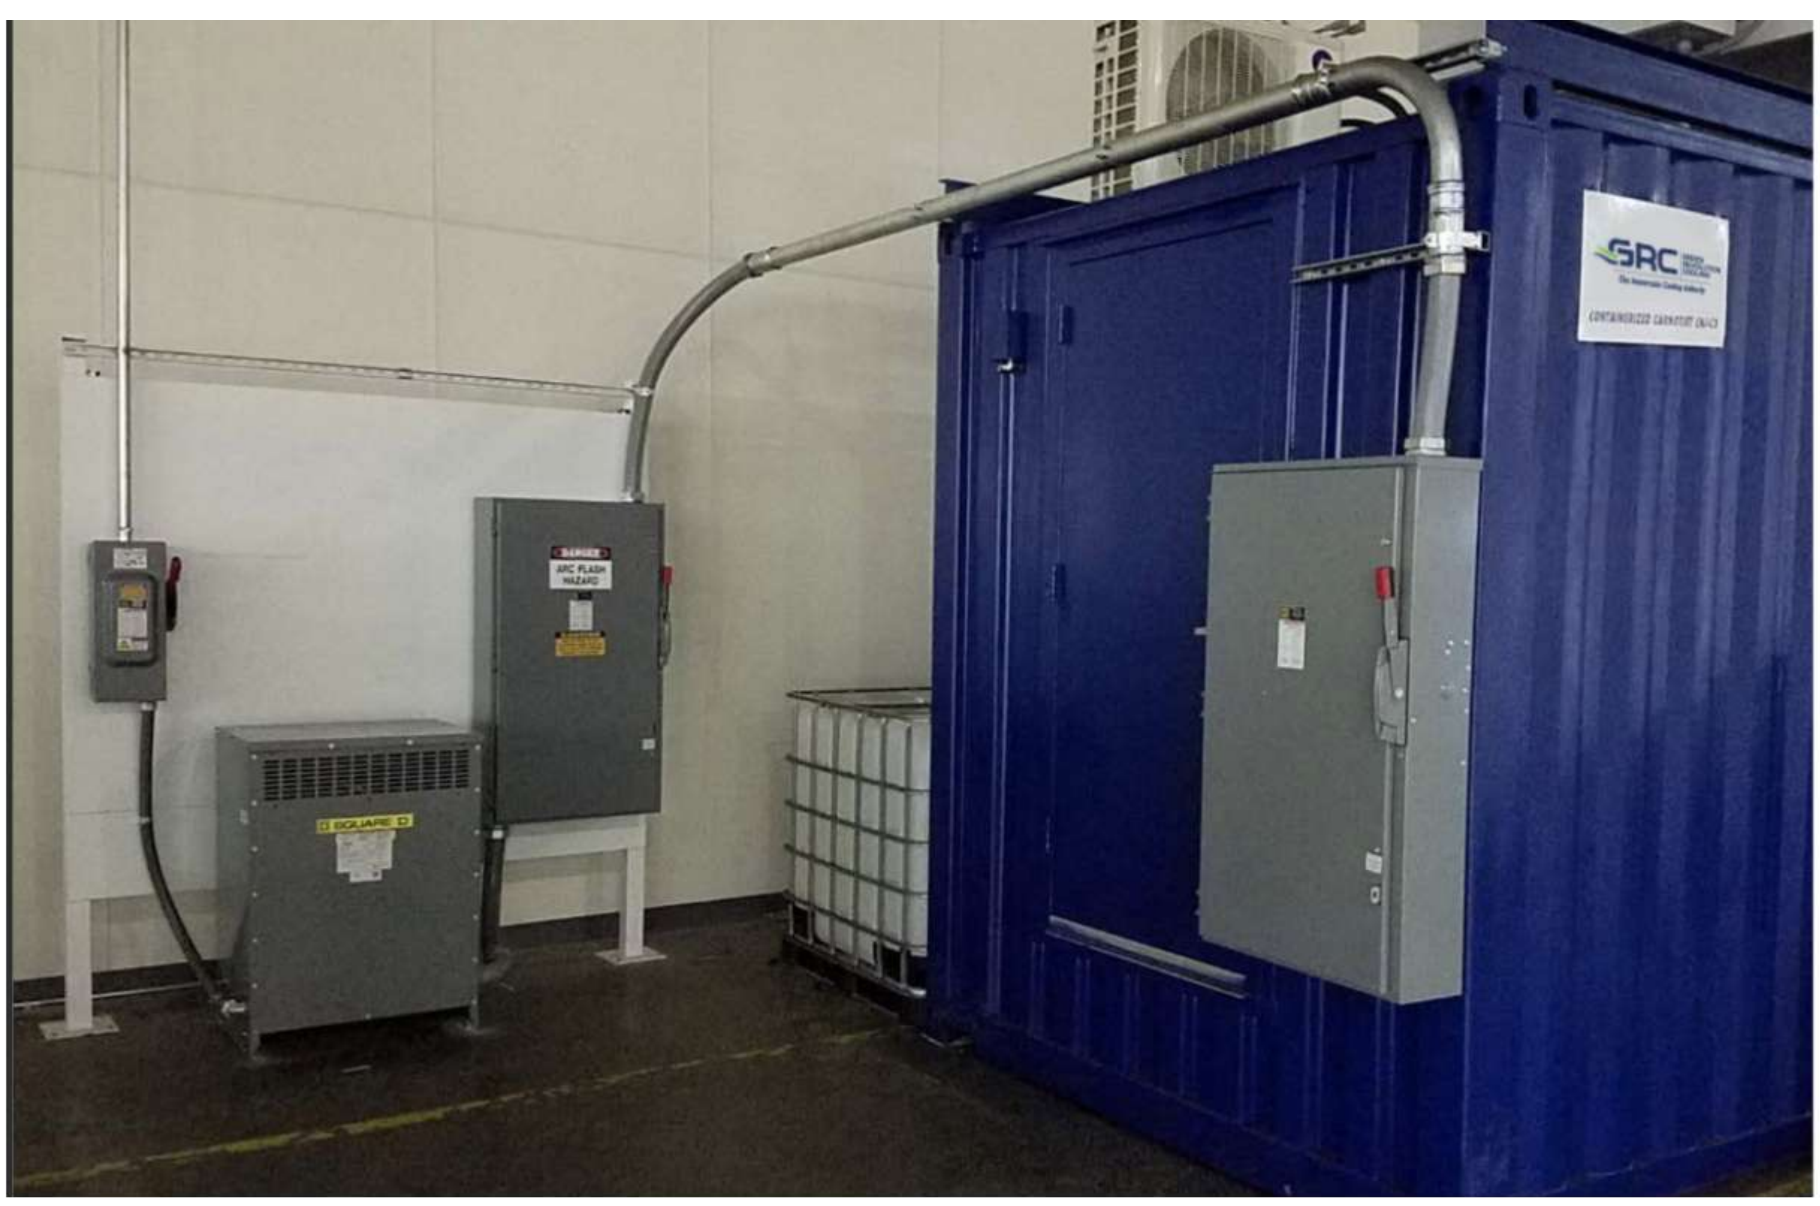

Image 2: (C3 Two – Tinker)

Task 2.4 Stress Analysis (Complete)

The containers and racks were designed to handle a seismic event, suitable for installation in seismic zones such as California.

Task 2.5 Connection Integration (Complete)

The container is designed for expedient installation, including making the necessary external connections: electrical, network, and water/drain. All electrical infrastructure downstream of the exterior cutoff is prewired, making for a plug-and-play installation. The only electrical connections required on site is wiring to the external cut off (arching conduit to the large box mounted on face on container) (See Image 2). Network connections are passed through the front wall.

In the container at Hill (C3 One), water and drain connections were made through the floor for secure water use and freeze prevention. The C3 Two container, with a dry cooler, has fewer connection requirements by not requiring a drain or water supply.

Task 2.6 Safety and Fire Detection (Complete)

GRC used the expertise of fire and code compliance professionals to create a safe environment. 2N redundant fire detection alerts both GRC and a local relay system. Various floors were tested to find a surface that maintains a safe, non-slip surface regardless of fluid being present. All lights are motion detecting and if power is lost, the exit sign and an emergency light illuminates the exit path to an easy egress path.

Task 2.7 Lift Mechanism (Complete)

A two-axis electric crane integrated into the container makes maintenance of the servers quick and easy. The crane is mounted onto the ceiling, allowing for full clearance when lifting servers. The hook (triangular bar) connects a server to the hoist where it can be lifted with a button push. This allows any size or weight of server to be used in the system without the need to be fully supported by personnel (See Image 3).

Task 2.8 Cooling Tower (Complete)

Image 3: (C3 Two – Tinker (Similar to C3 One – Hill)

Image 4: (C3 One – Hill)

A cooling tower and dry cooler model have been designed, constructed, and installed; both have advantages and disadvantages.

The cooling tower used in C3 One was designed to handle extreme temperature and has insulation, line heaters and an indoor storage tank. Early designs had minimum temperature of 10° F, final design can handle as low at -4° F. The cooling towers must also be able to perform reliably in extremely hot conditions, handling maximum expected temperatures, not just in Utah, but most any environment in the world.

GRC developed controls which previously did not exist in the market to fully automate and monitor the datacenter in a reliable manner. Most cooling towers rely on a traditional method of using a float valve which is prone to breakage. The design was changed to use an electronic water level sensor, fed into a control computer to monitors levels and activates a solenoid automatically to refill the tower when needed. In the event of electronic failure, a redundant analog system prevents overflow.

A pressure regulator at the container input sets the pressure to a consistent level. This allows a range of supply pressures to be connected safely without concern of connection availability.

Task 3: Fabrication

Two containers were fabricated; small design changes to the second container resulted in dramatic improvements in assembly time. Overall, assembly time was reduced by nearly 60%. Pieces were designed to be pre-assembled, then installed within the container.

Task 4: Testing and Installation

4.1: C3 One (Hill AFB) Functional Testing (Complete)

While GRC had existing quality control and testing for the standard rack and pump module, new functional testing procedures were developed for the new product to detect any quality control issues before leaving the manufacturing area. Before shipping, all pressure, temperature, voltage, current, etc. sensors are tested for accuracy and an operating test completed. Due to the advanced monitoring system, the multitude of sensors allow diagnostics of the system's functionality.

4.2 Installation of C3 One (Hill AFB) at DoD Facility (Complete)

The unit was shipped on a standard flatbed truck with an ancillary load of exterior insulation panels. It had some nonstandard sized items that were at risk for shipping damage. Container Two addressed these deficiencies by simplifying the design.

4.3 C3 One (Hill AFB) Operational Testing

A battery of tests were developed by the Air Force Testing Authority to definitively test specifications and energy savings. A brief summary of findings:

- Extremecoldtesting:teststheinsulation,automaticheatingelementsandremotewatersump can successfully operate in the winter environment on Ogden, UT. Operated successfully at 2 F.

- Hotweathertesting:acontainerisratedforaspecificheatrejection(kW)atacertain environmental condition. The validation test confirms the capacity. The container surpassed design specifications on heat rejection capability. At a world record temperature of 33.4° C WB the system is calculated to maintain a coolant temperature of 48.3° C; an acceptable temperature for liquid immersion.

- ElectricalSystemVerification:theelectricalsystemverificationtestedmaximumloadsofpumps and fans operating simultaneously. All systems operated without issue.

- Serverefficiency:aserverinairdatacenterranasoftwaretestusingLinpackanditspoweruse measured. The exact same server had its fans removed and immersed in the container's racks, ran the software test again using Linpack and its power use measured. Power usage decreased 7.2% on stress test, 18.1% less power in Air Force application.

- Coolingsystemefficiency:theaveragePUEofthecontainerwasmeasuredat1.037.The reduction in server power for the same server in ElectroSafe as compared to servers in air used 41.4% less total power compared to HEDC datacenter with a PUE of 1.45.

- SDI(ServersDesignedforImmersion):thesepurposebuiltserversaredesignedwithoutspacing for fans or air flow shrouds and eliminate the typical server case, allowing denser configurations reducing capital costs. A 40 x 8ft container with eight racks of SDIs can compute the equivalent of 153 racks of air-cooled servers, currently in use. Running the same computing task with existing servers and infrastructure would use 6.5x the power as what is consumed using SDIs in the container.

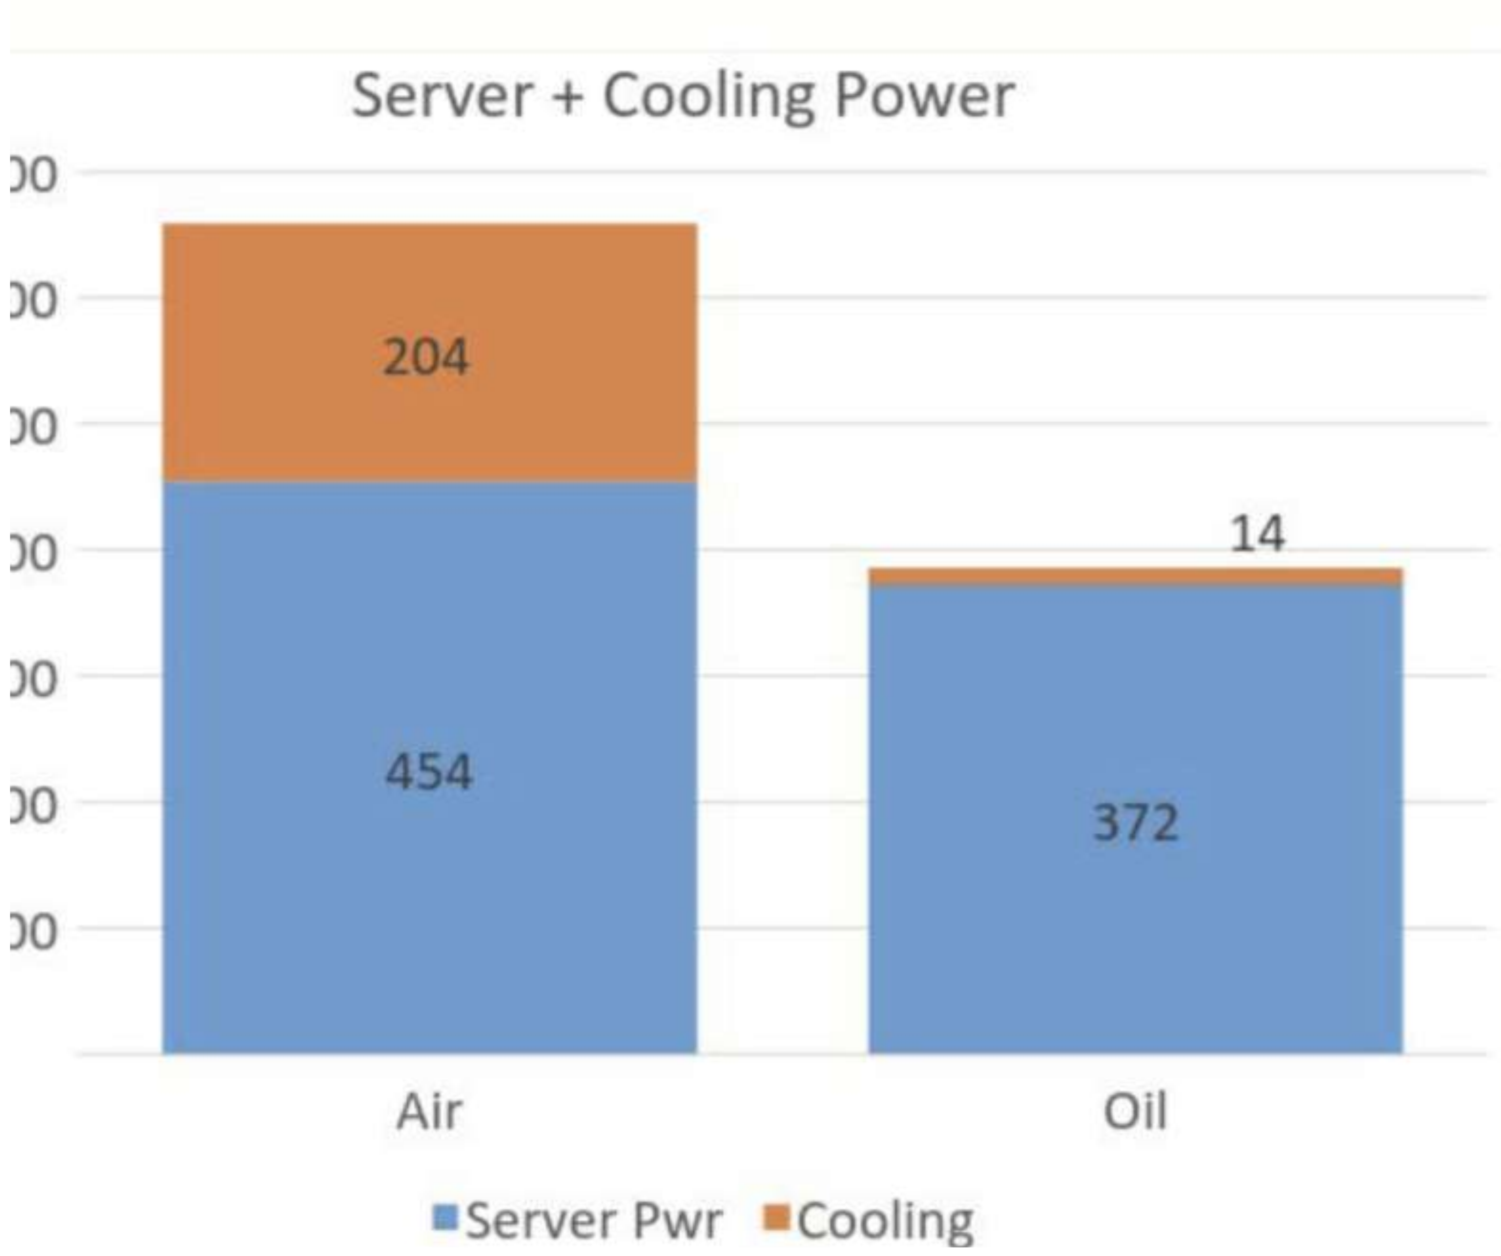

Figure 1: Server Power vs. Cooling Power (More details in appendix)

4.4 C3 Two (Tinker AFB): Functional Testing

The C3 two was functionally tested at GRC Headquarters. Due to the simpler system, and reduced part count, the quality control became easier. A new brand of PDUs were tested providing higher functionality including remote shutoff; useful to cut power if the temperature becomes too hot.

4.5 C3 Two (Tinker AFB): Installation and Operational Testing at DoD Facility

The container, final heat rejection, and ElectroSafe coolant was transported on a standard flatbed truck with no shipping damage.
Due to improvement in the design, the amount of time spent to install the container onsite was reduced by more than 90%+. In total, the unit installation at Tinker took a single person less than 20 man-hours total. The unit has had 100% uptime since installation.

Image 5: C3 Two-Tinker

In Summary:

These initial two ICEtank Data Centers have allowed for many design improvements to create the most automated, efficient, easy to transport and low cost modular data center. The improvements made between the two units provided dramatic improvements in simplicity, assembly time, and aesthetic appearance.

These improvements should also help with deployment within the Air Force. Without external insulation and with significantly reduced infrastructure on the roof, the unit can be transported just like any standard 40 ft. ISO container.

Image 6: Shipping the C3 Two using standard methods designed for ISO container movement.

Future Research:

The research and engineering done for the first two containers showed great progress towards making a viable production unit for the Air Force. Both containers are functional, but additional research is needed to make the unit as robust and fully featured as the Air Force would need for a production unit. Research topics for the third container include:

- Increased reliability of the cooling system and increasing environment operating range while increasing redundancy.
- Easier to maintain,including research to reduce water usage and treatment requirements.
- Improved manufacturing,further modularizing the manufacturing to increase output speed and quality.
- Reduced assembly time at destination with a target of one day installation.
- Easier shipping logistics with all components on ISO skid.
- Integration of UPS for servers to allow simplified servers and less redundant electrical infrastructure for same level of resilience.
- Integration of a fire suppression system.
- Increase reliability and functionality o fmonitoring system.

Appendix – Testing Results of C3 One (Hill)

Objective 1.1: Verify Functionality at Extreme Temperatures

1.1A Cold Temperature Testing

Metric: To avoid pipe freezing during expected winters, water temperatures must maintain at least 19.6° C higher than ambient temperature, with a 10kW or higher heat load, on the coldest day (-16.2° C dry bulb (DB)).

Result: PASSED. The lowest ambient temperature recorded was -16.7° C, and the corresponding water pipe temperature was 36° C higher, better than required 19.6 °C delta. The ICEtank remained functional and no damaged occurred.

Details: The ICEtank's specification states no freezing will occur at -16.2° C with a minimum 10kW load. Testing with less than 10kW indicates a better than specified performance. The coldest temperature recorded for this test was on 01 January 2016: -16.7° C. The heat load was 2.6kW, 74% lower than the specified 10kW. The lowest temperature recorded on the water piping was 19.3° C, a 36.0° C delta. The test results were that the ICEtank remained functional with no damage occurring.

1.1B High Temperature Testing:

Time	1/1/2016 7:53
Ambient (˚c)	-16.7
Basin Temp (˚c)	27.2
Exterior water line (˚c)	19.3
Heat Exchange Output (˚c)	28.3
Minimum Delta (˚c)	36.0
Heat Load (˚c)	2600

Table 1: C3 One Cold Temp Testing Results

Metric: Ensure system equivalently meets performance specification of 45° C (max) ElectroSafe temperature at 20° C peak wet bulb (WB). Performance is only attainable with cooling towers, heat exchangers, and entire system working in tandem. Therefore, all systems must function well for test to pass.
Result: PASSED. The ICEtank performs at or better than specification.

