Development of a Very Dense Liquid Cooled Compute Platform

Phil Hughes
CEO, Clustered Systems Company, Inc.
phil@clusteredsystems.com
415 613 9264

Goals

- Stated
  - Achieve a data center infrastructure efficiency (DCiE) of better than 0.85 vs today’s average of approx. 0.61
  - Be capable of effectively supporting HPC applications yet have a manufacturing cost that makes it cost effective for mainstream applications.
  - Have the capability of being cooled without the use of refrigeration systems for most of the year in the majority of US locations.
- Extended
  - 100KW power and cooling per rack
  - Open platform enabled for Compute, Storage, Networking
  - Upgraded power train, virtual I/O for additional energy saving

Team Members/Contributors

Team

Clustered Systems Company
extreme green electronics cooling

Liebert
Keeping Business in Business™

EMTEC

SLAC
NATIONAL ACCELERATOR LABORATORY

Funding

DEPARTMENT OF ENERGY
UNITED STATES OF AMERICA

STATE OF CALIFORNIA
ENERGY COMMISSION

Contributors

aqs

ONE STOP SYSTEMS

intel

EMERSON
Network Power

SMART
Modular Technologies

California
Machining
Center

Bodycote

TechShop
BUILD YOUR DREAMS HERE

The Choice…..

Air

Refrigerant

5500x volume →


©CoghillCartooning.com

90% Less Energy

1000% Higher Compute Density

- 👍 Small Pipes
- 👍 Few pumps

- 👎 **BIG** Ducts
- 👎 Many fans

The Choice (cont)….