Details: A 20 °C wet bulb is an extreme condition for the region of Ogden, UT and has not yet been reached this year. Hence, performance data from lower wet bulb temperatures and 25kW load was used to calibrate models to extrapolate performance at 20° C WB and 50kW as per specification. Wet bulb s based on temperature achieved after evaporation and is the temperature relevant to cooling towers. ASHRAE1 developed the model to correlate heat load to cooling tower performance with resulting water temperatures. Water temperature data, along with varying wet bulb metrics were recorded with this system at 25kW and showed a linear relationship which can be e xtrapolated to 20° C WB. The model correlated performance at 50kW load to data at 25 kW. Using the model, performance is better than required to meet specifications.

1.2 Electrical System Verification:

Metric: Breakers must not trip under max load.
Result: PASSED. No breakers tripped during the test. Remote monitoring alerts received upon PDU overloading.

Details: The purpose of this test was to ensure functionality of the ICEtank's electrical systems at maximum electrical load. Without enough resistive elements to test the entire system at 100% load, the test was conducted one immersion rack at a time. The metrics for this test were a pass/fail based on breakers tripping or not using the following procedures:

- Connect all available heating elements into a single rack, loading the Power Distribution Units (PDU) between 80-100% of breaker rating.
- For each PDU branch, balance phases within +/- 2A, as is common practice.
- Distribute heating elements to ensure at least 25 kW of heat, the maximum rack rating.
- Record electrical loads using meters in PDU.
- Switch load to other PDU electrical feeds. Run for at least 4 hours.

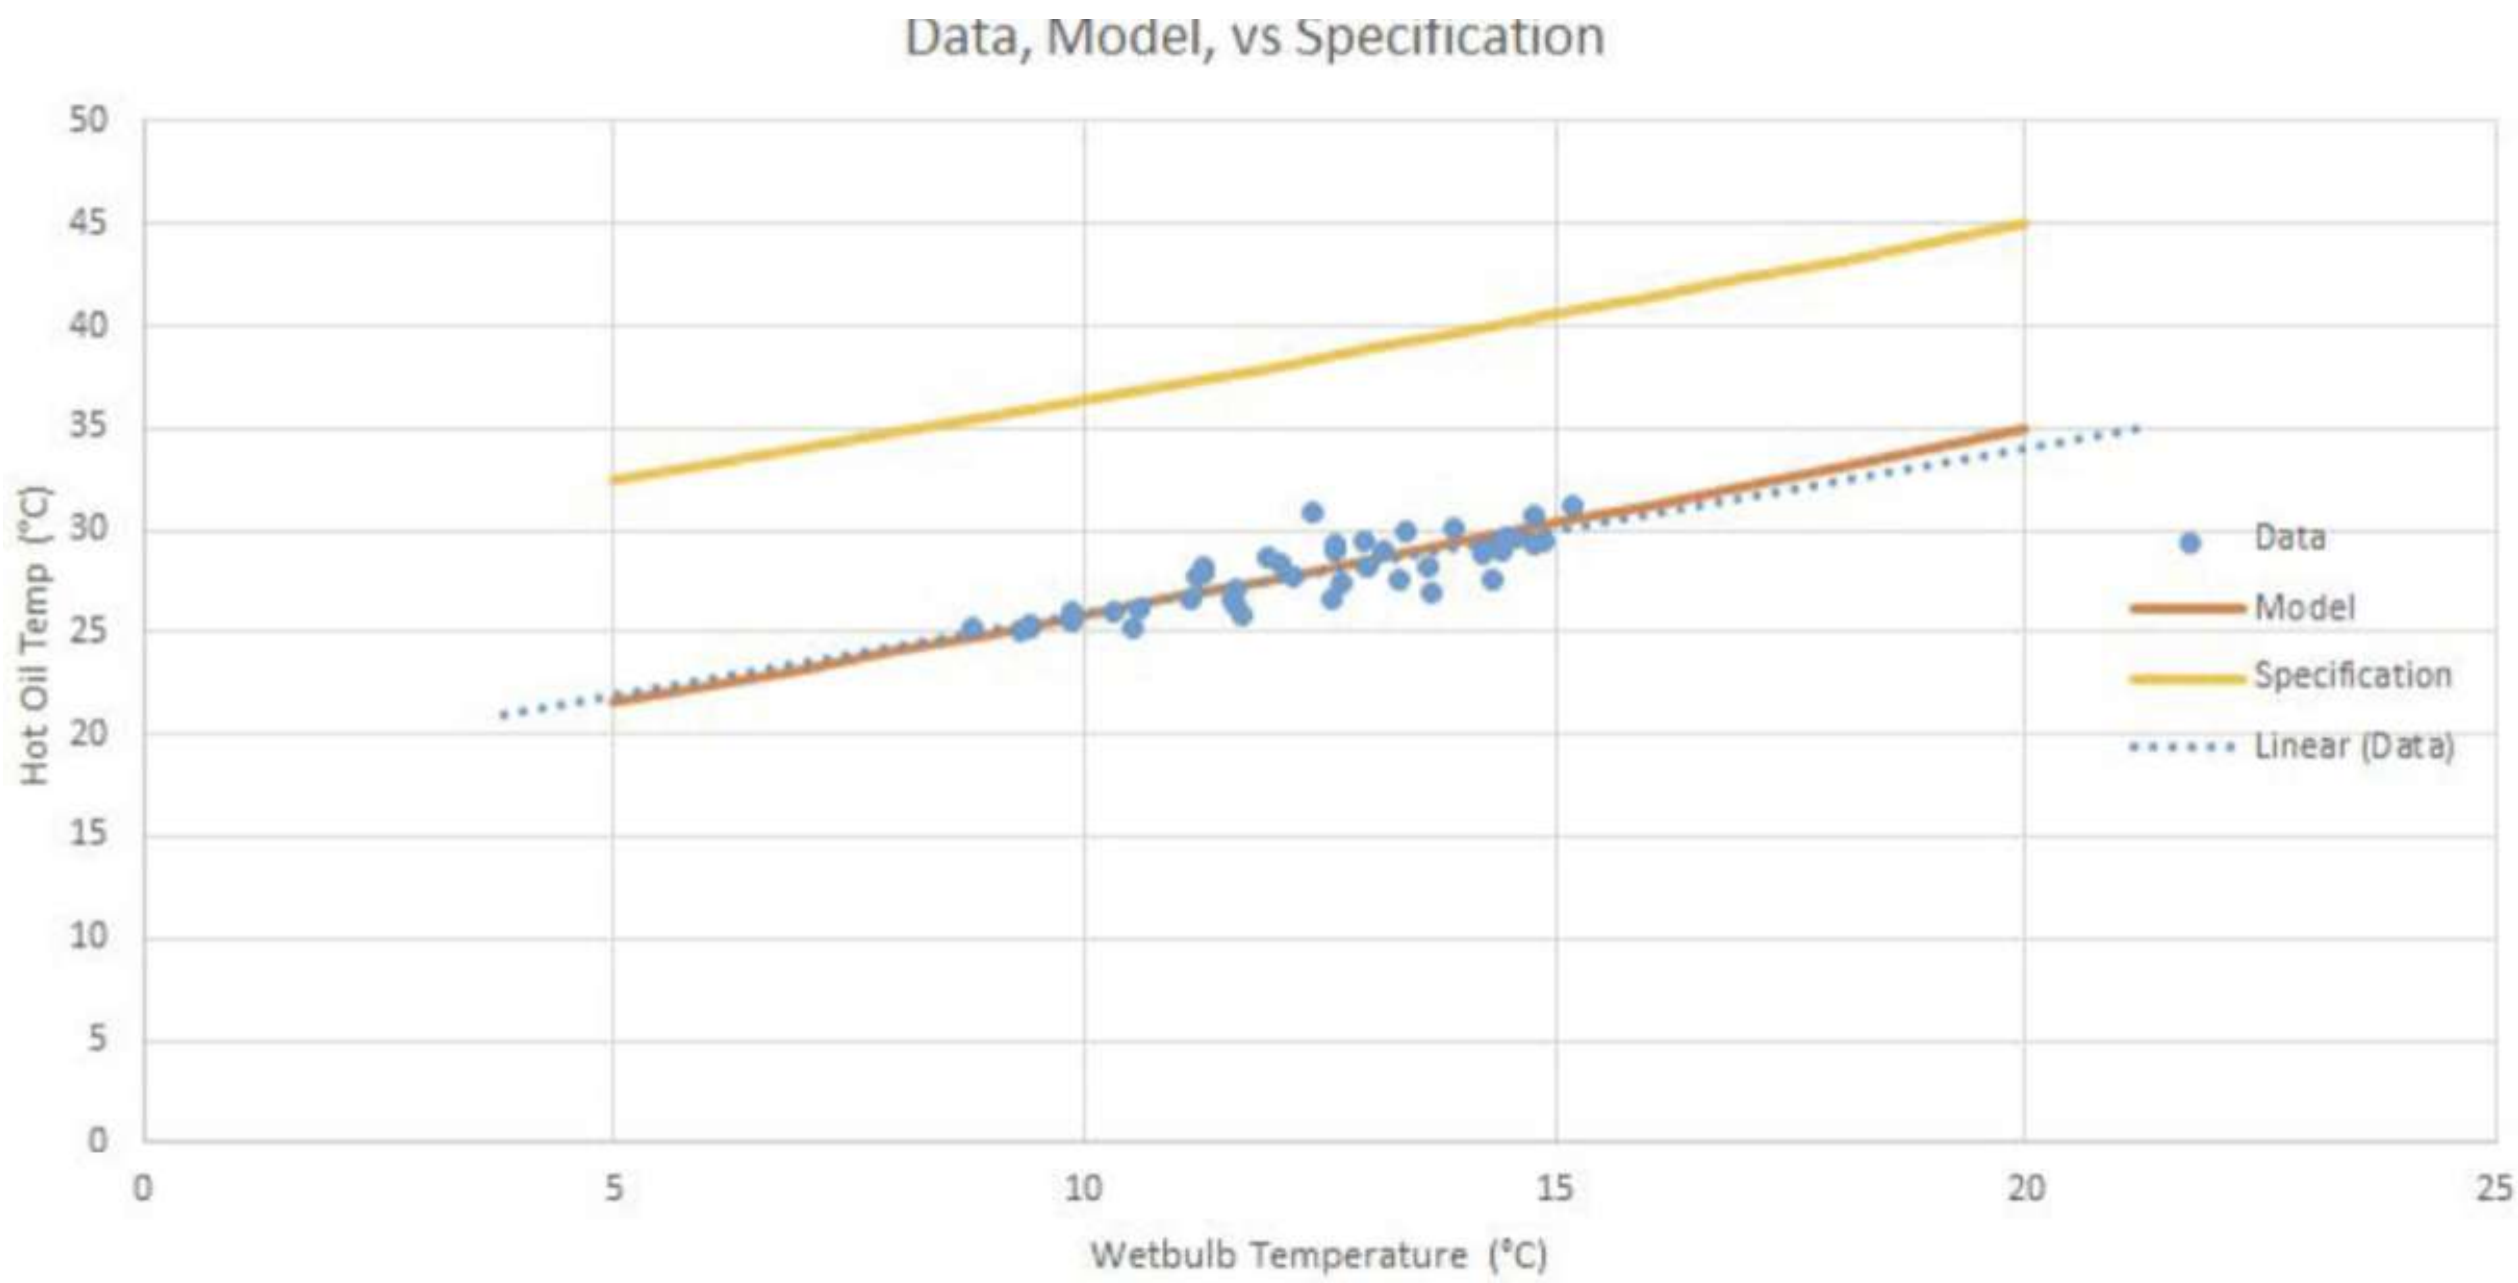

Figure 1: C3 One Hot Temp Testing Results

The ICEtank's CDU, and evaporative cooling tower ran at 100% capacity while overloading the PDU to test if the breakers would trip.

Breaker	Breaker size (amps)	80% of breaker per NEC (amps)	Peak amps measured at max possible load
PDUs	60	48	52.3-53.9 (measured per phase)
Primary CT Fan and Pump	20	16	16.6*
Secondary CT Fan and Pump	20	16	17.1*
Primary Coolant Pump	15	11.25	9.0
Secondary Coolant Pump	15	11.25	5.8

*Peak only possible with manual control.Automatic control maxes at fewer amps than 80% rating

Table 2: C3 One Electrical Testing Results

2.1 Server Efficiency

Metric: Server must use equal or less power per computation cycle immersed in ElectroSafe as compared to air.

Result: PASSED. Servers use 18.1% less power.

Details: Initially, several air-cooled servers were to be tested to compare immersion to air. However, preliminary measurements show server power varies significantly between adjacent servers.

	Power (watts)
Air-Cooled Server #1 in Air at Idle	454
Air-Cooled Server #2 in Air at Idle	480
Air-Cooled Server #3 in Air running production load	446

Table 3: C3 One Power use Test Results

With this discrepancy in mind, power measurements were taken from an air-cooled server; then the exact same server was modified (remove fans and thermal paste) for immersion in ElectroSafe and retested in same manner. The same “Wattsup” power meter is connected to the same power strip connected to the same power supplies in the same server running the same program; only difference is the cooling environment.

According to Air Force staff, the servers are programmed to maintain CPU readiness and therefore idle power is not minimized. A server used in Air Force production was measured and its power is approximately the same at idle. The measured power savings at idle will be representative of savings from converting the current datacenter.

The important metric is power per computational cycle and not just power usage, the software Linpack was used to calibrate computational speed. Linpack was run with the following configuration/settings; NB: 128, N 16,384. The problem required approximately 500 seconds to run. Computational speed in air vs immersion was equal within 0.2% so the computational speed is not affected per cooling type. Typically, computational speed is unchanged except for overclocking scenarios. Power savings identified Linpack represent a reduction in peak power, which can be used to downsize electrical infrastructure.

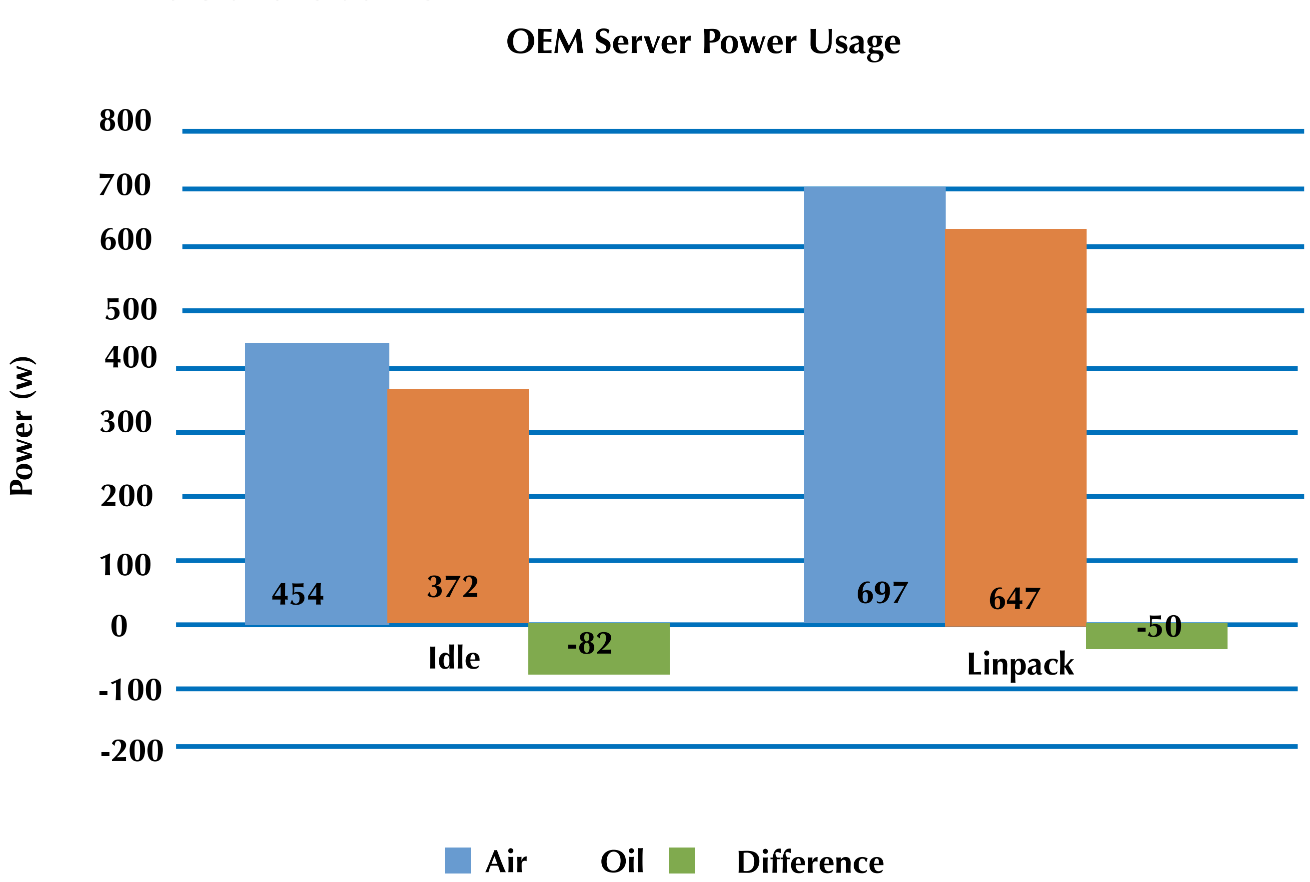

Figure 2: C3 One Power Usage Test Results

The power usage difference was especially significant at idle as fan power is a higher percentage of total power at low loads. Most data centers operate at near idle, the majority of time.	Air (W)	Dielectric Fluid (W)	Difference (W)	Difference (%)
Idle	454.4	372.1	-82.3	-18.1%
Linpack	697.4	647.3	-50.1	-7.2%

Table 4: C3 One Electrical Testing Results

2.2 Container Cooling Efficiency

Metric: Cooling system must have a Power Utilization Effectiveness (PUE) of 1.15 or lower. Result: PASSED. 93.1% reduction in cooling power from 1.45 PUE to a measured PUE of 1.037.

Details: Total cooling power usage includes CDU (including control computer, coolant pump, water pump) and cooling tower. Power measurements were taken over a 24-hour period than averaged.

Load	Average Power (W)watts
Control Computers	50
Coolant Pump	398
Water Pump	276
Cooling Tower	233
Total Cooling Power (sum of controls, fan, pumps)	958
IT Load	25687
PUE	1.037

Table 5: C3 One Power Consumption

2.3 Total Power Usage

Metric: Reduce total power.

Result: PASSED. Total power savings of 41.4% compared to same server in current datacenter.

Details: The reduction in server power and the low PUE combine to form the most important metric - total power usage. Total power usage was reduced by 41.4% compared to the air cooled datacenter; the total power usage is found by multiplying PUE by the individual server power.

Table 6: C3 One Total Power Usage

* An immersion-cooled server plus cooling power is less than the air-cooled server alone.

	Air (W)	Oil (W)	Δ (W)	Δ (%)	Equation
Server Power	454	372	-82	-18.1%	Measured (low utilization)
Cooling Power	204	14	-191	-93.2%	Measured: Server Power * (PUE-1)
PUE	1.45	1.037			(Cooling Power + Server Power)/Server Power
Total Power	659	386	**-273**	**-41.4%**	Server Power * PUE

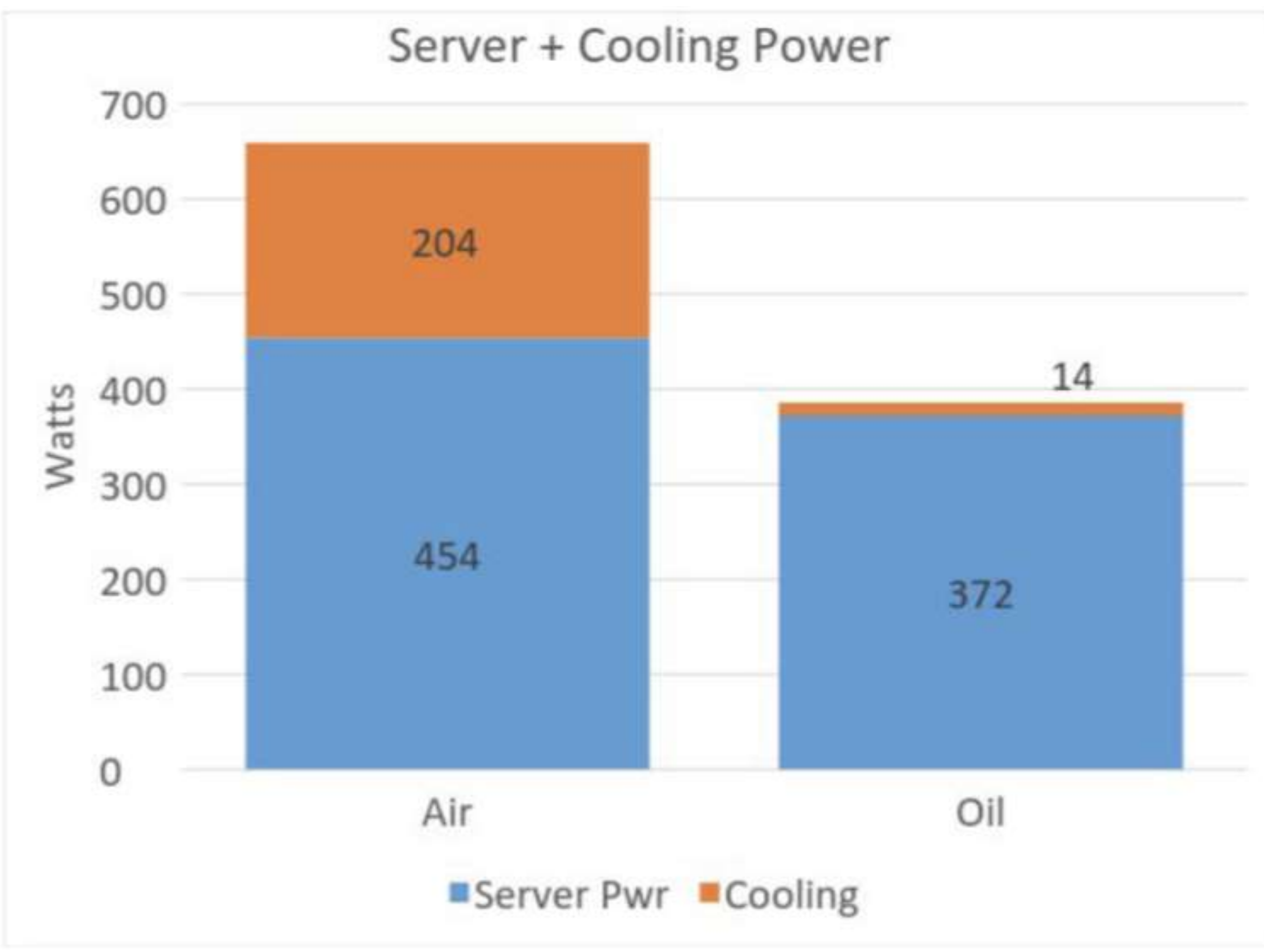

Figure 3: C3 One Server vs. Cooling Power

Addendum A - SDI Testing

Servers Designed for Immersion (SDIs) were chosen to reduce purchase price through eliminating the need for fans or a complicated chassis. Removing the unnecessary air components, the servers can be made denser, using less data center space for the same computational power.

New servers have more computational power in many metrics; they are packaged tighter, taking up less space and increasing rack density. The processing speed is increased per component, and the power to complete a computation is reduced.

Server Model	Immersed OEM vs. SDI	Immersed OEM	SDI	OEM in Air	OEM in Air vs. SDI
Space (U)	-88%	4	0.5	4	-88%
Linpack (W)	-44%	647	361	697	-48%
Idle (W)	-73%	372	99	454	-78%
FLOPS	139%	298	712	298	139%
Power/U	346%	162	722	174	314%
FLOPS/U	1811%	75	1,424	75	1811%
W/FLOP	-77%	2.17	0.51	2.34	-78%
Power/ 42U Rack	346%	6,794	30,324	7,319	314%
PUE	0%	1.04	1.04	1.45	-28%
Total: PUE* W/FLOPS	-77%	2.3	.53	3.4	**-84%**

Table 7: C3 One Immersion vs. Air Comparison of Servers

The power per U increased 314% so a 42U using the old servers would reject 7.3kW, SDIs will reject 30.3kW. With the increase in Power/U, new methods of cooling are necessary. The density of the new servers surpasses air cooling capabilities.

With increase in Floating Point Operations/Unit (FLOPS/U), the datacenter can become smaller. For the same task, modular data centers could provide equivalent computational power as large data centers have provided. To complete an equivalent task, the SDI servers require **1811% less space (i.e. air-cooled servers use 19 times the space). A 40 x 8ft container with 8 racks of SDIs can compute the equivalent of 153 racks of air-cooled servers, currently in use.**

With the increase in efficiency, Floating Point Operations per second/Watt (FLOPS/W), the total power to compute a task is reduced. As the SDI do not come with fans to reduce purchase price, they cannot be run outside of an immersed environment and therefore cannot test the same air vs immersion comparison. The OEM server demonstrated an 18.1% decrease in idle power, when immersed. Idle power of an SDI compared to the OEM server in air was reduced by 78%.

The efficiency upgrade by using the SDI in ElectroSafe compared to the the OEM server in air is significant. To compute the same task, the **SDI require 78% less energy** than the the OEM server in air.

Tying all the efficiency components together; the air-cooled server running in standard infrastructure will use 3.4W/FLOPS, an SDI in the immersion datacenter used 0.5W/FLOPS. Changing from the air-cooled data center using the existing server to SDI in ElectroSafe will use **84% less power (i.e. 6.5x the power for the OEM servers)** for the same compute.

Glossary

ICEtank: Liquid immersion cooling system for servers designed and manufactured by GRC.

C3: ICEtank (C3-one is at Hill AFB; C3-two is at Tinker AFB) ISO compliant

COTS: Commercial Off the Shelf. Standard servers designed for air-cooling that may be optimized for immersion.

DELTA: Difference between two things. E.g. the delta between 12 and 18 is 6.

ElectroSafe: GRC's proprietary coolant. ElectroSafe is an odorless, non-toxic, single phase coolant that is both electrically and chemically inert. The proprietary blend of high-performance fluids is the result of years of development, testing, and deployments across the globe. ElectroSafe ensures maximum performance and material compatibility, enabling your servers to run efficiently and reliably.

DRY-BULB: (DB) Ambient temperature, our normal understanding of temperature. Traditionally measured by expansion of mercury. Assume a temperature implies dry-bulb unless noted.

FLOPS: Floating-point operations per second is the number of operations the computer can execute. The most common measure of computational output.

HEDC: Hill Enterprise Data Center. The reference data center.

LINPACK: The most common benchmarking software to measure FLOPS. It measures how fast a computer solves a dense n by n system of linear equations Ax = b, which is a common task in engineering. Scores consistently when programmed the same but typically ran multiple times to ensure caching does not affect performance.

PDU: Power Distribution Unit. A big power strip that may have power metering capability.

PUE: Power Usage Effectiveness: Anything that isn't considered a computing device in a data center (i.e. lighting, cooling, etc.) falls into the category of facility energy consumption. PUE= (IT Equipment Energy + Facility Energy)/ IT Equipment Energy. PUE can be misleading: reducing energy use of a server while keeping the same cooling power calculates to a worse PUE, even though total energy use is lessened. The average PUE of US datacenters is 1.73, HEDC 1.45, GRC ICEtank 1.037, ideal PUE is 1.0.

SDI: Server Designed for Immersion (sometimes referred to as Built for Immersion)

WET-BULB: (WB) Coldest temperature reached by evaporation. Most relevant temperature to cooling towers. WB=DB at 100% humidity. WB is lower than DB at less than 100% humidity. Traditionally measured by measuring a wet temperature probe with air moving across it.

Tokyo Institute of Technology

The Immersion Cooled TSUBAME-KFC: From Exascale Prototype to The Greenest Supercomputer in The World

Abstract—Modern supercomputer performance is principally limited by power. TSUBAME-KFC is a state-of- the-art prototype for our next-generation TSUBAME3.0 supercomputer and towards future exascale. In collaboration with Green Revolution Cooling (GRC) and others, TSUBAME-KFC submerges compute nodes configured with extremely high processor/component density, into non-toxic, low viscosity coolant with high 260 Celsius flash point, and cooled using ambient / evaporative cooling tower. This minimizes cooling power while all semiconductor components kept at low temperature to lower leakage current. Numerous off-line in addition to on-line power and temperature sensors are facilitated throughout and constantly monitored to immediately observe the effect of voltage/frequency control. As a result,TSUBAME-KFC achieved world No. 1 on the Green500 in Nov. 2013 and Jun. 2014, by over 20% c.f. the nearest competitors.

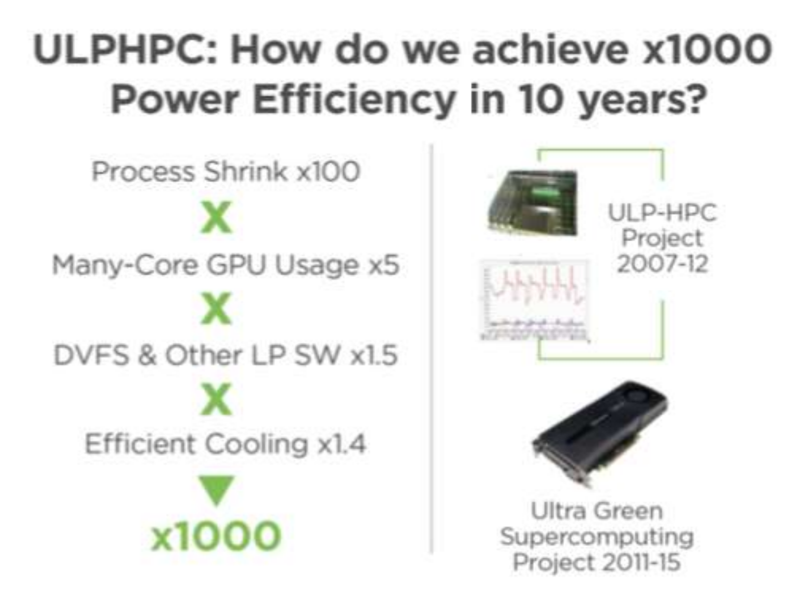

Fig. 1. The breakdown of 1000 times improvement in power efficiency in the ULPHPC project

I. INTRODUCTION

The most predominant issue towards future exascale supercomputers is power consumption, and the importance of being "green" has been, and still is a focus of many research and systems in supercomputing for the past 10 years at the least. Early systems employed low power embedded or ultra portable notebook processors such as Green Destiny [1], Blue-Gene/L[2], and MegaProto[3]. The DARPA exascale report established the goal of 20 megawatts for future exascale machines, analyzed the extreme challenges involved, and concluded that comprehensive low-power design would be required to even come close to the goal by 2018-20 time frame. This was followed on by the IESP (International Exascale Software Project) group, whose white paper was published to supplement the software requirements for achieving the 50 gigaflops/W goal[4], [5]. Nonetheless, the goal is still considered very difficult to reach without comprehensive research and engineering, with power reduction being the ultimate goal over all other objectives.

We have been extremely cognizant of the low power requirements supercomputers; there have been a number of research projects we have conducted, including ULP(Ultra Low-Power) HPC project conducted during 2007-2012, with the goal of achieving the 1000-fold improvement of power-performance efficiency(Figure 1). Many novel research achievements that harness substantial opportunity for HPC power saving. They include the use of new devices, new architectures, DVFS techniques, as well as extremely aggressive cooling strategies. As an example, the use of extreme many core processors, in the form of GPUs and Xeon Phi, have seen initial adoption as general purpose processors in the HPC arena, allowing up to several factors power-performance improvement compared to conventional CPUs [6], [7], [8]. There are now numerous supercomputers that employ many-core processors as their principle compute components, including the top two on the June 2014 edition of the Top500[10].

Tokyo Tech.'s TSUBAME2.0 supercomputer was designed ground up to accommodate GPUs as the principle compute engine, commissioned in Nov. 1st, 2010 to become #4 in the world on the Top 500. Approximately 90% of the FLOPs and 80% of the memory bandwidth are provided by over 4000 NVIDIA Fermi M2050 GPUs in the system. TSUBAME2.0 became #3 in the world in the Green500[11] which ranks the Top500 supercomputers in terms of their power efficiency, with 958 megaflops/W. Also it was recognized as the "greenest production supercomputer in the world" for Nov. 2010 and June 2011 respectively, as other machines on the list were small and non-production in nature. TSUBAME2.0 has recently been upgraded to TSUBAME2.5 by replacing the GPUs to the latest Kepler K20x, improving the Green500 power efficiency to 3069 megaflops/W, more than tripling the power efficiency and being ranked 6th in the world.

Nonetheless we are still more than a factor of 15 away from the 50 gigaflops/W goal for the 20MW exaflop machine. Moreover, we have expended the onetime architectural "leap", in that we have maximized the use of the architectural transition to extreme many-cores. As such, we are required to exploit other power conservation methods, such as DVFS and novel cooling methods, as well as further enhancing the many core usage, towards exascale. Thus, the follow-on project to ULP-HPC, the Ultra Green Supercomputing was proposed and funded directly by MEXT (the Japanese Ministry of Education, Sports, Culture, Science and Technology), to integrate the results and experience obtained from both the series of basic research and production TSUBAME supercomputers, and further deploy innovative power saving elements.As a main deliverable and research platform of this project, TSUBAME-KFC was designed, constructed, and deployed in the fall of 2013. TSUBAME-KFC serves the following

First, TSUBAME-KFC serves as a prototype for TSUBAME3.0 to be commissioned in 2016 as a successor to TSUBAME2.5, and features extensive exploitation of many core architecture in extremely dense packaging, numerous sensors and control features of power and thermals. It relies even more on many-core GPU for providing performance both in compute and memory. As such TSUBAME-KFC features over 600 Teraflops of single precision performance in a single rack, and could easily scale to a petaflop for TSUBAME3.0 while maintaining the per-lack power consumption to the level of TSUBAME2.0, or approximately 35 KWatts per rack.

Secondly, TSUBAME-KFC facilitates extremely efficient cooling technology via warm liquid immersion, developed by GRC, in an unmanned container environment as described in Section II. This method largely reduces power consumption for cooling since power hungry chillers are removed; as demonstrated in Section IV, the power consumption for cooling is less than 10% of IT power (power usage effectiveness is less than 1.1).

The overall scheme for achieving the 1000 fold power efficiency during the period of 2006-2017 is shown in Figure 1. Process scaling (the Moore's law) allows us to achieve nearly x100 increase in power performance the use of extreme many cores such as GPUs, plus proper software adaptation such as reduced precision, will give us x5; extensive active power and thermal control of the machine could net us as much as 50%; finally, efficient cooling, with possible provisions for energy recovery, could provide 40%; these factors are independent and thus could be multiplicative, netting x1000 improvement overall.

Although the experiments with TSUBAME-KFC will continue until the spring of 2016, it has already demonstrated world's top power efficiency. On the Green500 and the Green Graph 500 lists announced in November, 2013, TSUBAMEKFC became number one in the world in both, surpassing the second ranked machines on both lists by 24% respectively. This indicates that TSUBAME-KFC represents the best of the state-of-the-art in power efficiency in supercomping, and the results could even be carried over to IDCs for their improvements.

A. Discussion on Cooling

Ahead of overview of TSUBAME-KFC's cooling technology with warm liquid immersion, we discuss the cooling methodologies; While immersion cooling has been deployed in the past in machines such as the Cray-2, the Florinate coolant utilized was extremely expensive, and moreover evaporated at low temperature of 56 degrees Celsius, and in fact the vapor was collected to be re- condensed, requiring airtight packaging. In fact all the follow-on supercomputers except for ETA10 resorted to either low temperature air or water (nonsubmersive)cooling.

the MCS rack, and inside the rack there is a forced circulation of cooled air [9], and the server inside is air- cooled. For TSUBAME2.0, the inlet water temperature is approximately 7–8 degrees Celsius typical, while the outlet is 15–18 degrees. The inlet air to the server matches that, while the server outlet temperature is approximately +10T. Although the cooling solution is far more efficient than conventional air cooling in SC and IDC centers, due to the chilled water and rack/node fan requirements, with the observed PUE (power usage effectiveness) of 1.29 on the year average basis, in that we are losing more than 20% of energy towards cooling.

AVERAGE AMBIENT TEMPERATURE IN TOKYO

Month	Average Temperature (°C)	Average Wet-Bulb Temperature (°C)
Jan	6.1	2.1
Feb	6.1	2.1
Mar	9.4	5.0
Apr	14.6	10.3
May	18.9	14.9
Jun	22.1	18.5
Jul	25.8	22.4
Aug	27.4	23.0
Sep	23.8	21.1
Oct	18.5	14.9
Nov	13.3	9.2
Dec	8.7	5.2

One of the largest sources of power consumption of TSUBAME2 was identified to be cooling and semiconductor leakage power according to earlier prototype experiments and measurements on the production machine. In particular, average power of TSUBAME2 is slightly under a megawatt, but requires 8-9 degrees chilled water to be supplied by the chillers. Except for winter, ambient temperature of Tokyo is well above such (Table I), and thus involves operating power-hungry compressors. Also, when GPUs are in full operation,their temperature rises to nearly 80-90 degrees even with our chilled water enclosed air cooling. Through preliminary studies comparing air cooling to immersion cooling, we observed about 15% increase in power.

B.TSUBAME

In order to reduce overall total power usage, removing power-hungry compressors while keeping temperature of processors lower, we have decided to build a liquid immersion cooled, highly dense cluster with extensive power and thermal monitoring, with the requirements as depicted already in Section 1, called TSUBAME-KFC (Kepler Fluid Cooling). TSUBAME-KFC was designed and built in collaboration with NEC, NVIDIA, and most importantly, GRC that provided the liquid immersion cooling technology. Figure 2 is the overview of KFC, and Figure 3 the external view of the container and the cooling tower. TSUBAME-KFC submerges the servers in a warm liquid; although previous systems such as SuperMUC[14] have employed warm liquid cooling, since the latter is water-cooled, it required custom design of liquid pipes

TSUBAME-KFC: Ultragreen Supercomputer Testbud

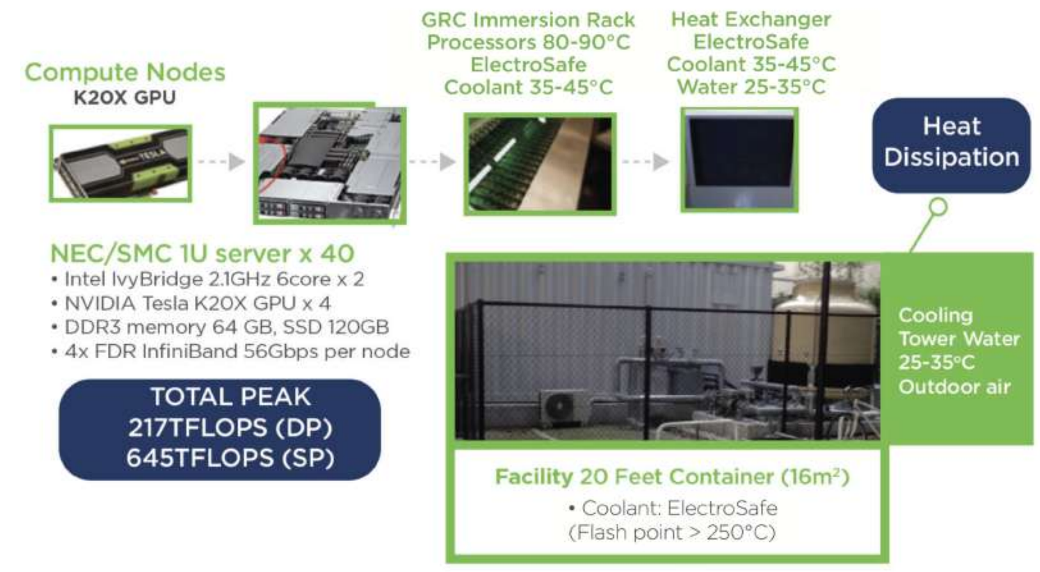

Fig. 2. TSUBAME-KFC cooling overview: the heat emitted from the server is transferred to GRC's ElectroSafe coolant, which then transfers the heat to a warm water loop, which in turn is cooled by an evaporative cooling tower.

With sufficient thermal capacity to cool the server. GRC's ICEraQ liquid immersion cooling (Figure 4) allows us to use standard servers with smaller degree of customization as is described later. Although the amount of coolant required is substantial, over 1000 liters compared to a few liters for standard pipe-based water cooling, it has the advantage of significant thermal capacity for easier control and resilience to power fluctuations. Immersion also has the advantage of effectively cooling all the components in the system, not just CPUs/GPUs or power supply, adding to increased cooling efficiency as well as long-term component stability.1 Figure 5 shows how all the nodes are completely submerged in the ElectroSafe coolant.

In order to cool the coolant itself, the there is a heat exchanger that transfers heat to the secondary water loop right next to the rack. The water in turn is cooled by ambient air / evaporative cooling tower right outside the TSUBAME-KFC 20-feet container. The cooling tower is a standard drip cooler where the water is slowly flowed to the bottom, cooling the water with ambient air through radiation and evaporation in the process. Although the cooling temperature cannot go below the dew point, because the maximum outlet target temperature of the TSUBAME-KFC rack is 35 degrees Celsius, as well as substantial capacity of both coolant and water in the cooling loop, allowing for averaging of the thermals throughout the machine without any excessive hotspot, preliminary analysis indicates that we do not require any chillers, even in the hottest summer Tokyo weather exceeding 35 degrees Celsius with high humidity. We believe the problem in the worst case scenario can be largely overcome with appropriate load control of the machine to not to reach maximum TDP. As such, the required cooling power consists of two pumps to circulate the coolant and water loops respectively, as well as a large fan internal to the cooling tower, plus replenishing of evaporated water.

C. Compute Nodes and Their Customization

TSUBAME-KFC employs 40 nodes of a customized version of standard highly-dense compute server NEC/SMC 104Re-1G (Supermicro OEM) that embodies 2 CPUs and 4 GPUs in a dense 1U form factor:

- Intel Xeon E5-2620 v2 (IvyBridge) 6 Cores2.1GHz ×2
- DDR3 Memory 64GB
- NVIDIA Tesla K20X GPU ×4
- SATA3 SSD 120GB (Expanded with 2 × 500GB SSD March 2014).
- 4x FDR InfiniBand HCA ×1
- CentOS Linux (x86 64) 6.4
- GCC 4.4.7, Intel Compiler 2013.1.039 •CUDA 5.5
- OpenMPI 1.7.2

The theoretical peak performance of each node is 15.8 Teraflops in single precision and 5.3 Teraflops in double precision floating point respectively. With 40 nodes comprising a single rack, the combined performance is approximately 212 Teraflops in double precision and reaches over 632 Teraflops in single precision, approaching

Fig. 3. Exterior view of TSUBAME-KFC: evaporative cooling tower right next to the 20-feet container for complete lights-out operation and low waterpump power

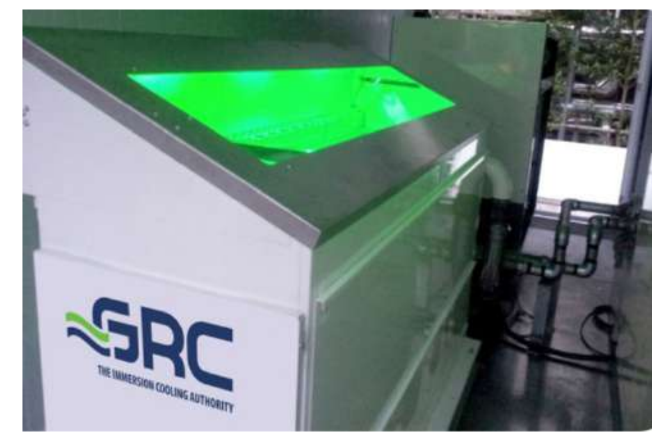

Fig. 4. The GRC ICEraQ system installed inside the container

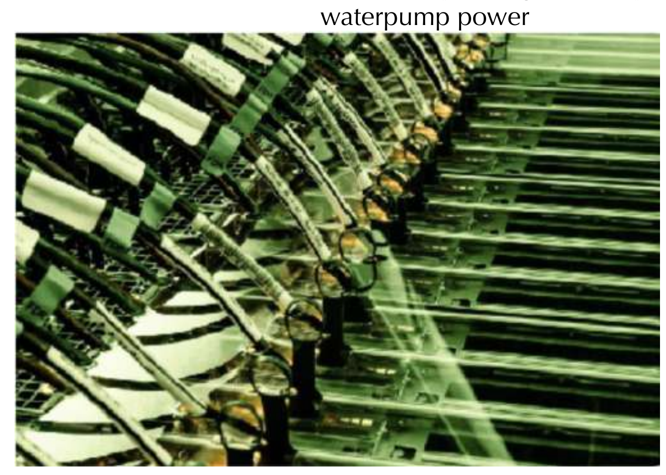

Fig. 5. The compute nodes submerged in the GRC ICEraQ system

the petaflop/rack metric for exascale. Although standard servers were used as a baseline, the following customization were performed on the nodes jointly with GRC and NEC:

- Removal of moving components — In order to submerse to high viscosity liquid, all moving parts such as server fans (12 units) were removed, as well as employing SSDs for storage. This has the additional benefit of lowering the node power requirements
- Thermal paste replacement — Since silicone grease between the liquid coolant processor and the passive cooler will dissolve in the liquid coolant, it was replaced with thin metallic sheets.