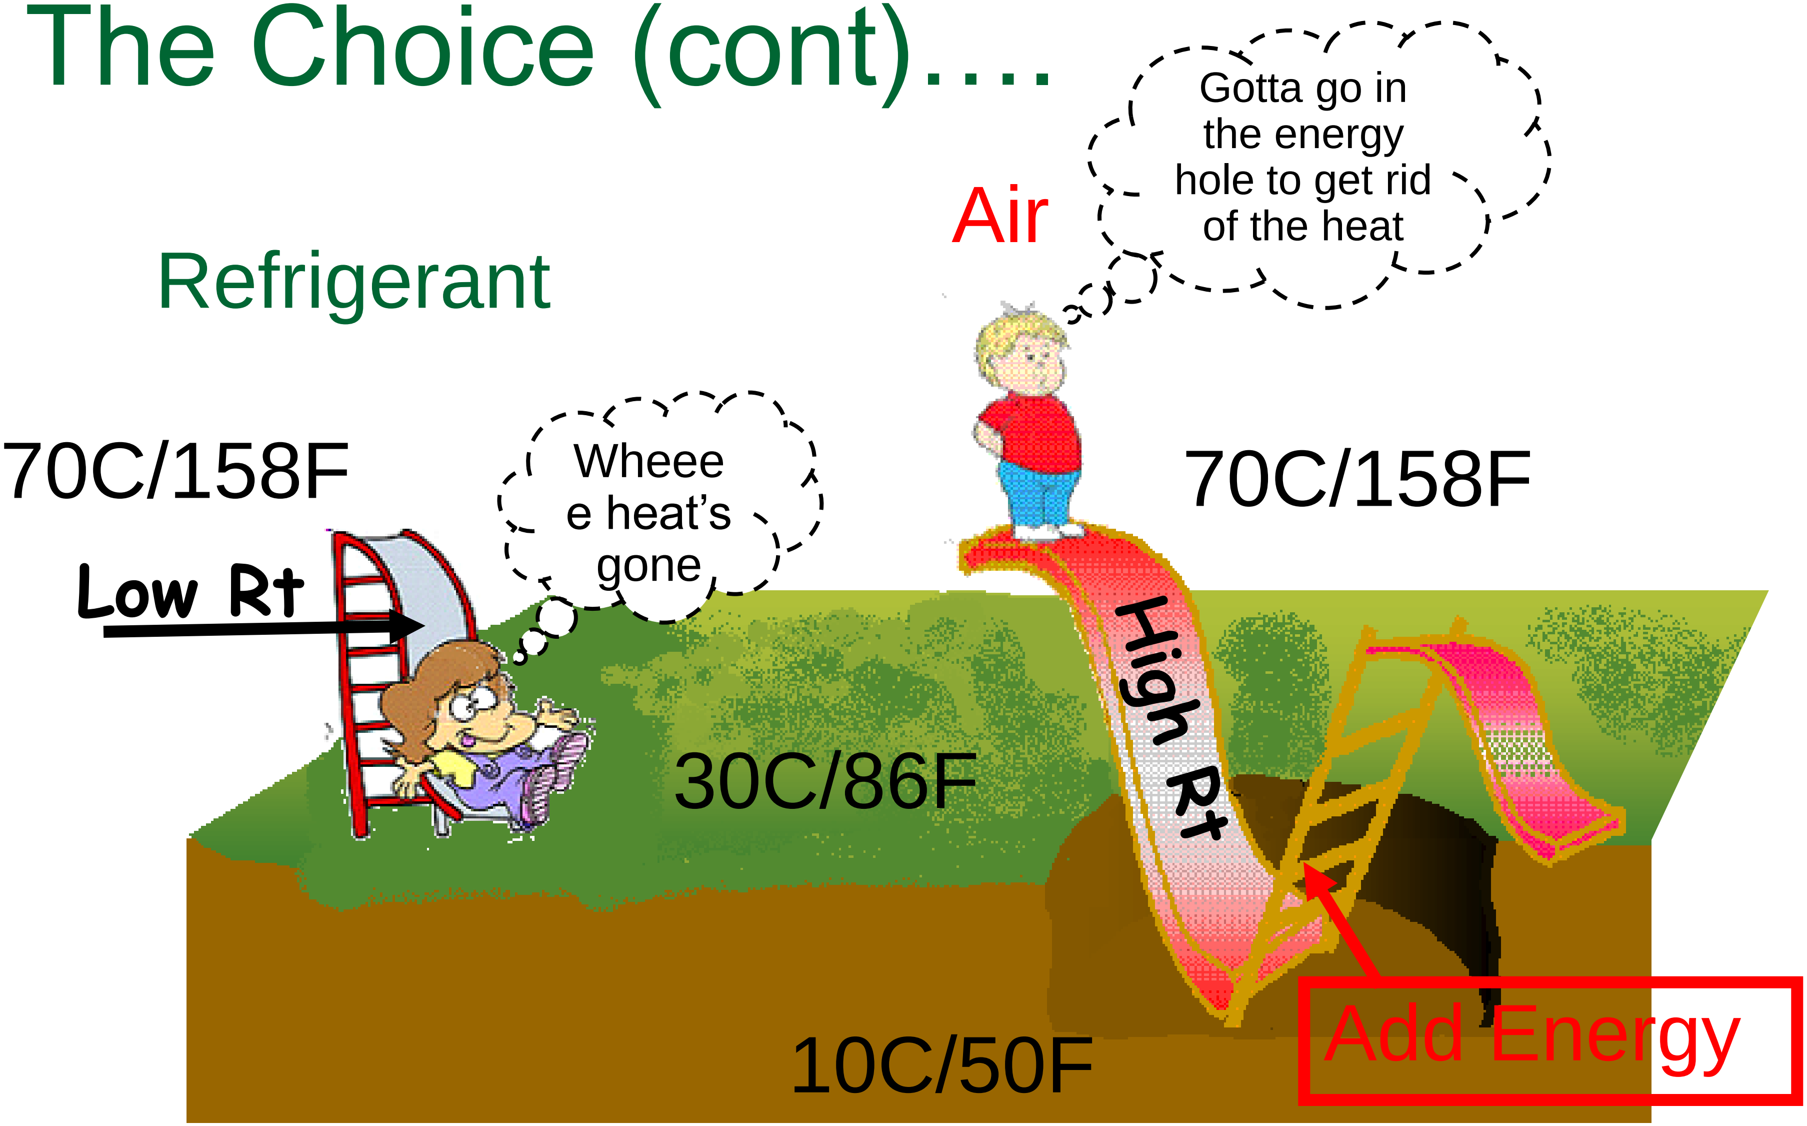

Control – Or Not

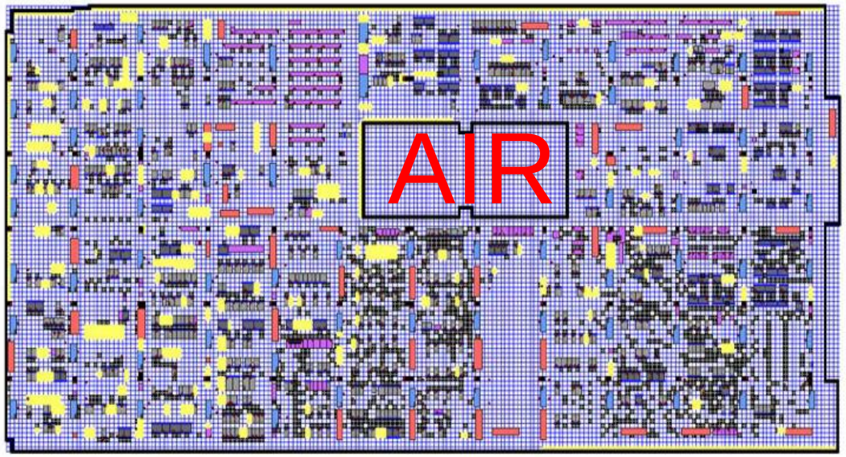

922 thermal sensors
62 pressure sensors
72 flow sensors
42 humidity sensors

87 in plenum
72 discharge + 72 return
492 @ inlet (class #1)
283 @ network

Refrigerant

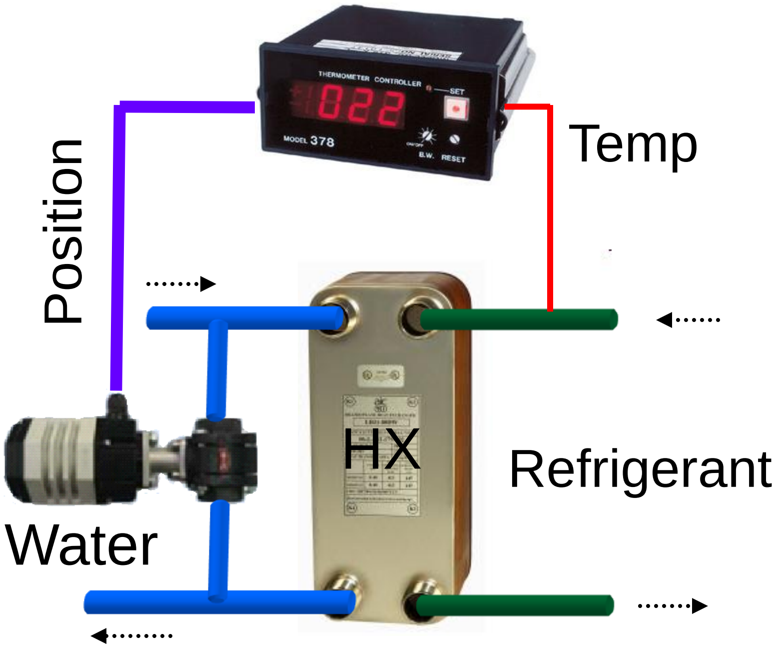

Simple proportional control is sufficient