D. Coolant

The 40 servers are submersed in a 42U horizontal rack that is placed sideways instead of vertically. The theoretical peak power consumption of the nodes are approximately

40 kilowatts, far exceeding the standard 5˜15 kilowatts per rack in IDCs, and slightly greater than TSUBAME2.0's 35 kilowatts, although the ICEraQ system has the capability to handle up to 100 kilowatts per rack.2 In order to conform to regulations in Japan, the coolant had to be optimized carefully jointly with GRC and NEC. Regular mineral oil based coolants are non-toxic with low viscosity and fairly resistant to long-term oxidization, but nonetheless have the flash point of 177 degrees Celsius, and as a result considered as a flammable material under the Japanese fire laws subject to strict fire regulations equivalent to gasoline stands, with fairly stringent measures and licenses required for both the installation (such as comprehensive fire extinguishers and periodic inspections) as well as the operators (licenses large-scale supercomputer center operations. After extensive research, GRC's ElectroSafe+ coolant was proposed, with flashpoint of 260 degrees, well above the threshold of 250 degrees avoiding such complications.3 The selected coolant is also reasonable in the aspect of price.

E. Power and Thermal Measurements

TSUBAME-KFC embodies numerous power and thermal sensors, both on-line integral to the servers and other IT equipment, as well as independent off-line sensors. All the sensor streams are aggregated and can be observed in realtime as well as archived collectively. For example, the node power sensors are off-line, with individual PanasonicKW2G sensors and AKL1000 data logger, allowing server and switch measurement to be done in real time without any performance intrusion every second (Figure 6). Also infrastructural powers such as coolant and water pumps, as well as cooling tower are measured.

The list of sensors and their measurement intervals are described in Table II.

Although the sensors might seem too extensive for a production supercomputer, we believe most of the sensors will be incorporated into production TSUBAME3.0, to realize finegrain control of the machine for power minimization.The entire TSUBAME-KFC system was completed and began operation in October 2013, and will continue with various experimentations leading up to TSUBAME3.0, until Spring of 2016

III. POWER EFFICIENCY METRICS

For completeness we describe the power efficiency being used by the Green500 list, as well as the PUE (PowerUsage Effectiveness). Admittedly there have been numerous discussions regarding the metrics in the recent years, in that they only capture certain facets of power efficiency of a given machine, and should not be taken as ultimate decisive values.

Nonetheless for brevity we will not be controversial in this paper, but rather accept these as established metrics with well reported results that allows us to compare the power.

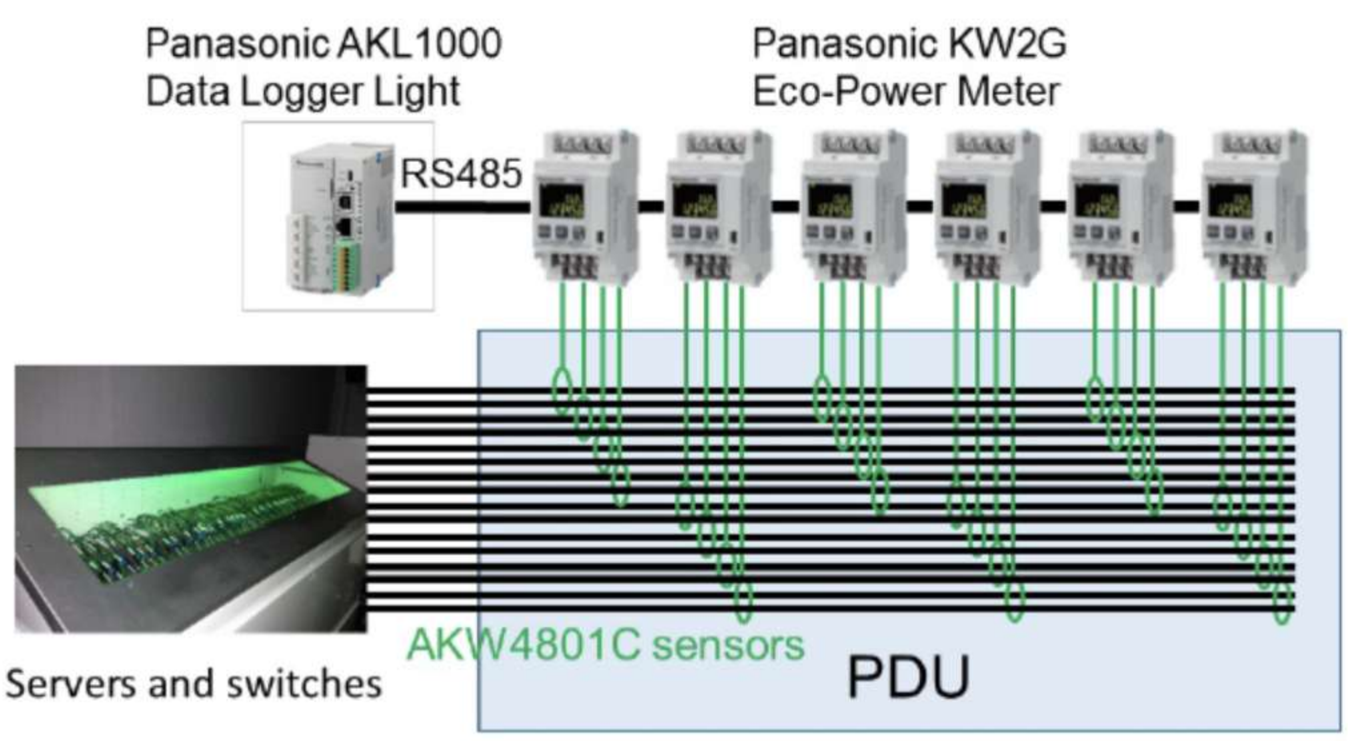

Fig. 6. Power measurement system of TSUBAME-KFC nodes

TABLE II.

Measured Component	Type	Provided By	Interval	Resolution
Computer node	Power	Panasonic data logger	1 sec.	0.1W
Network	Power	Panasonic data logger	1 sec.	0.1W
Cooling tower	Power	Panasonic data logger	1 sec.	0.1W
Pump (coolant)	Power	Panasonic data logger	1 sec.	0.1W
Pump (water)	Power	Panasonic data logger	1 sec.	0.1W
Outdoor air	Temperature	Panasonic data logger	10 sec.	0.1°C
Outdoor air	Humidity	Panasonic data logger	10 sec.	1%
Indoor air	Temperature	Panasonic data logger	10 sec.	0.1°C
Indoor air	Humidity	Panasonic data logger	10 sec.	1%
CPU, GPU	Temperature	IPMI on nodes (BMC)	<1 sec.	1°C
Pump (coolant)	Speed	GRC controller box	20 sec.	2%
Pump (coolant)	Power	GRC controller box	20 sec.	0.1W
Coolant	Temperature	GRC controller box	20 sec.	0.1°C
Water	Temperature	GRC controller box	20 sec.	0.1°C

efficiency with other systems at reasonable degrees of accuracy.

A. PUE

PUE (power usage effectiveness) is widely used in data centers as a metric to measure the power efficiency of cooling, and is given by:

$$PUE = \frac{(ITEquipmentPower + InfrastructurePower)}{ITEquipmentPower}$$

In this article, we equate cooling power to infrastructure power. PUE=1 indicates the ideal case of no power required for cooling, while PUE=2, which is a fairly common value in classical data centers, indicate that cooling is spending as much power as the IT equipment. As indicated,TSUBAME2.5 PUE is measured to be 1.29 on they early average, which is far superior to such values, but still is spending (1 − 1/1.29) = 23% power to cooling, and more if we account for the chassis fans. One caveat with PUE is that, low PUE may not necessarily indicate overall power efficiency. This has been a problem pointed out recently with the so called ambient air cooling in which natural outside air is used instead of cooled air by CRAC units. Due to the higher ambient temperature, the operational temperature of machines become considerably higher than traditional cooling methods, increasing the overall IT power with increased fan speeds and higher leakage current, both of which are accounted for as IT power, thus enlarging the denominator — the end result being lower PUE but higher overall datacenter power.

Thus it is important to measure PUE in terms of baseline power consumption under traditional cooling methods, not what is measured with new cooling solutions.

B. The Power Efficiency Metric of the Green500

Green500 basically measures the power efficiency of a machine during the Linpack run as dictated by the run rules of the Top500 Linpack benchmark[12], and the measured Linpack Flops/Watt is used as a metric to rank the machine, up to the 500th on the Top 500 list. A machine has to be listed on the Top 500 list to qualify to be ranked on the Green500, but a more power efficient run could be conducted and reported; otherwise the submitted Top 500 power result will be used.

The power measurement in the denominator of the Top500 is measured under the following conditions:

1) Overall power of the compute node and the network are measured and included.

2) Storage power are not included, unless it is integral to the node such as local disks.

3) Power for cooling are not included, unless it is integral to the node such as chassis fans.

4) For the Level 1 measurement, minimum 20% or 1 minute timespan of the duration of the entire Linpack run, whichever is longer, needs to be measured and averaged. The first 10% of the beginning and end the end of the Linpack

As specified in condition 3 above, one of the criticisms of the Green500 is the failure to include the overall cooling power requirements, i.e., advances in efficient cooling technologies do not have direct relevance to the Green500 measurement and ranking. Nonetheless, efficient cooling could have indirect effect such as elimination of chassis-internal cooling components, as well as lower thermal operations for components leading to lower power.

Also, condition 4 indicates that, for Level 1 (and Level 0) measurement only a subset duration needs to be measured. Due to the nature of Linpack, whose power consumption gradually drops in modern architectures to the nature of Linpack, whose power consumption gradually drops in modern architectures as the unsolved matrix becomes smaller and thus more communication intensive, the reported number is typically smaller than the overall average (Level 2 measurement) or during peak. For instance, our November 2013 submission of the TSUBAME2.5 supercomputer recorded 2.831 Petaflops, and the average power consumption for the entire run was 1125 kilowatts. On the other hand, the Level 1 measurement of the 70–90% duration of the run was 922.5 kilowatts, resulting in 3.069 GFlops/Watt submission measurement ranked 6th in the world.

Although there are some controversy as to whether the run rule constitutes a valid power measurement for Linpack, nonetheless as a ranking metric could be considered “fair” if they are measured the same way across all the machines. However, due to the one minute rule above, small machines whose power degradation timespan is much shorter than one minute is disadvantaged in principle.

TABLE III.

COMPARISON OF TSUBAME-KFC SUBMERSED NODE AND AIR-COOLED NODE

Cooling	Air-Cooled (26°C)	Immersion Coolant (29°C)	Immersion Coolant (19°C)
Temp (°C)			
CPU1	46	42	33
CPU2	50	40	31
GPU1	52	47	42
GPU2	59	46	43
GPU3	57	40	33
GPU4	48	49	42
Node Power (W)	749	693	691

IV. TSUBAME-KFC EVALUATION

A. Effect of Liquid Immersion Cooling on the Servers

We first measured how liquid immersion cooling affects the power and thermals of individual, densely configured TSUBAME-KFC node. We configured an air-cooled node with exactly the same CPU/GPU/memory hardware for comparative purpose. As mentioned earlier, the air-cooled node has the original 12 fans, while the immersed node have the fans eliminated.

Table III shows the comparison between air-cooled with inlet 26 degrees Celsius, versus 29 and 19 degrees inlet for the ElectroSafe coolant. The servers are continuously running double precision matrix multiply using CUBLAS to incur the highest power and thermal load.

Comparing the air-cooled versus immersion, although the former has lower temperature input, the latter exhibits substantially lower temperature, especially GPU2 and GPU3 where difference is more than 10 degrees, or ΔT of 33 degrees for air compared to 20 degrees for ElectroSafe. This is due to much higher thermal capacity of ElectroSafe, especially since these GPUs are inline to the airflow path of GPU1 and GPU4, being affected by the already warmed air. The result for 19 degree coolant inlet is even more significant, ΔT being fairly consistent 24 degrees.

Comparing the server power consumption, liquid immersion is approximately 7.8% lower than air, while inlet temperature difference has very little effect on the overall power. As mentioned earlier, this is largely the combined effect of fan removal and lower semiconductor temperature suppressing leakage current. That there is small difference between the two temperature points of liquid immersion could indicate that the former is more dominant, but the prior experiments have indicated that the latter effect is also significant. This could be possibly explained by the exponential effect of leakage current versus temperature being exponential in nature, and thus that even at 29 degrees the component temperature was too low to exhibit the difference not hitting the rising "knee" of the curve, while at higher temperature it would quickly rise. We plan to conduct more thorough experiments during summer months to investigate the temperature point at which the component temperature will start to exhibit noticeable increase in server power consumption, a valuable data for power control in that we would want to control the coolant temperature just under this point to minimize cooling power.

B. Power and PUE Measurement

In order to measure TSUBAME-KFC power consumption and PUE, we stressed the server with the highest load of CUBLAS matrix multiply as in the previous subsection for all the nodes. Figure 7 shows the results. The PUE number for TSUBAME is derived from the actual measurements from the real-time power sensors, while the air cooling was extrapolated from real power measurement of the server, with the assumption that the state-of-the-art air cooling would be as efficient as TSUBAME2.0's PUE of 1.29.

According to the measurements, power consumption of the TSUBAME-KFC IT equipment was 28.9 kilowatts while the pump and cooling tower power combined were2.60 kilowatts for apparent PUE of 1.09. However, as mentioned earlier, the PUE comparison here is misleading, as the node server power consumption itself has gone down significantly. The total power usage of 28.5 kilowatts is22% smaller than 40.7 kilowatts in the air-cooled case. In fact, the total power usage of 31.5 kilowatts is essentially equivalent to the IT-only power usage of air-cooled machine, being 31.2 kilowatts. As such TSUBAME-KFC efficiency cannot merely be judged by comparing PUE values alone.

C. Green500 Measurement

In order to measure the efficiency of TSUBAME-KFC under more realistic setting, we challenge the Green 500for utmost efficiency. Although both benchmarks do not directly measure contributions from cooling, nonetheless we expect higher efficiency through improved node efficiency as described above, as well as improved tuning for power efficiency rather than absolute performance. As an initial comparative measure, we targeted the previous#1 system for the June 2013 edition of the Green 500, namely the 3.209 Gigaflops/Watt record of the CINECA/ Eurotech Aurora machine.

In order to achieve the maximum power efficiency, we employed the following strategies at the software and hardware control levels:

• At the architecture level we increased the ratio of GPU to CPU ratio from 1:1 to 2:1, thus decreasing the overhead of CPU and other peripheral power consumption(The GPU was also slightly slower, being NVIDIA K20 instead of K20X for KFC, but the overall effect on the performance is believed to be low for theGreen500.)

• We employed a new, more efficient in-core Linpack kernel provided by NVIDIA for both TSUBAMEKFC and TSUBAME2.5 measurements. This version only computes Linpack using the memory space of GPUs, not of the whole node, for efficiency and better power- performance. However, this also results in a much shorter runtime, and for a small machine such as TSUBAME-KFC, this hinders Level-1 measurement due to the one minute rule as described earlier.

• We tuned the HPL parameters to the maximum extent by exhaustive search of the parameter space. This involved not only the standard tuning of HPL parameters such as the block size (NB), and process grid (P&Q), but also

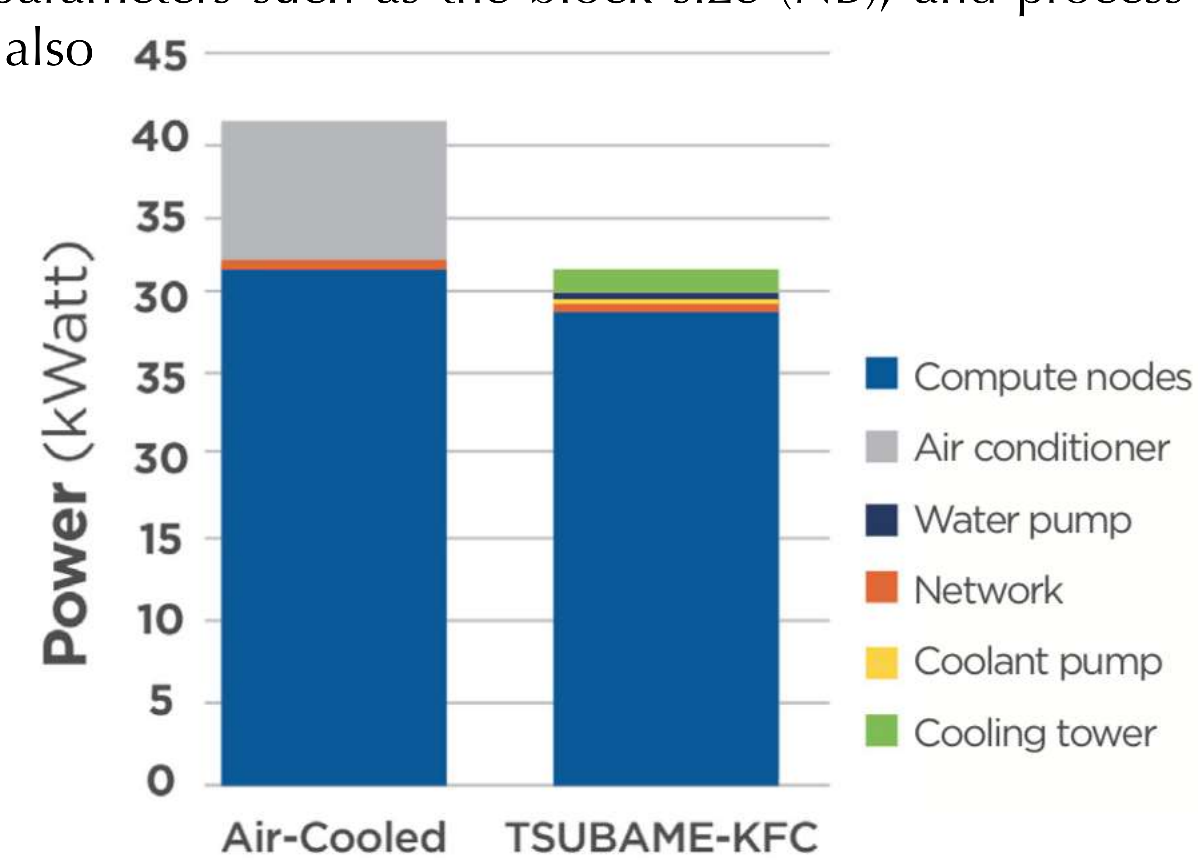

Fig. 7. PUE Evaluation of TSUBAME-KFC versus Air-Cooled Machine

Adjustment of the GPU clock and voltage, where the slowest clock of 614 MHz proved to be the most power efficient, compared to the default of 732 MHz (Available GPU clocks rates were 614(best), 640, 666, 705, 732 (default), 758, 784 (MHz)).

• We conducted measurements at the most thermally low and stable night hours.

During the tuning phase, we noticed that the best performance does not equal best power efficiency. Figure 8 shows power efficiency of Linpack runs with various configurations; we observed the best power efficiency is 24% better than the case with the best speed performance.

As a result, Flops/Watt of reached 4.503 GFlops/Watt,improving the CINECA record by over 40% (Table IV). On November 18th, 2013, the Green500 list was announced, in which TSUBAME-KFC was ranked #1 in the world, with 24% lead over Wilkes, the second ranked machine. In fact, the top three machines were similarly configured with 2:1GPU ratio of NVIDIA K20X GPUs versus Intel Ivy Bridge Xeon CPUs and Infiniband interconnect. We attribute the difference to better cooling as well as more extensive tuning of the parameters. In the latest list announced in June 2014, TSUBAME-KFC was ranked #1 again; the efficiency number is slightly changed to 4.390GFlops/Watt,since we had to choose another Linpack run for submission in order to keep the system in the latest Top500 list. Also TSUBAME-KFC obtained #1 in another ranking,namely Green Graph500 list [13] in November 2013, designed for competition for power efficiency in big-data analysis area, although we omit details for want of space.

D. Evaluation with the Phase-Field Simulation

We also evaluated power efficiency of a real application, namely the stencil based “phase-field” simulation, which was awarded the 2011 Gordon Bell prize[15] by achieving 2 Petaflops on the TSUBAME2.0. This application simulates the micro-scale dendritic growth of metal materials during solidification phase. We have used it for evaluation of power efficiency for multiple years in our projects and have comprehensive power performance records executing on TSUBAME1.0 machine commissioned in 2006.

When we run the application on a single TSUBAMEKFC

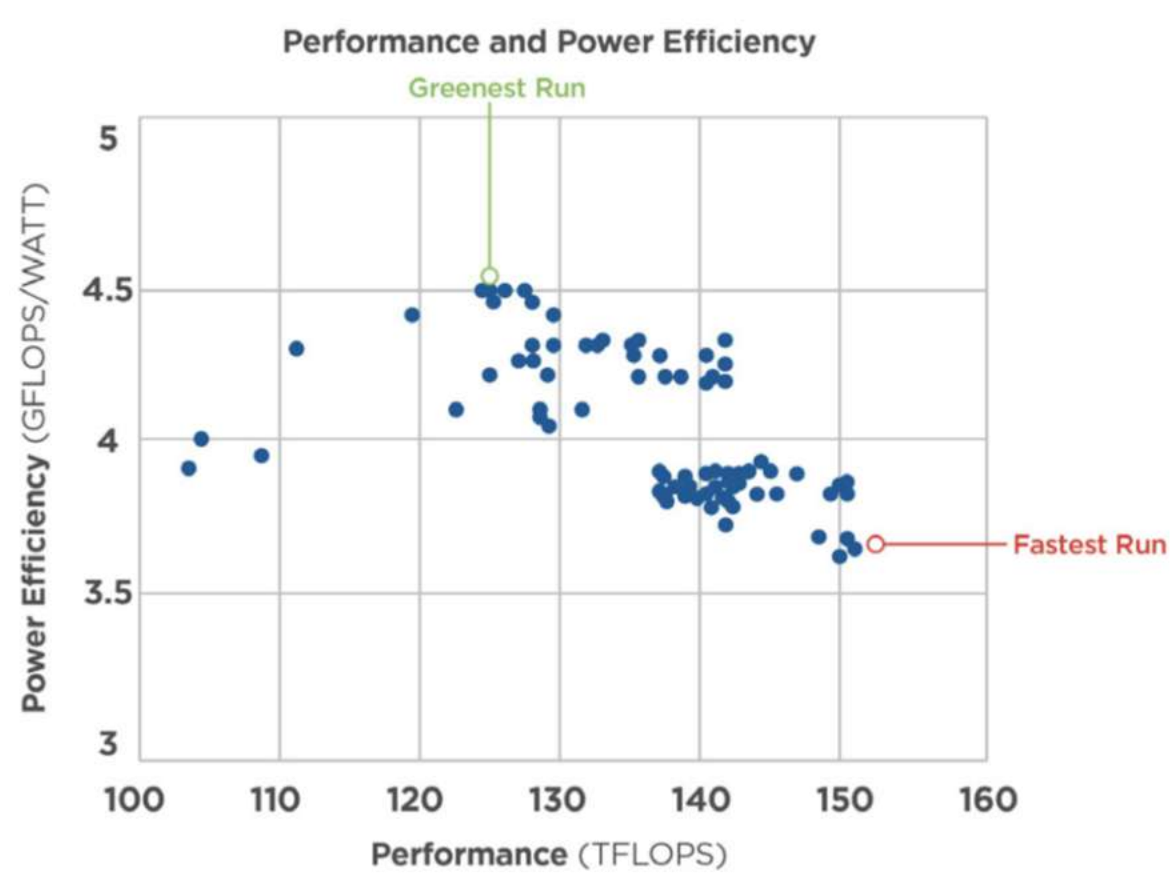

Fig. 8. Results of Linpack benchmark runs with various configurations. Each dot corresponds to a single Linpack run.

TABLE IV.

POWER EFFICIENCY IN THE GREEN500 METRICS

System	Time	Speed (TFLOPS)	Power (KW)	Power Efficiency (GFlops/W)
TSUBAME2.0	NOV. 2010	1192	1244	0.958
CINECA	Jun. 2013	98.51	30.70	3.209
TSUBAME2.5	Nov. 2013	2831	922.5	3.069
Wilkes	Nov. 2013	191.1	52.62	3.632
TSUBAME-KFC	Nov. 2013	125.1	27.78	**4.503**

node with four GPUs, we observed 3.62 TFlops (single precision)with the power consumption of 652 Watt; as such the power efficiency is 5.55 GFlops/Watt. Table V compares the result with that on an air-cooled machine with the same configuration, demonstrating that the immersion cooled machines provides 8.5% higher efficiency, consistent with our Green500 measurement.Figure 9 also shows the development of power efficiency by comparing several machines in different generations since 2006. Here we observe a gap between the 2006 (CPU only) and 2008 (with GPUs) numbers, which is the one-time performance leap with the many-cores transition.By extrapolating the results of multiple generations of GPU machines, we estimate that the expected performance circa 2016 being as 15.5 GFlops/Watt, which is 1,200 times more efficient than the 2006 number. This result supports our initial assessment of our ability to achieve the target improvement depicted in Figure 1, i.e., x1,000 in 10 years.

V. CONCLUSION

We demonstrated a prototype supercomputer that combines most of the known the state-of-the-art architecture in both hardware and software, materialized as TSUBAME-KFC, achieving the world’s top power efficiency in Green500 rankings circa November 2013 and June 2014. TSUBAMEKFC improved the previousGreen500 #1 record by over 40%, and was 24% better than similarly configured machine ranked #2. The most notable technology deployed by TSUBAMEKFC is the liquid immersion cooling; the total

TABLE IV.

POWER EFFICIENCY OF THE PHASE-FIELD SIMULATION ON TSUBAME-KFC

System	Speed (SP TFLOPS)	Power (KW)	Power Efficiency (GFlops/W)
A KFC node	3.62	0.652	5.55
An air-cooled node	3.58	0.701	5.11
40 KFC nodes	126	25.5	4.93

Fig. 9. Power efficiency of the phase-field simulation on machines in different generations

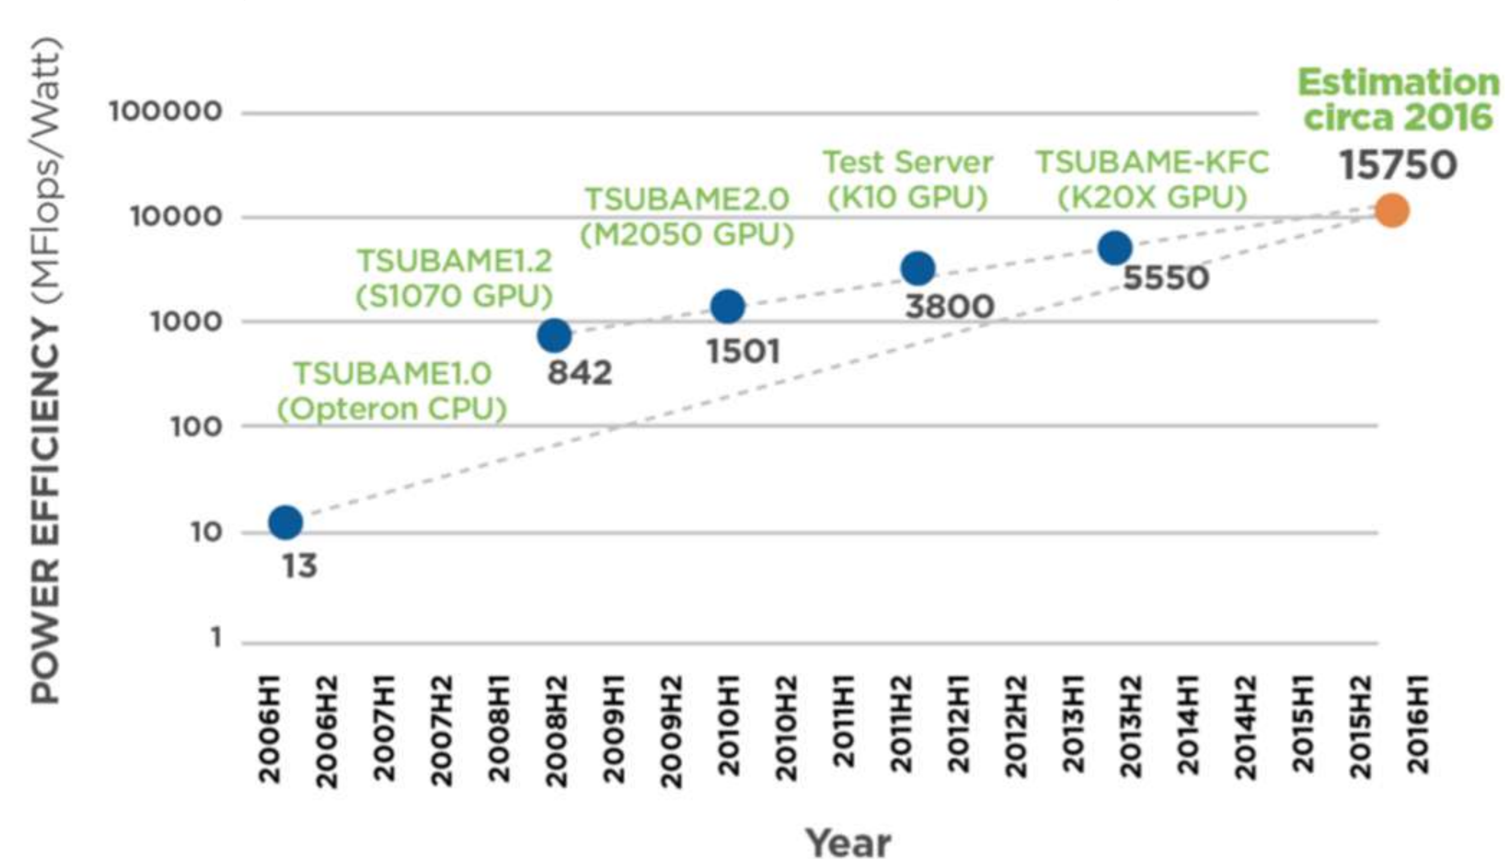

power consumption is reduced by 29% over compute nodes with the same configurations. These power saving features of TSUBAMEKFC will be incorporated into TSUBAME3.0, as is or with some pragmatic adaptations.

Not only as a prototype of TSUBAME3.0, TSUBAMEKFC also intends to be a platform for reproducible experiments regarding power saving. It is well known that the rise in semiconductor temperature results in substantial increase in power requirements, due to the increase in leakage current. As such, it is very difficult to conduct a reproducible experiment, maintaining constant thermal conditions. By immersion cooling with liquids of massive thermal capacity, we can control the thermals more easily; as is described later, TSUBAME-KFC is immersed in 1200 liters of liquid, allowing high thermal stability thus reproducibility all year round.We will continue our experimentations of TSUBAMEKFC, some of the most up-to-date-result only obtainable during summer in the camera-ready version of the paper, as well as longer term data such as long-term component faults. As an experimental platform, we will conduct further customization updates either in software or hardware if affordable, to affect the design of TSUBAME3.0 as the bleeding-edge power efficient and big data supercomputer of the era. We also hope to open up TSUBAME-KFC for uses by our collaborators so as to obtain reproducible results in power efficient computing on the state-of-the-art architecture.

ACKNOWLEDGMENTS

This research is funded by Japanese Ministry of Education,Culture, Sports, Science and Technology (MEXT) Ultra Green Supercomputing and Cloud Infrastructure Technology Advancement. We also deeply thank the extensive collaborations with NEC, NVIDIA, GRC, Supermicro, Mellanox, as well as GSIC/Tokyo Institute of Technology.

REFERENCES

[1] W. Feng: Making a Case for Efficient Supercomputing. ACM Queue, 1(7):54-64, October 2003.
[2] IBM Journal of Research and Development, special double issue on Blue Gene, Vol.49, No.2/3, March/-May, 2005.
[3] Hiroshi Nakashima, Hiroshi Nakamura, Mitsuhisa Sato, Taisuke Boku, Satoshi Matsuoka, et. al. (2 more authors), "MegaProto: 1 TFlops/10 kW Rack Is Feasible Even with Only Commodity Technology", Proc. IEEE/ACM Supercomputing 2005, the IEEE Computer Society Press, Nov. 2005.
[4] International Exascale Software Project, http://www.exascale.org/ iesp
[5] Japan Science and Technology Agency, CREST, Development of System Software Technologies for Post-Peta Scale High Performance Computing. http://www.postpeta.jst.go.jp/en
[6] K. Datta, M. Murphy, V. Volkov, S. Williams, J. Carter, L. Oliker, D. Patterson, J. Shalf, and K. Yelick, "Stencil Computation Optimization and Auto-tuning on State-of-the-Art Multicore Architectures", IEEE Supercomputing (SC08), Article No. 4, pp. 1-12, 2008.
[7] Song Huang, Shucai Xiao, and Wu-chun Feng, "On the Energy Efficiency of Graphics Processing Units for Scientific Computing", IEEE International Symposium on Parallel & Distributed Processing (IPDPS 2009), pp. 1-8, 2009.
[8] Shiqiao Du, Takuro Udagawa, Toshio Endo and Masakazu Sekijima, "Molecular Dynamics Simulation of a Biomolecule with High Speed, Low Power and Accuracy Using GPU-Accelerated TSUBAME2.0 Supercomputer", Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC 2011), pp. 1-5, 2011.
[9] Satoshi Matsuoka. "TSUBAME 2.0 Begins – The long road from TSUBAME1.0 to 2.0 (Part One) –", The TSUBAME E-Science Journal, Tokyo Tech. GSIC, Vol. 2, pp. 2-10, Nov. 2010
[10] Top500 Supercomputer Sites, http://www.top500.org
[11] The Green500 List, http://www.green500.org
[12] A. Petitet, R. C. Whaley, J. Dongarra, A. Cleary, "HPL - A Portable Implementation of the High-Performance Linpack Benchmark for Distributed-Memory Computers", http://www.netlib.org/benchmark/ hpl/
[13] The Green Graph500 List, http://green.graph500.org
[14] Leibniz Supercomputing Centre, "SuperMUC Petascale System". http://www.lrz.de/services/compute/supermuc
[15] T. Shimokawabe, T. Aoki, T. Takaki, A. Yamanaka, A. Nukada, T. Endo, N. Maruyama, S. Matsuoka. "Peta-scale Phase-Field Simulation for Dendritic Solidification on the TSUBAME 2.0 Supercomputer", IEEE/ACM Supercomputing (SC11), pp. 1–11, Seattle, November 2011.
Note: The K20X GPUs were changed to K80s in 2015. The name of the system is now TSUBAME-KFC/DL.

Vienna Scientific Cluster's The Immersion Supercomputer: Extreme Efficiency, Needs No Water

VIENNA SCIENTIFIC CLUSTER

Abstract

The Vienna Scientific Cluster (VSC), created to satisfy the demand for High Performance Computing (HPC) of a consortium of five Austrian universities, released a competitive tender to build the most pow- erful supercomputer in Austria. However, this was not your regular competitive bid process where sys- tems that meet the minimum requirement compete purely on the bid price. The VSC took a long-term perspective towards cost and performance, they defined very specific targets and tests that reflected the long-term cost of the system. GRC came together with ClusterVision, Intel, and Supermicro to create a custom solution that came out well on top of the competition. This paper will explore how GRC's technology helped dramatically reduce the upfront cost and energy consumption of the system. Thereby allowing reallocation of capital towards more computing power.

Introduction

When the five Austrian universities came together to create the VSC-3 their aim was to build the most powerful mainframe computer in Austria, while ensuring minimal cost of ownership. Getting the most computing power for a limited budget is something every data center planner strives for. However,
the VSC went a step further and changed the very way in which systems are procured. Traditionally, organizations define a minimum level of computing power, they decide on a ballpark budget, and then invite bids from vendors to meet those requirements at the lowest possible cost.
VSC on the other hand created a procurement process that would reflect the total cost of ownership for the system over a longer period of time. They defined very specific metrics and tests for the systems and then compared their performance scores along with their upfront and operating costs. By measuring what they truly wanted, the VSC was able to get a system that was custom tailored to their needs. Green Revolution Cooling (GRC), in partnership with the leading European commercial HPC cluster specialist, ClusterVision, as well as Intel and Supermicro, created a solution that beat every other competitor, including big brand OEMs and other precision cooling solution providers.

Selection Criteria

VSC defined a point system on which each competing system would be tested. All entries were put through an acceptance test to ensure that they meet the minimum requirements for computing power, energy efficiency, network performance, and storage performance. The acceptance testing included a stability test followed by an endurance test.
Once accepted, each of the systems were tested on:

- ComputingPower
- EnergyEfficiency
- Price

Description	Value	Units	Source
Upfront Cost	€6,000,000	Euro	Estimated
Computing Power	500kW	Kilo Watt	Estimated
Total PUE	1.15	-	Predicted Target
Cost of Power	€0.13	€/kW-hr	Average Cost of Industrial Electricity in Austria[1]

*Assumptions - Not reflective of final installation, these are GRC's estimates for a typical installation.

Let's take a deeper look into how such a criteria reflects total cost of ownership of a system:

Results

Calculating the total cost of the system based on the above assumptions, over a five year period shows that the cost of energy for even an extremely efficient system would be about 35% of the total cost of ownership.

Costs Over 5 Years

Energy	€3,274,050.00	(35% of Total Cost)
Ownership	€9,274,050.00	

Hence, including energy efficiency as a procurement criteria and giving it the appropriate weightage helped the VSC get a system that not just had a lower upfront cost but lower operating cost as well.
Further, the comprehensive testing on each of the metrics ensured that they confirmed the true cost of the system. For example, the test for energy efficiency measured the power drawn by the complete system, including that of the IT hardware. Additionally, the VSC estimated the power requirements of supplying any cold water that the cooling system may use. This power requirement was calculated based on the water temperature requirements of the cooling system, and is reflective of the cost of procuring and operating an appropriately sized chiller plant.
The ICEraQ (formerly known as CarnoJet) system's ability to operate with water as warm as 50°C gave it a considerable edge over most of the competition. This meant that the system would not require a chiller even in the peak of summer where water temperatures approach 43°C. Thereby, completely eliminating the capital and operating expenses related with a chiller plant. Let's take a deeper look into some of the other factors that helped GRC come out on top.

GRC's Winning Solution Solution

The winning solution, pictured above, was installed in its Vienna home in July 2014. The Cluster is made up of 2,020 nodes, each with 16 processor cores housed in the ICEraQ system. The more than 600 teraflops of computing power takes up a mere 540 kilowatts of power and is packed into a little over 1000 square feet of white space. Making it one of the most efficient non-GPU based supercomputers in the world, while ensuring extremely efficient white space utilization.

Area	1,104	Sq. ft.
Total Power	540	kW
Number of Nodes	2,020	16 processors /each
Computing Power	>600	TFLOPS
Efficiency	0.8^2	kW/TFLOP
Power Density	490	W/Sq. ft.

Another key highlight of the deployment was how GRC's technology enabled server cost reductions through hardware optimization. This reduction in server costs coupled with infrastructure savings helped the VSC get more compute from their budget.

To put things into perspective, 600 TFLOPS of power means that the cluster can perform more than 600 trillion additions or multiplications in just one second. For example, in one nanosecond, VSC-3 could take the speeds and ranges of every pass and shot in the entire FIFA World Cup 2014 tournament and calculate the corresponding launch angles.

Key Benefits of GRC's Solution

Here are some of the key benefits of GRC's technology that gave it the competitive edge:

Lower Energy Requirements

Lower Cooling Energy — The ElectroSafeTM coolant used in the ICEraQ system enables extremely efficient cooling. The coolant is a dielectric mineral oil blend, does not conduct electricity but is a good conductor of heat, making it ideal for cooling of IT equipment. This superior heat conductivity, drastically improves the system's ability to both extract heat from the servers and subsequently to expel that heat out of the system.
It is this property of ElectroSafe that allows the ICEraQ system to maintain core temperatures well below conventional cooling methods even with water as warm as 50°C. The ICEraQ system boasts a PUE of 1.02-1.033, as was reported by Intel based on tests carried out for over one year.

Lower Server Energy

Besides reducing the energy required for cooling, the ICEraQ system also reduces the energy consumed by servers themselves. The system completely eliminates the need for many components such as server fans, as is explained in greater detail further in this paper, this reduction of components along with superior thermal management helps reduce the energy drawn by the IT load itself.

The ElectroSafe coolant helps eliminate hot spots, and reduces current leakages from components. These savings added up to a 10-20% reduction in energy drawn by the cluster itself.

More for Less

The price of a system is usually directly proportional to its computing power. So an increase in computing power would result in a higher priced system. However, GRC's technology enabled dramatic reductions in equipment and infrastructure costs, allowing a bigger portion of the budget to go towards computing power. Here's where these savings came from:

Simpler Infrastructure — Apart from eliminating the need for any kind of ambient cooling such as chillers, and air handlers, the lower power rating of the data center helps downsize the complete infrastructure, including backup generators, UPS batteries, and power distribution.

Optimized Hardware — The use of the ICEraQ system enables the optimization of hardware for cost and energy reductions. There are a few obvious ways wherein the hardware can be optimized to cut down the costs of the servers. Most OEMs like Supermicro®, in this case, give customers the option to customize hardware to their specific needs. Customizing hardware design and configuration can help further trim procurement and operating costs of equipment for use in the ICEraQ system.

Standard OEM systems are designed for a variety of applications and usually come with an abundance of expansion ports and/or memory channels on the motherboard. These are usually populated with additional and unnecessary components that use more power and therefore require more cooling as well.

GRC's technology enables the removal of all these unnecessary components and channels on the mother board. This mixing and matching of hardware can help reduce both, the cost and power consumption of the system. The lower power rating also helps downsize the power supply, further reducing the cost of the hardware. An optimized power supply can save over $50 per power supply, while optimizing the motherboard itself can save over $100 per system. Additional savings are also available by picking memory without heat spreaders. Overall, choosing the right combination, and size of optimized hardware can result in savings close to $300 per unit, adding up to a total of up to $12,000 per rack.

Besides the cost savings, the ICEraQ system allowed system designers to pick the optimal hardware for the application, without concerns regarding their form factor or thermal restraints.

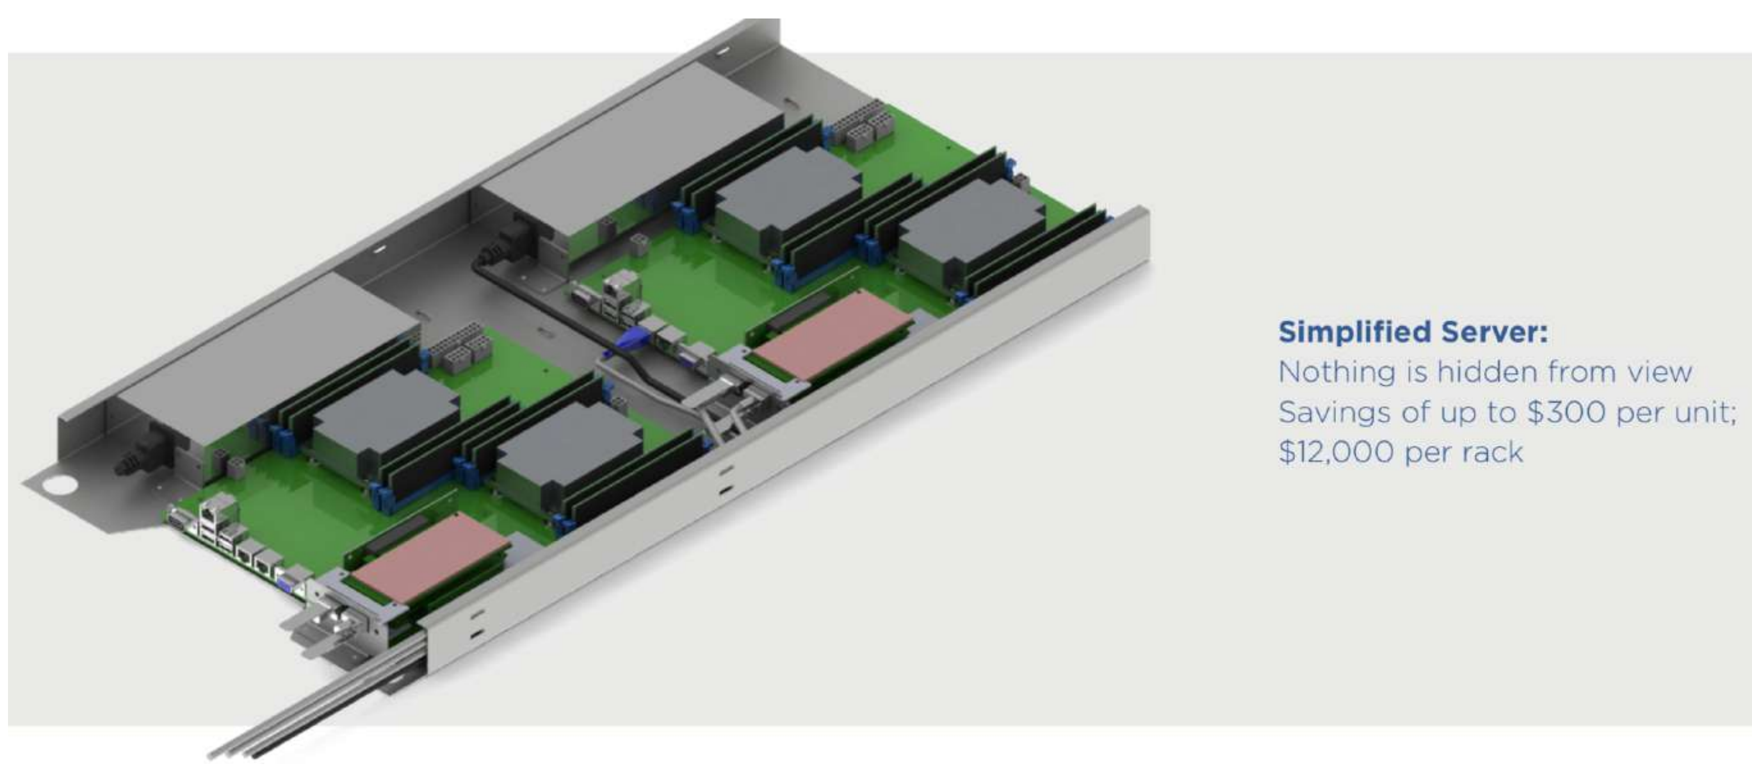

Zero Water Consumption

Another key feature, of the installation was its ability to reduce water consumption to zero. The ICEraQ's ability to accept warm water, combined with the climate in Vienna allowed for the use of a dry-cooler instead of a cooling tower, thereby delivering water-free cooling around the year.

Conclusion

Performance and efficiency have always been seen as a trade-off for upfront costs. However, GRC's technology is changing that view, many customers like the VSC are realizing the true benefits of the ICEraQ system. Beyond the lower build and operating costs, the ICEraQ system is giving data center planners more flexibility to get more of what they want by making the whole system leaner, right from the data center architecture down to the servers, everything is optimized for performance and efficiency. GRC encourages innovation in hardware, and the ICEraQ system enables just that. For more information on how GRC can help design custom solutions for your next build

References

1 Austria, S. (2014). STATISTIK AUSTRIA - Energy Prices. [online]Statistik.at. Available at: http://www. statistik.at/web_en/statistics/Prices/energy_prices/074082.html [Accessed 31 Oct. 2014].

2 Tuwien.ac.at, (2014). Technische Universität Wien : Österreichsneuer Supercomputer ist grün. [online] Available at: http://www.tuwien.ac.at/en/news/news_detail/article/8880/ [Accessed 31 Oct.2014].

3 Data Center Knowledge, (2012). Intel Embraces Submerging Servers in Oil | Data Center Knowledge. [online] Available at: http://www.datacenterknowledge.com/archives/2012/09/04/intel-explores- mineral-oil-cooling/ [Accessed 31Oct. 2014].

eVolution networks

About eVolution

Market Leader In Data Networks Energy Efficiency

Mobile & Data Centers

Software Only Solution

Sprint
Telefónica
Tier-1 Customers

Berkshire Hathaway Energy
GE
Blue Chip Investors

WE Belive In

ENERGY
One of the most fundamental and Crucial elements of life

EFFICIENCY
The way service providers operate contrasts all energy efficiency pricpiles

SOLUTION
Managed by using an adaptive,enegry-minded approach

AI
Artificial Intelligence drives the evolution of energy consumption forward

Market Trends

5% of global energy consumption originates from data centers

- IT load energy constitutes 80% of total energy consumption for data centers
- By 2020, data centers will consume 140 billion kilowatt- hours annually

Between 2015 and 2020, the NFV and open code market is expected to grow at a compounded annual growth rate of 42%

The Artificial Intelligence market is estimated to grow at a CAGR of 53.65% from 2015 to 2020

Smart Energy Solution

- Reduces up to 50% of DC Energy Bill (Depends on DC Load)
- Dynamically live-migrates nodes to available physical servers in order to minimize energy consumption of customer's data center facility
- Based on eVolution's AI and Deep Learning Prediction (DLP)© technology
- Supports multiple virtualization platforms: VMWare, OpenStack, Oracle

Potential Savings

- Data centers are provisioned for peak-demand while different service workloads fluctuate throughout the day
- Most data centers are optimized for handling the peak workload, rarely are they optimized for off peak workloads.

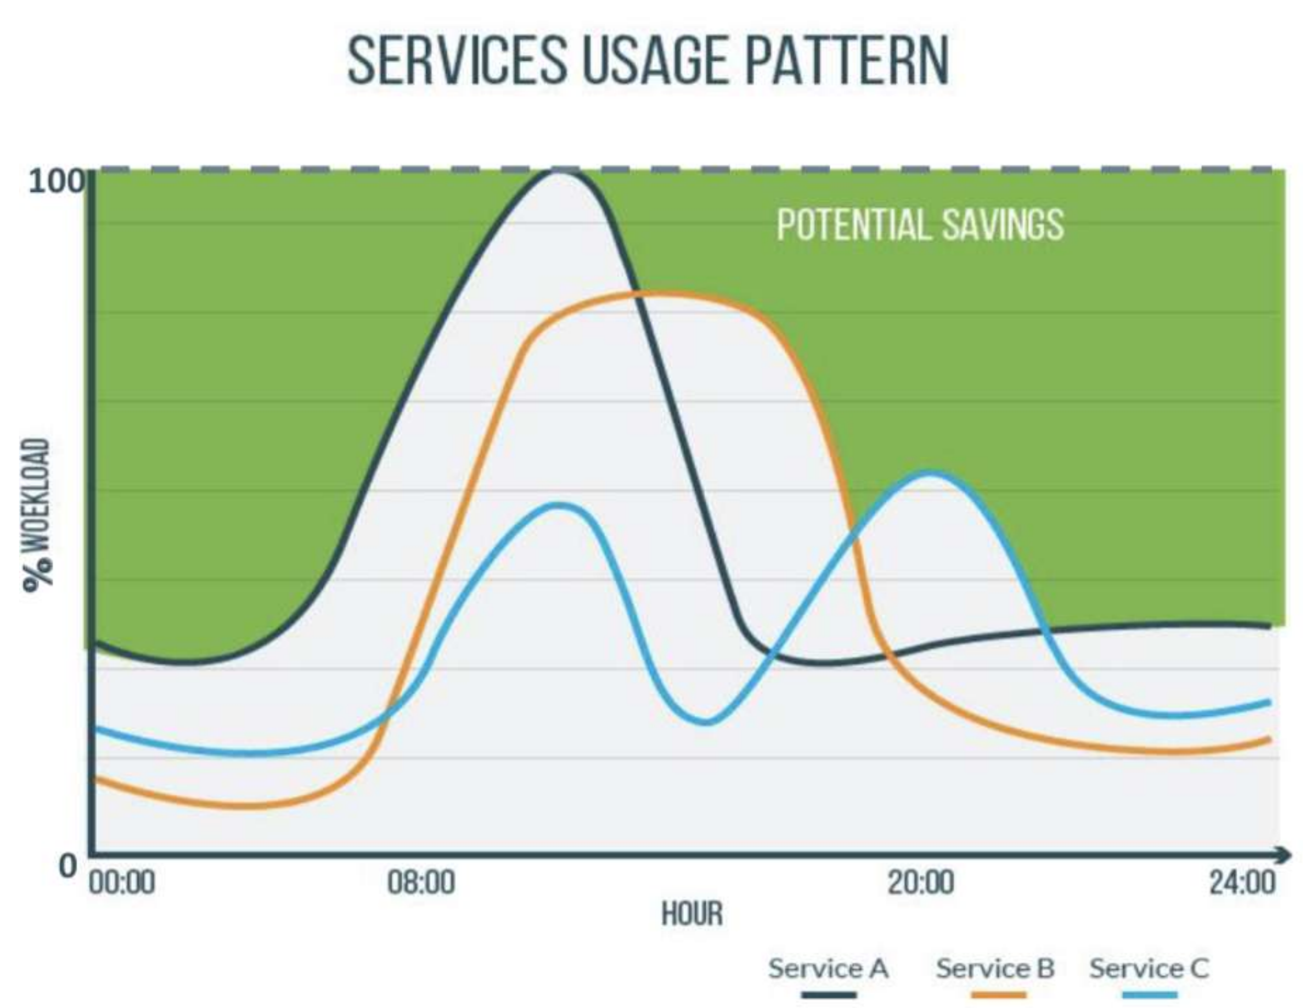

How it Works?

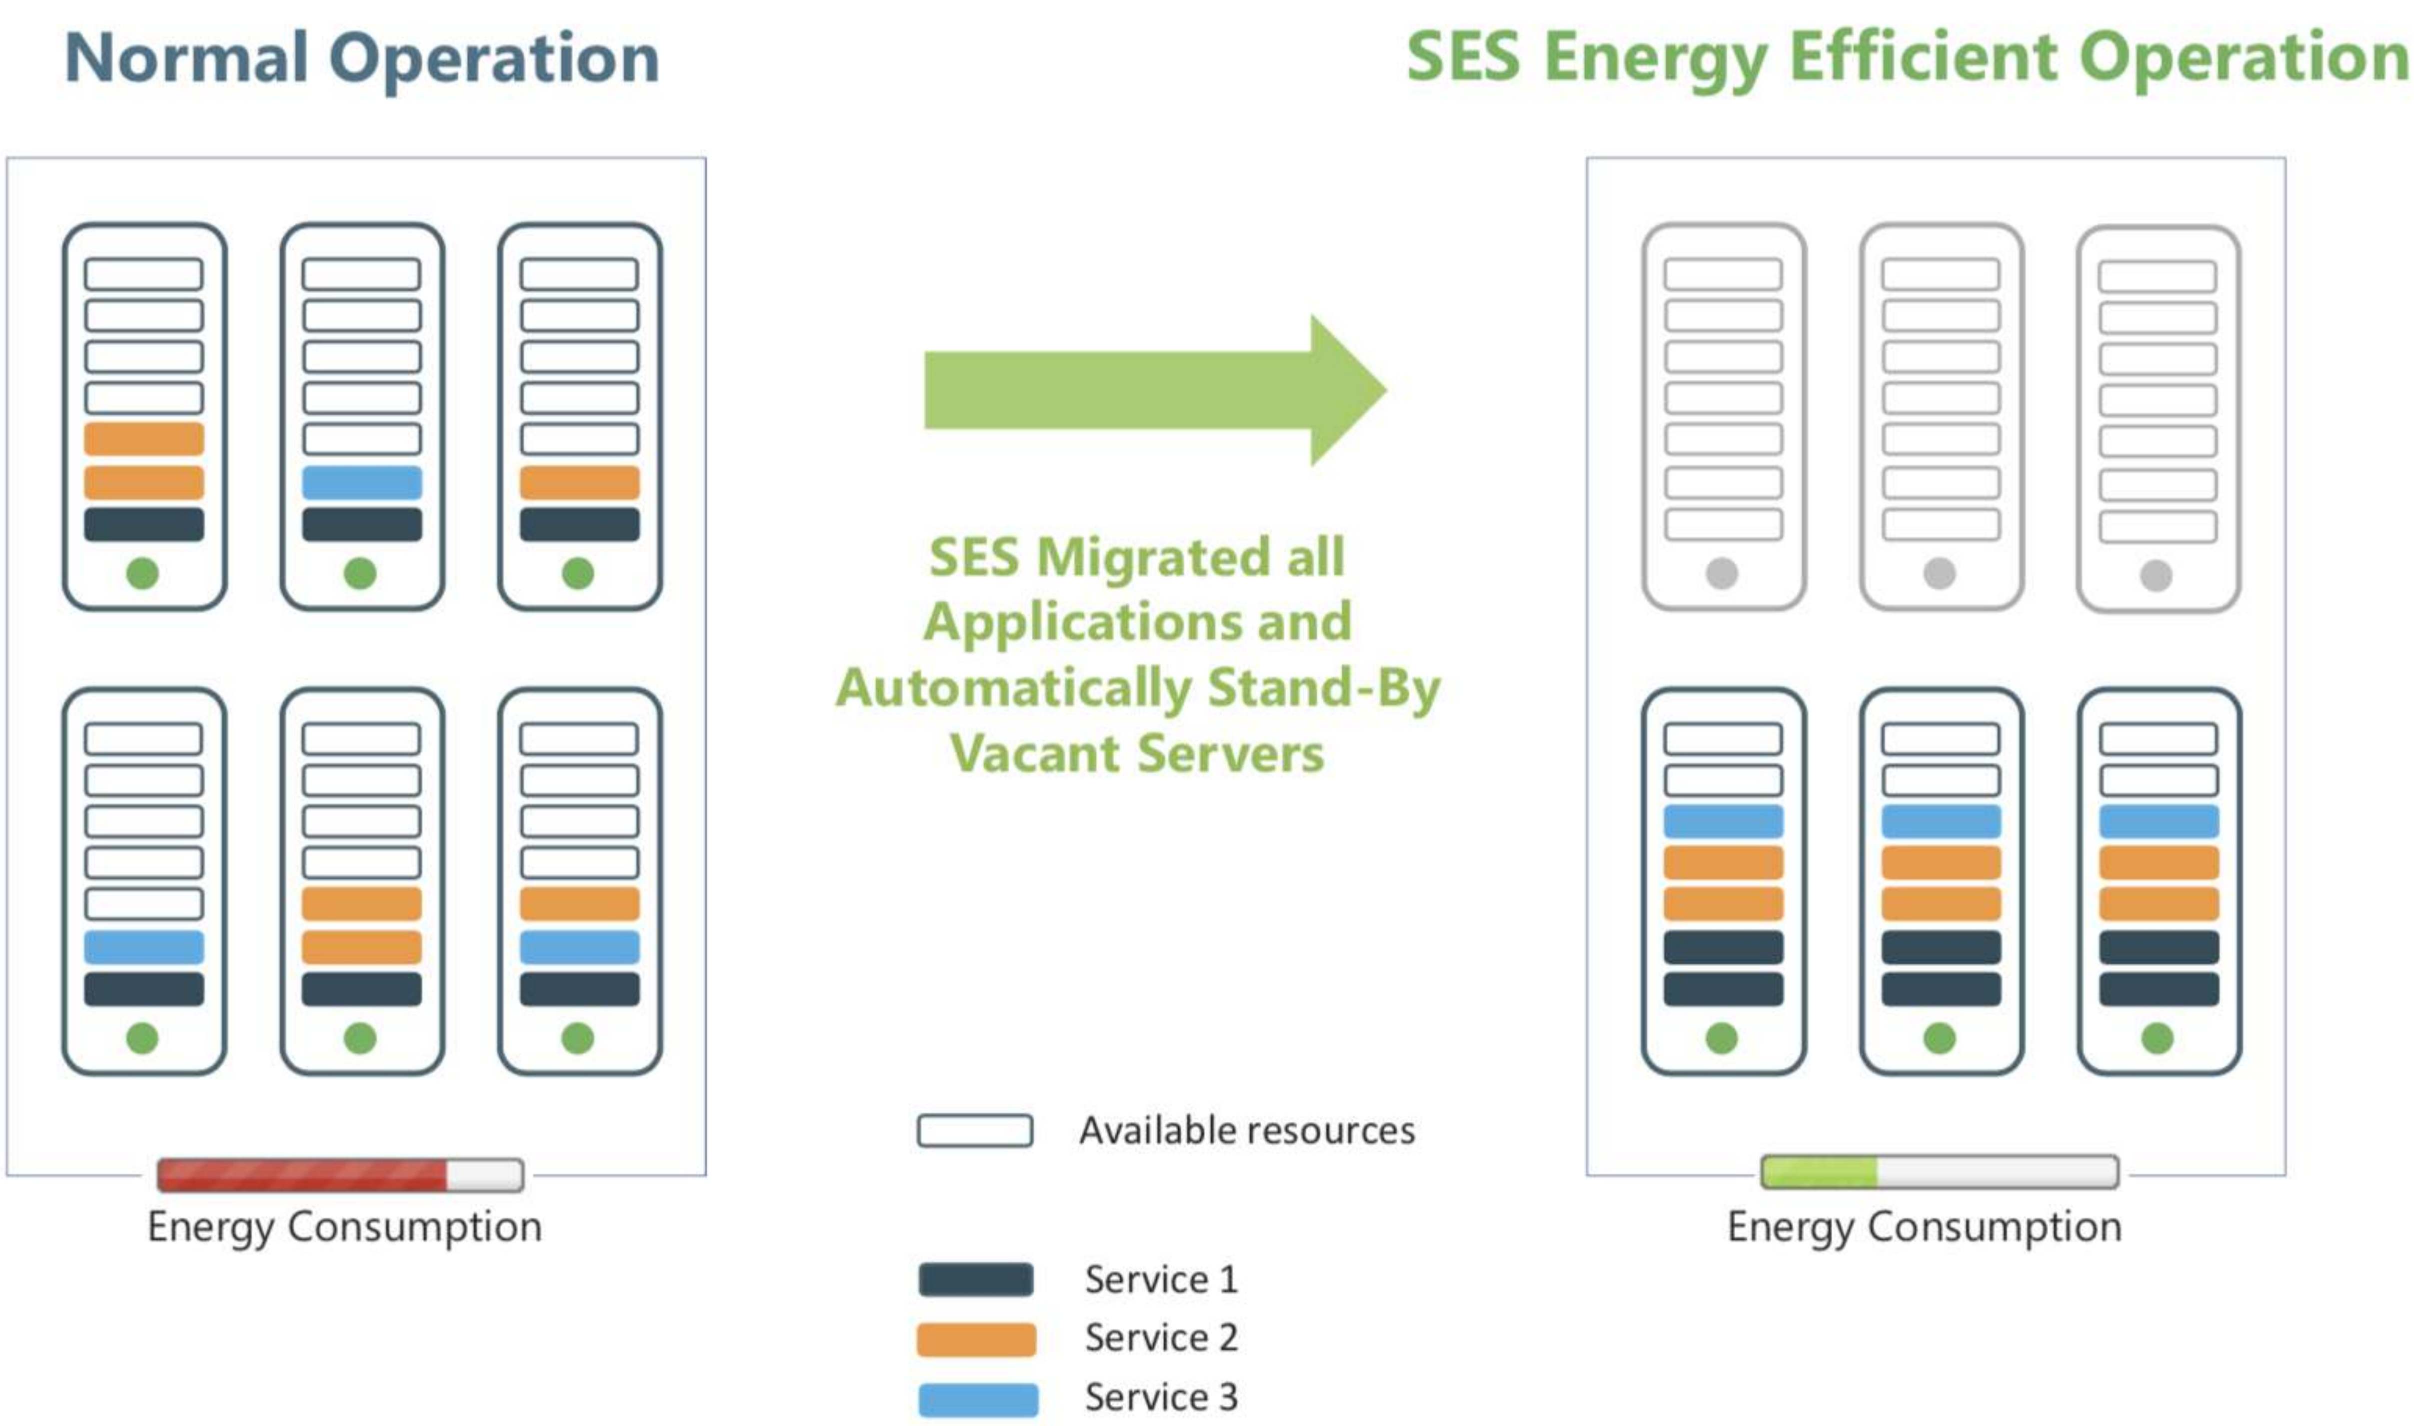

SES INSIGHTS

- Intense Resource Cluster Identification
- Cluster Optimization
- Climate Awareness
- Anomaly behavior Recognition
- Hardware Suggestion

SES Insights

Better results could be achieved by configuring the IPMI

REACTIVE Outage SYSTEM

- Prolong critical services time
- Only critical services stays up
- Real-time detection

Net Income Increase over time

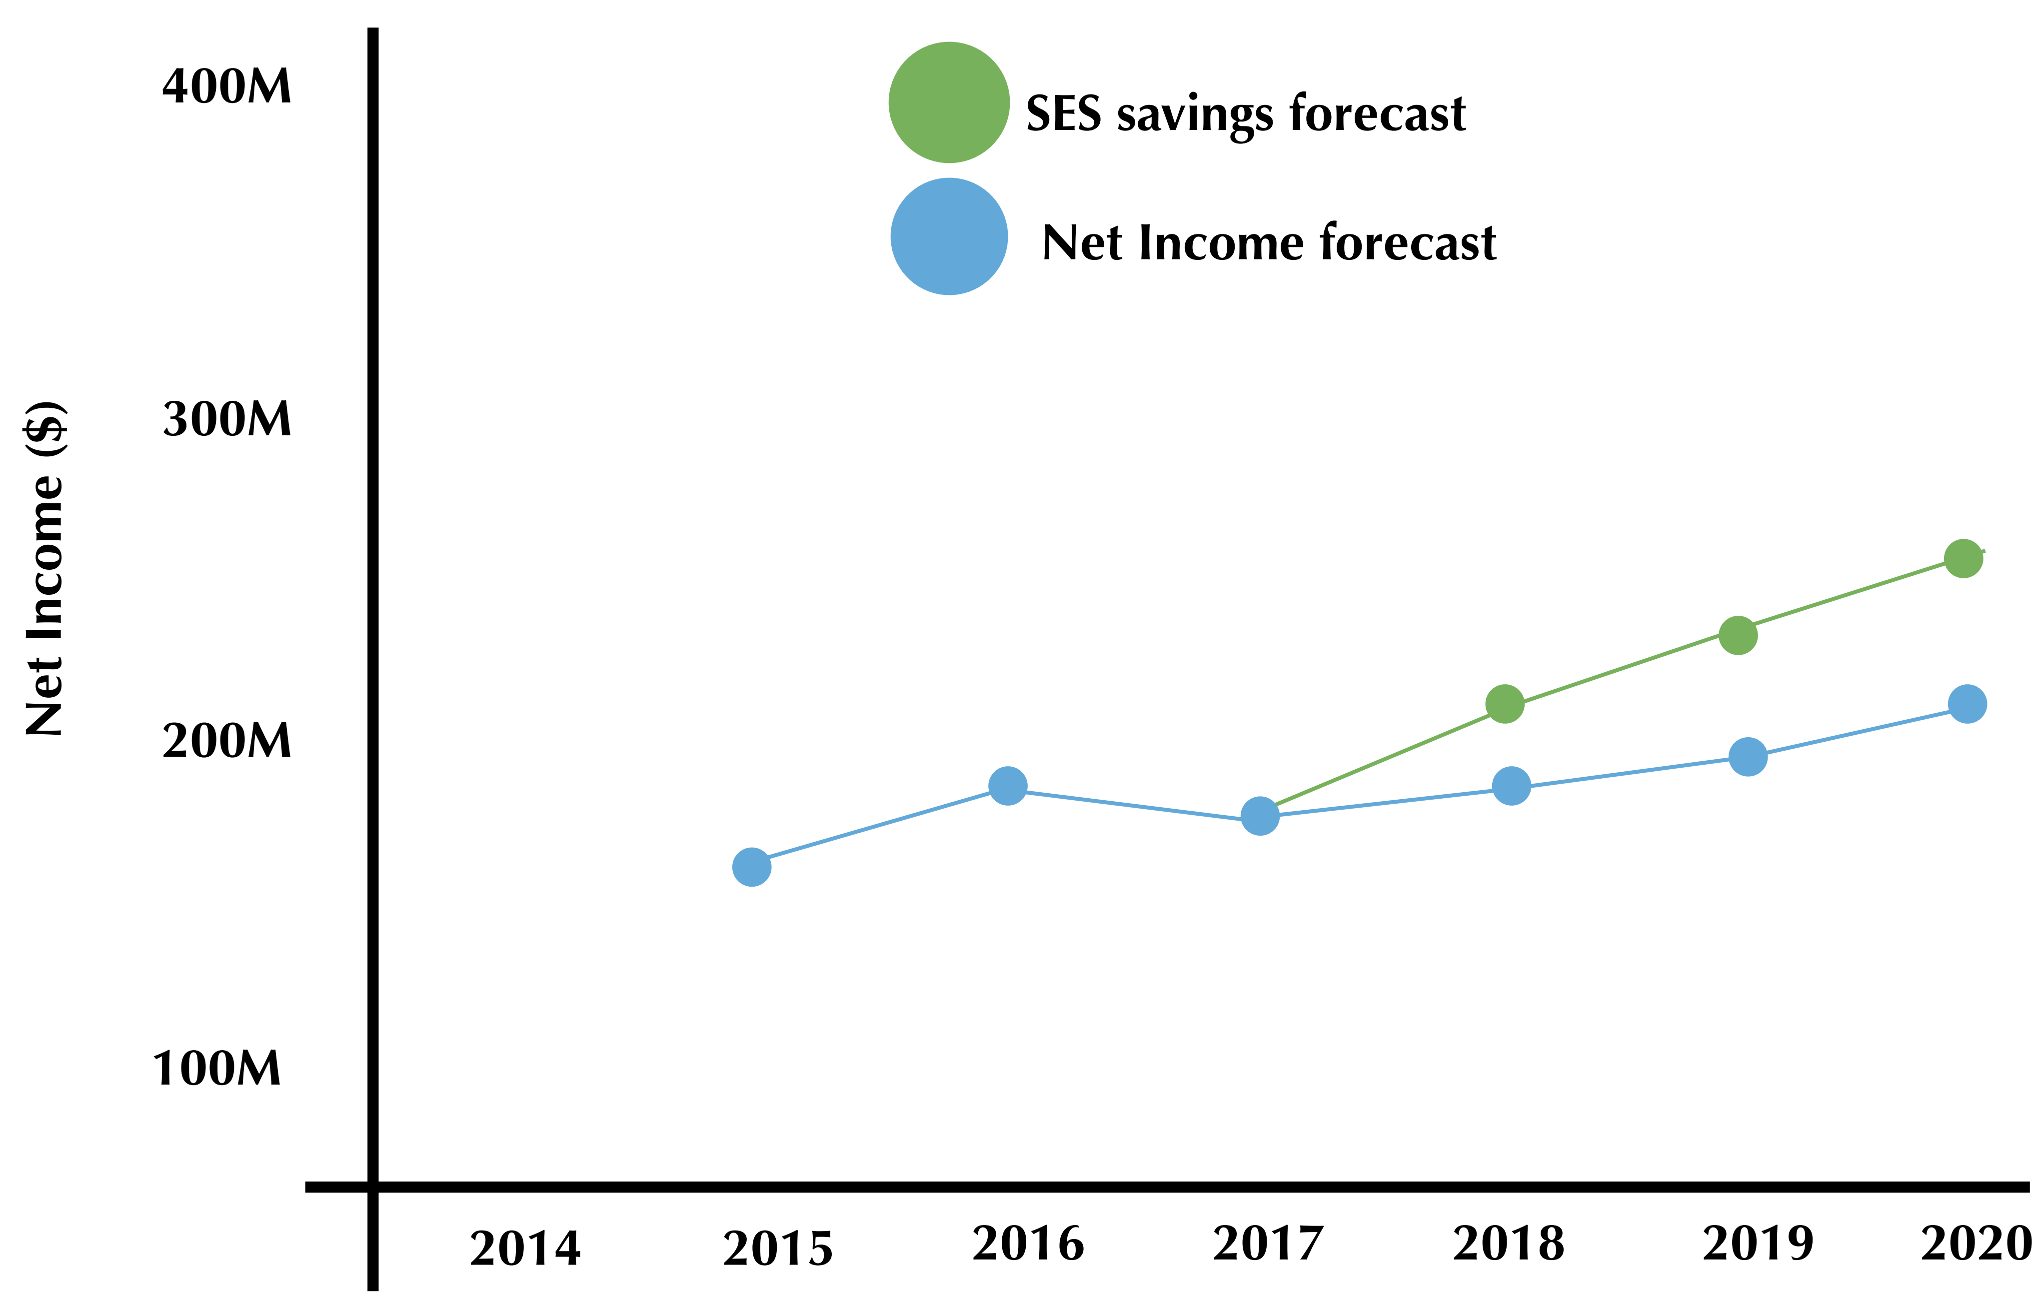

SES User Interface

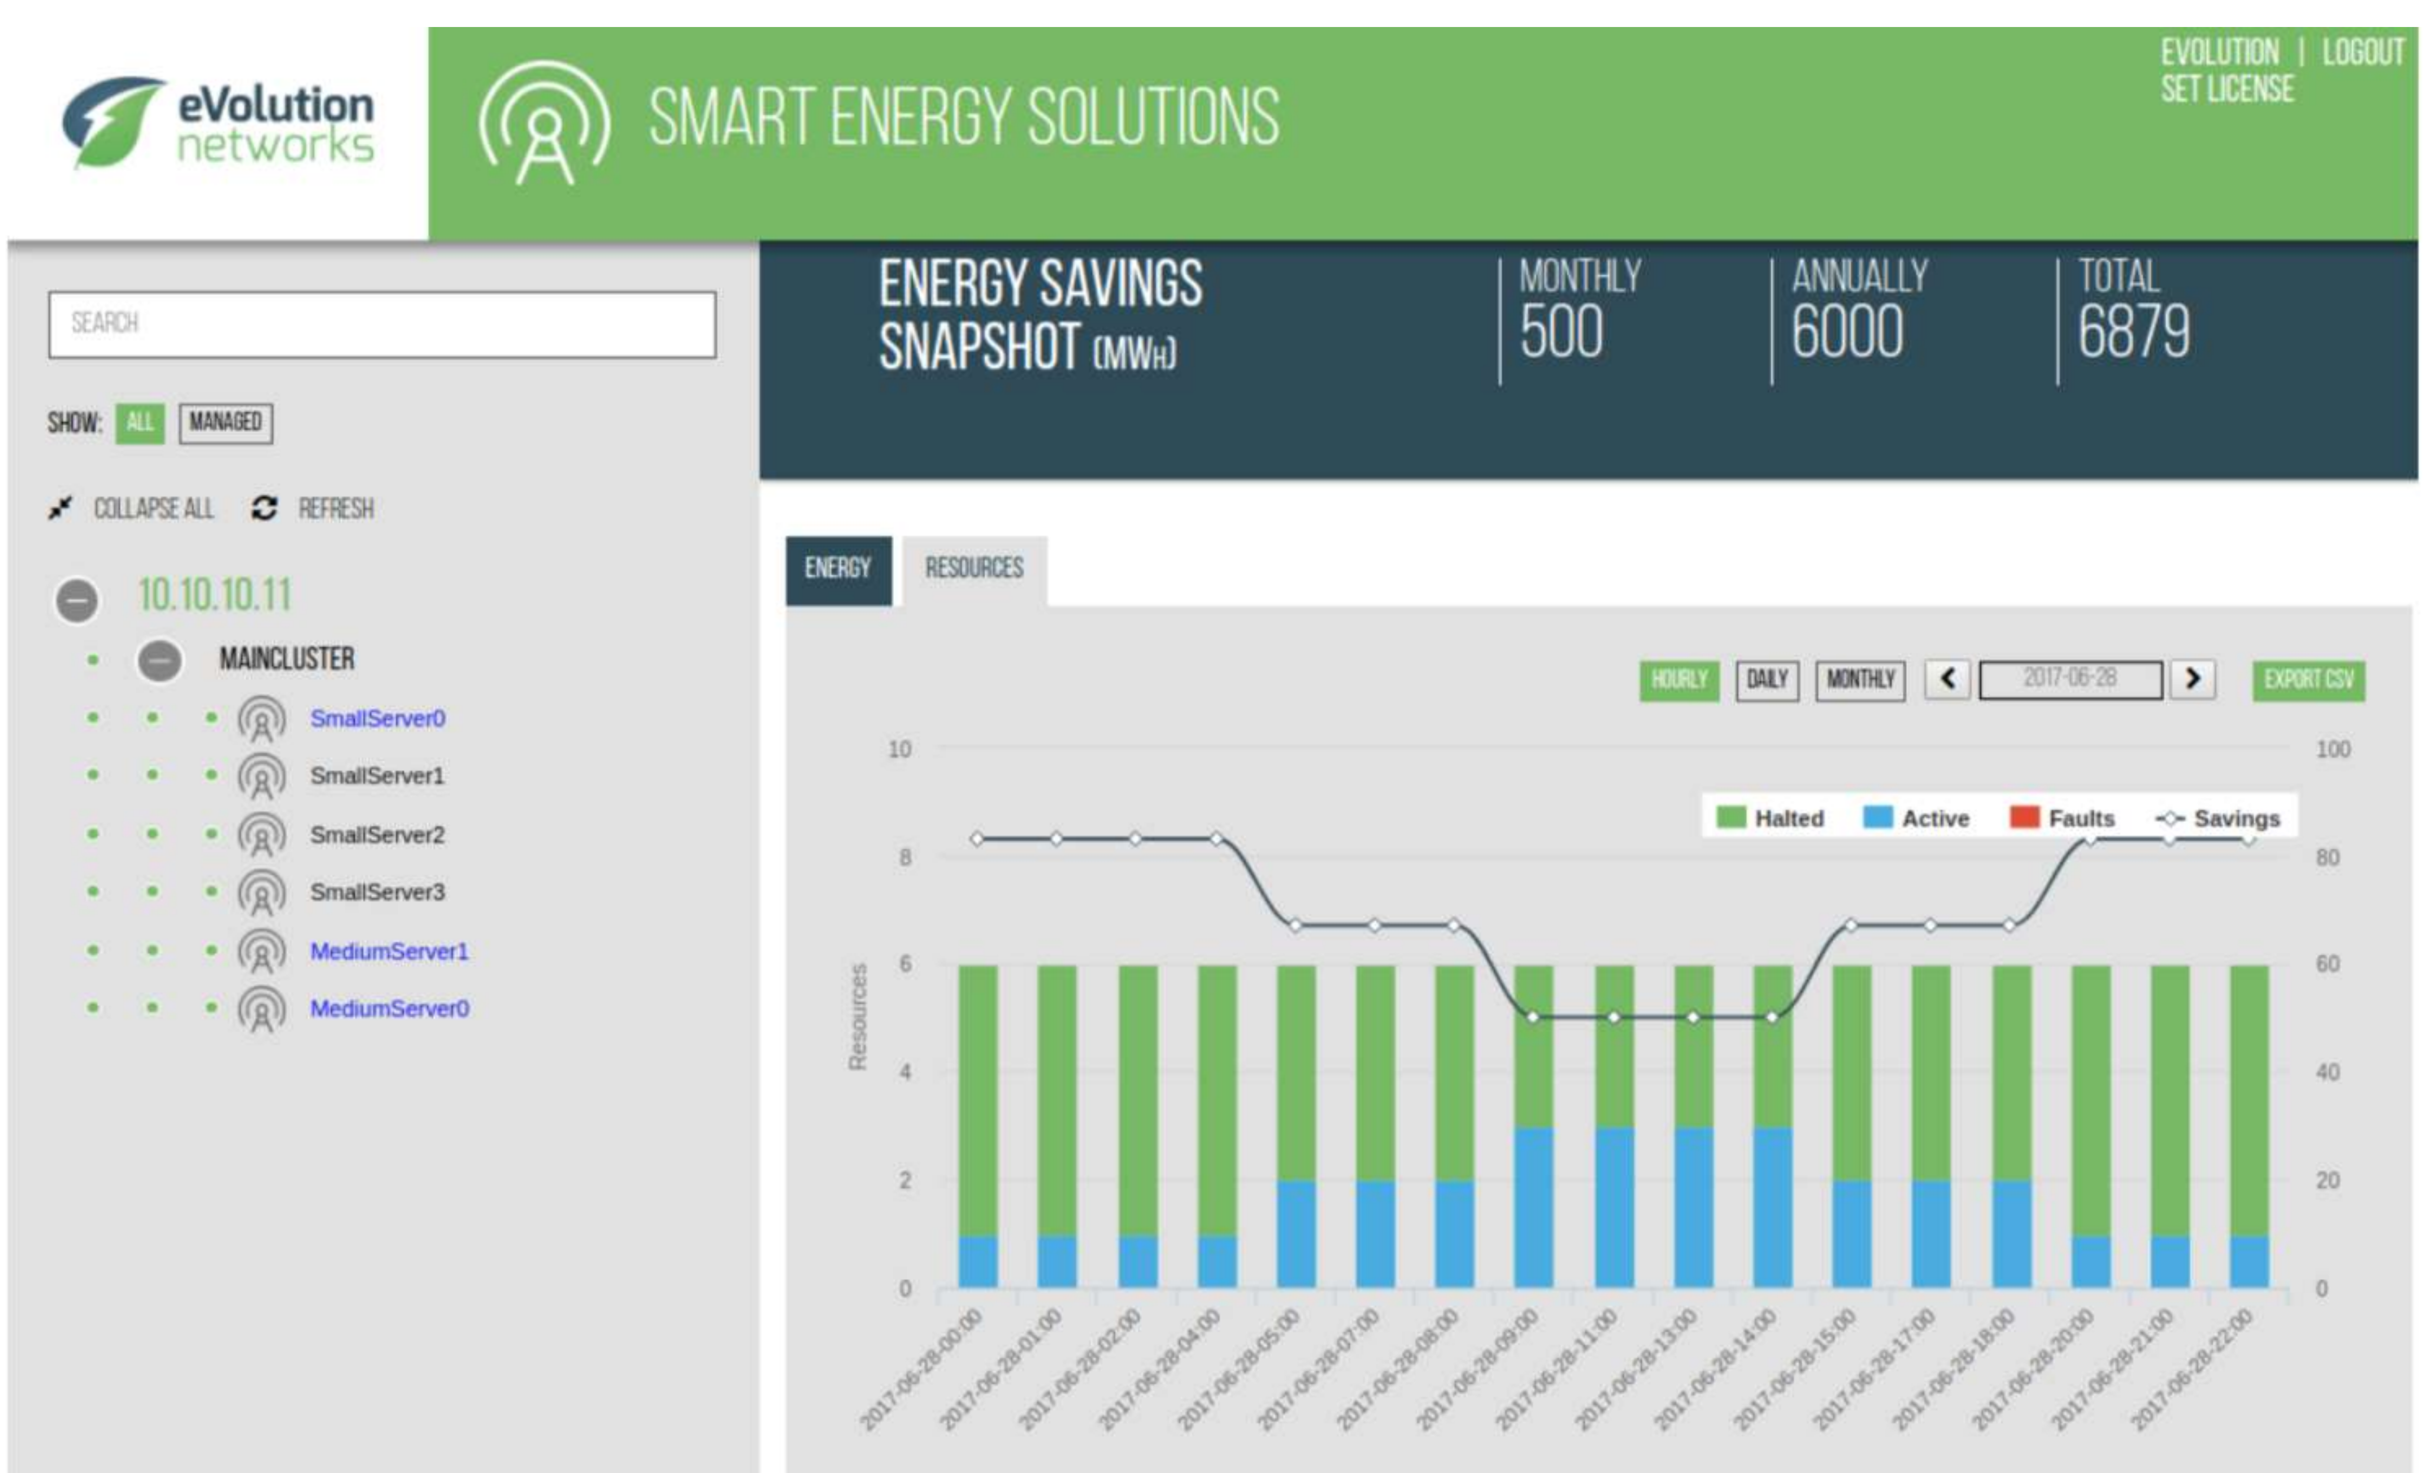

Example Report

eVolution networks

Customer Logo

Energy Savings Report

Account Name	
Number of Clusters	251
Number of Physical Servers	4,076
Number of Virtual Servers	31,588
Days System Running	43

License Type	Operational
Valid Until	30/08/2019

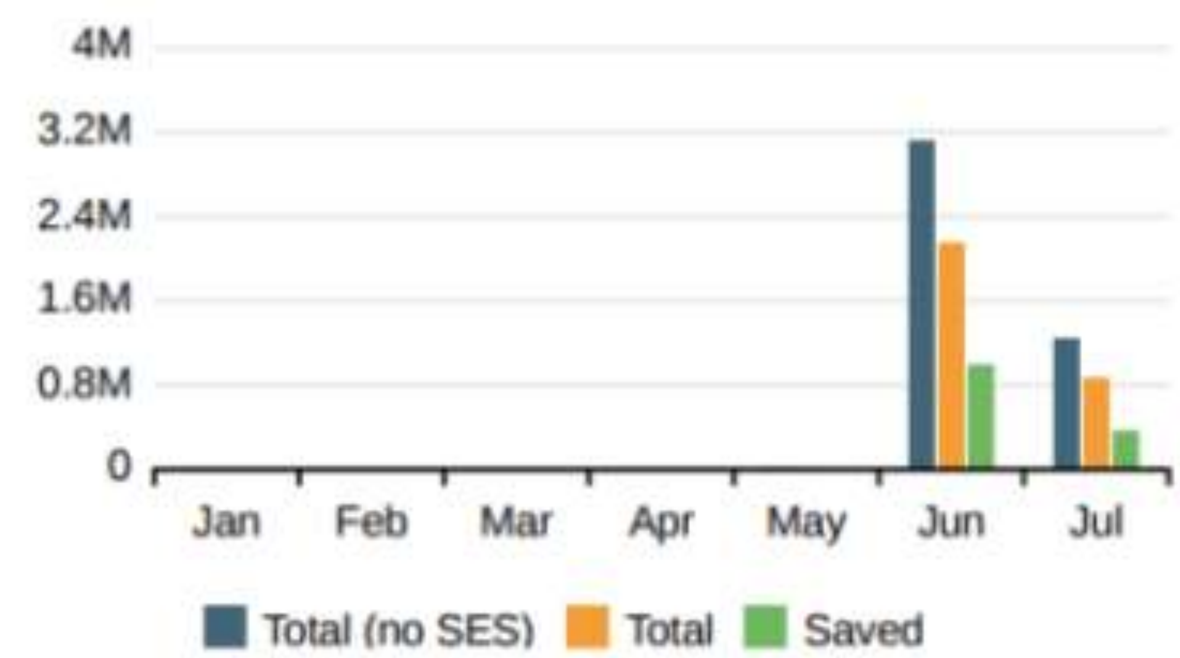

Estimated Annual Data	
Energy Consumption	10,711,419 KWh
Energy Consumption with SES	6,533,965 KWh
SES Consumption Savings	**4,177,453 KWh**
Energy Cost	1,285,370 USD
Energy Cost with SES	784,075 USD
SES Savings	**501,294 USD**
Cooling saved	156,815 USD

Consumption & Savings (Last 43 days)	
Energy Consumption	1,253,236 KWh
Energy Consumption with SES	764,473 KWh
SES Consumption Savings	**488,776 KWh**
Energy Savings (%)	39%
KWh Price	0.12 USD
Energy Cost	150,388 USD
Energy Cost with SES	91,736 USD
SES Savings	**58,653 USD**
Cooling saved	11,730 USD

aeeb88a1f3639b186773aa2c8caa717e

Page 1 of 1

SES Licenses

Assessment License

- ✓ Read-only mode
- ✓ Analyzes estimated savings
- ✓ Dashboard and Report

FREE

Operational License

- ✓ Live operation
- ✓ License by Server
- ✓ Dashboard and Report
- ✓ 24/7 Support

PREMIUM

Up to 40% Energy Savings for your DC facility

TimeLINE

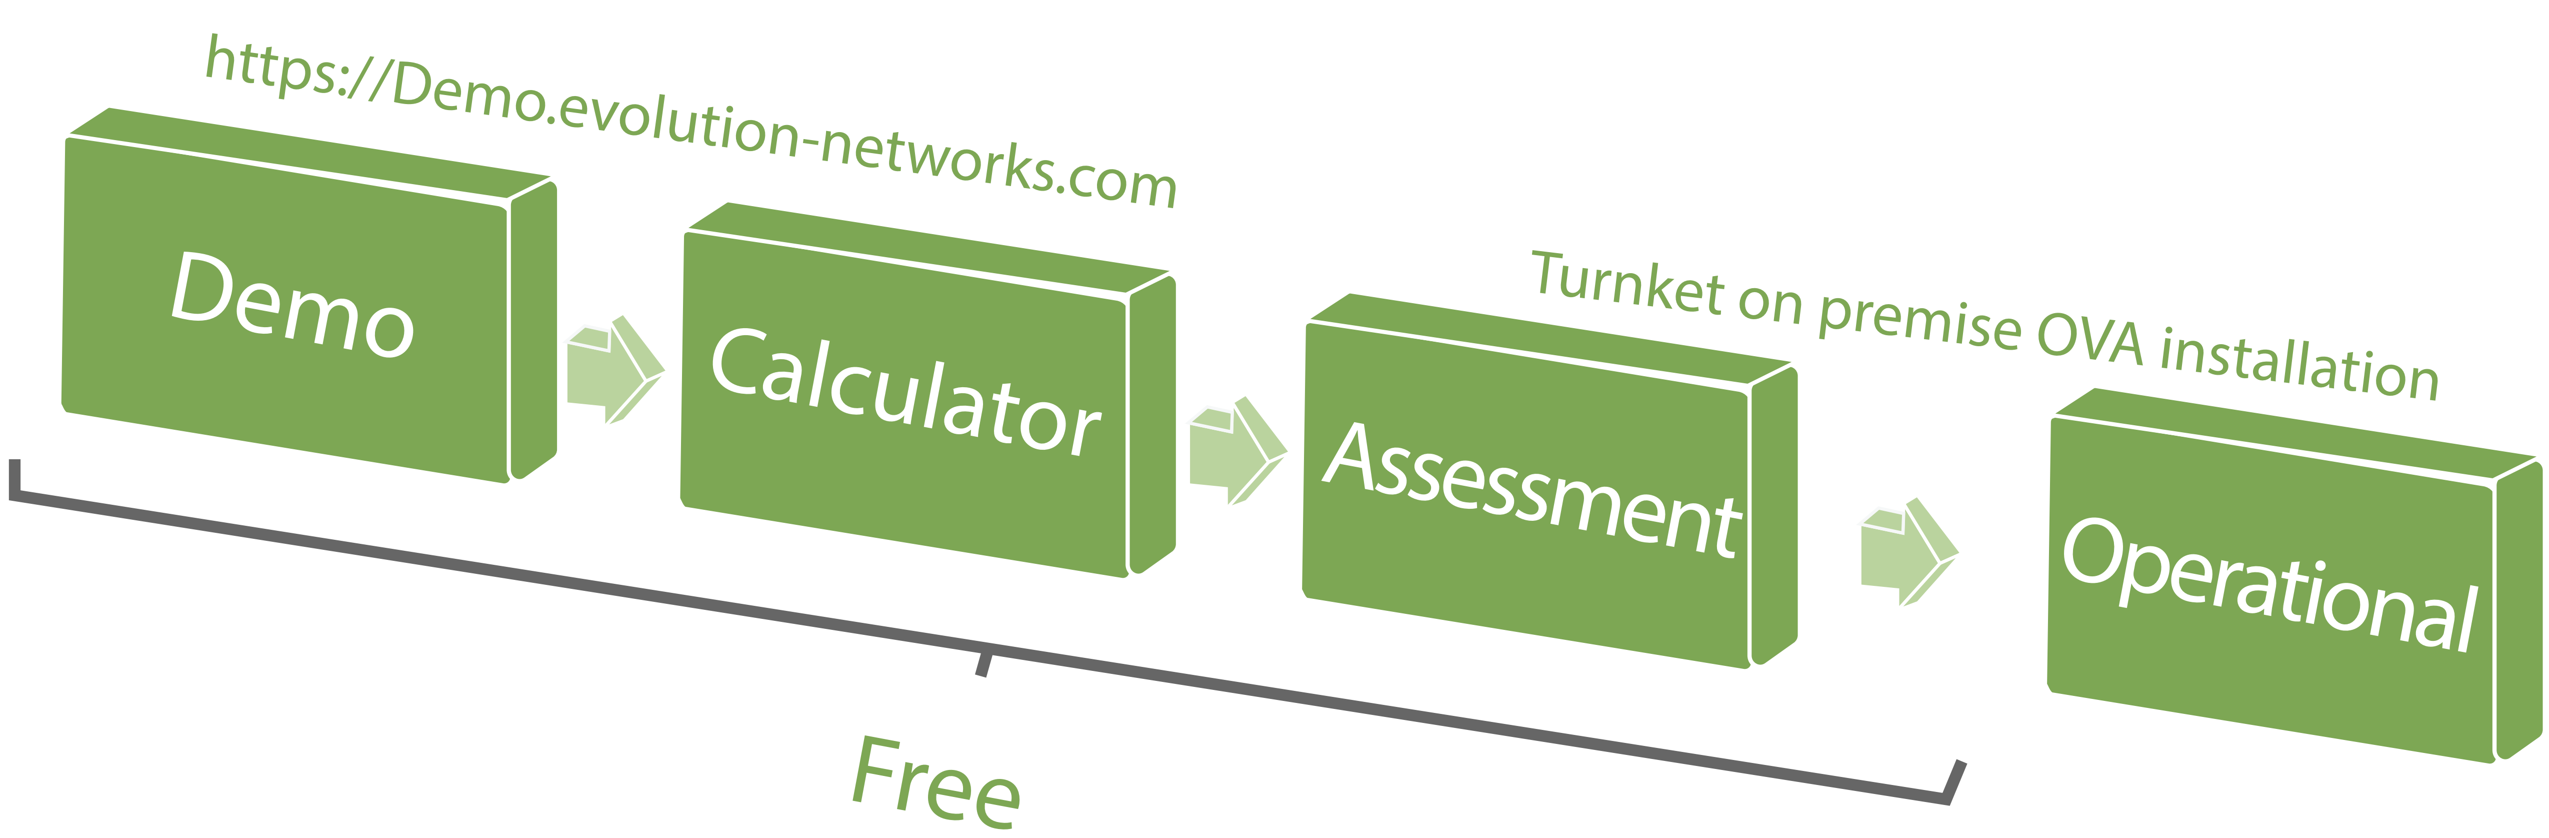

COLO offering

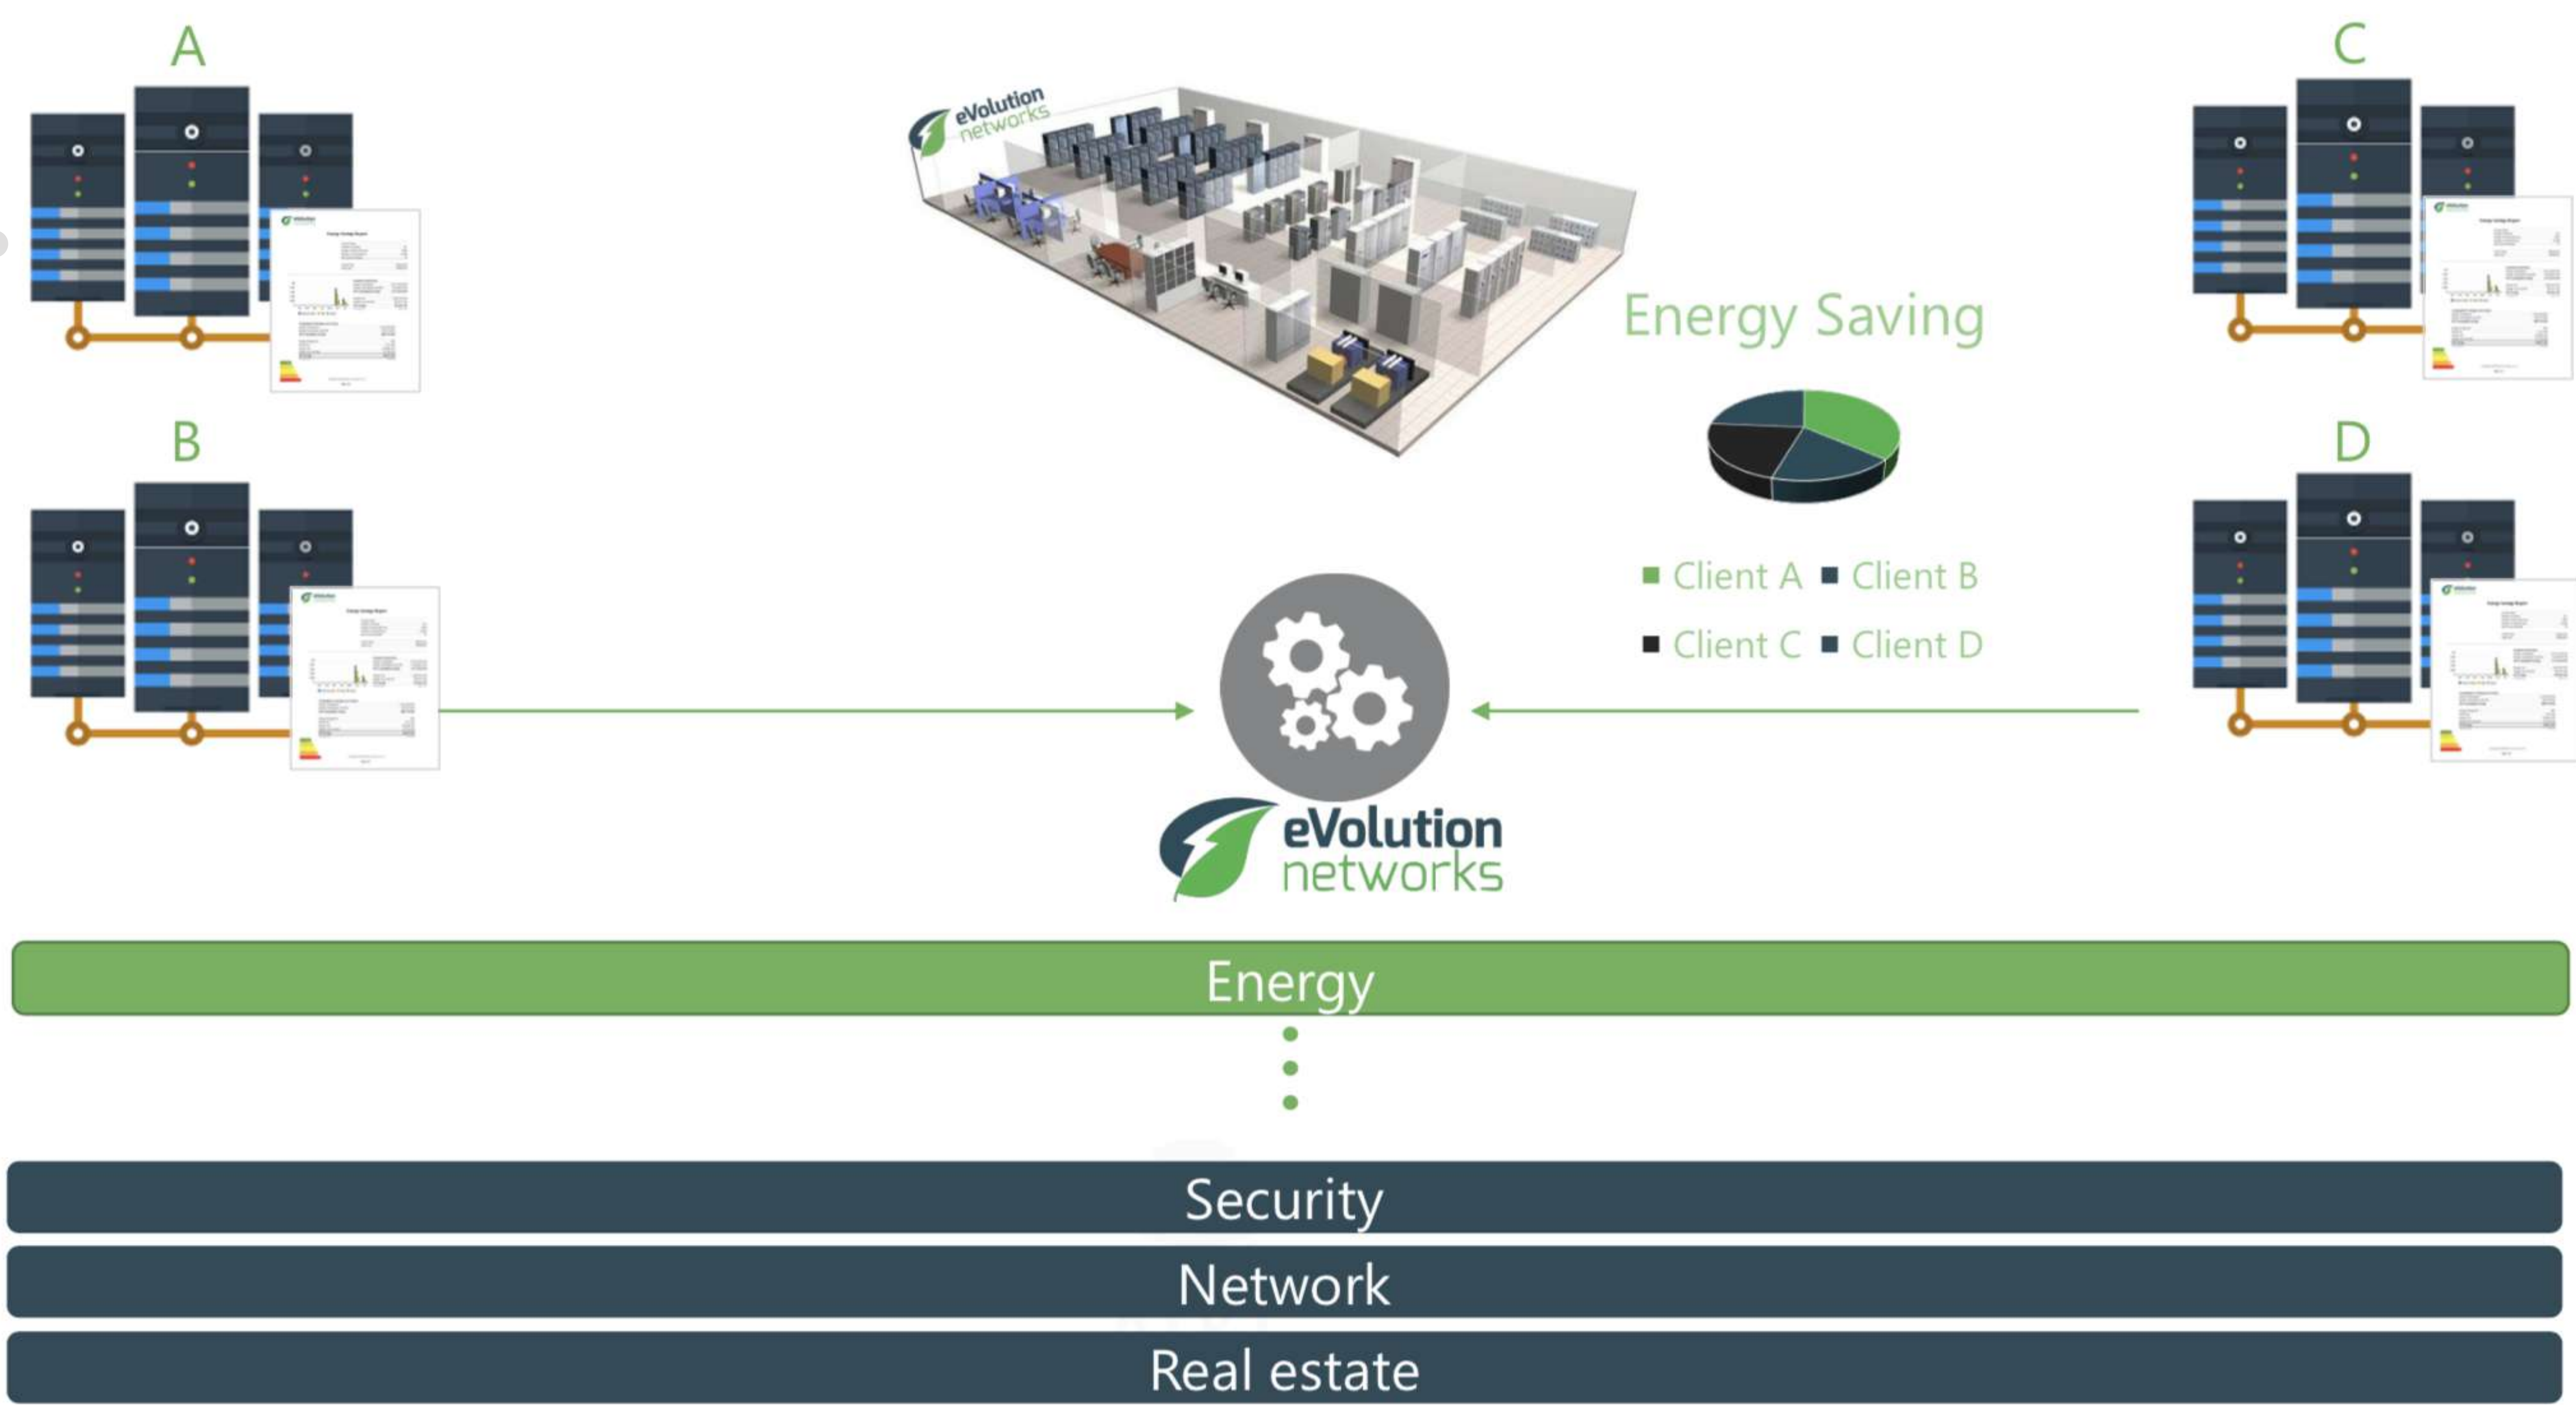

SES
Multi Tanat version

SES MT allows DC to become eVolution compatible where tenants reduction is acknowledged by the data center

- **Customized Reports**
- **Per Customer Price setting**
- **Aggregated Reporting View**
- **Optional Management View**

Harness the power of AI and Machine Learning to save 40% of your Server Energy bill while keeping your SLA's

Reduce your carbon footprint significantly

Accredited by vmware and openstack COMPATIBLE

Start saving from day 1
Save upwards of 80K US$ annually for every 1000 servers
Payback in less than 6 months

Self-Deploy in minutes with a free assessment to view potential savings

Connected only to Native Certified API Runs locally, no outbound connectivity needed Highly secure local implementation

---

Intuitive web-GUI
Realtime alerts
Built-in Report generator and Savings analysis tools Localized user interface

SMART ENERGY SOLUTION

SILVER
PACKAGE

LICENSE KEY

40% REDUCTION of your data center energy costs

Check out our online demo
https://demo.evolution-networks.com

YOUR FREE ANALYSIS LICENSE

TRY IT

Level 7 Al Murjanah Tower
Price Sultan St. Ar Rawdah
Corner of Prince Sultan St. and Alkayyal St.
PO Box 10113 Jeddah 21433,KSA
+966 (12)601-7790
info@almarjeia.com.sa
sales@almarjeia.com.sa
www.almarjeia.com.sa