Refrigerant section is isothermal

Shadows

AIR

30C >37C >40C

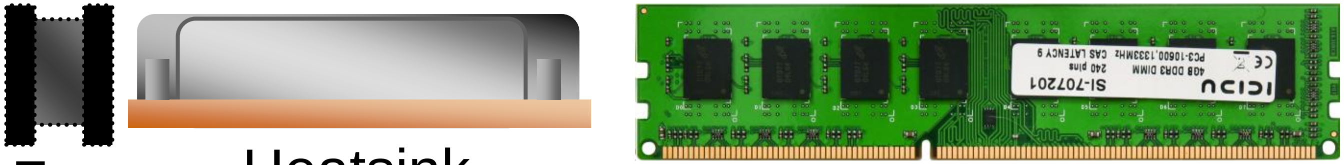

Fan

Heatsink

DIMM

Hot air kills DIMMs

Refrigerant

30C 30C 30C

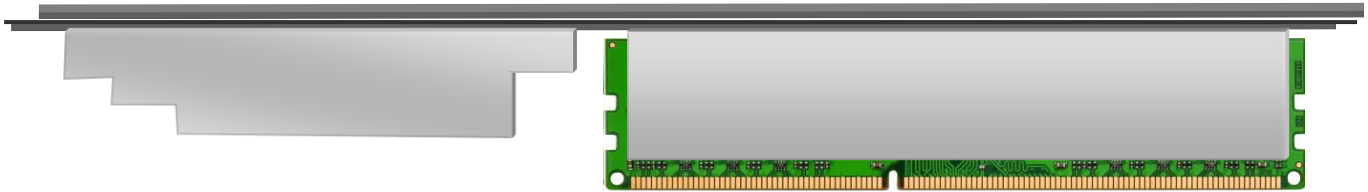

Cold Plate
TIM
Lid
TIM

Heat Riser

Jacketed DIMM

System is Isothermal

Where the Power Goes: Air

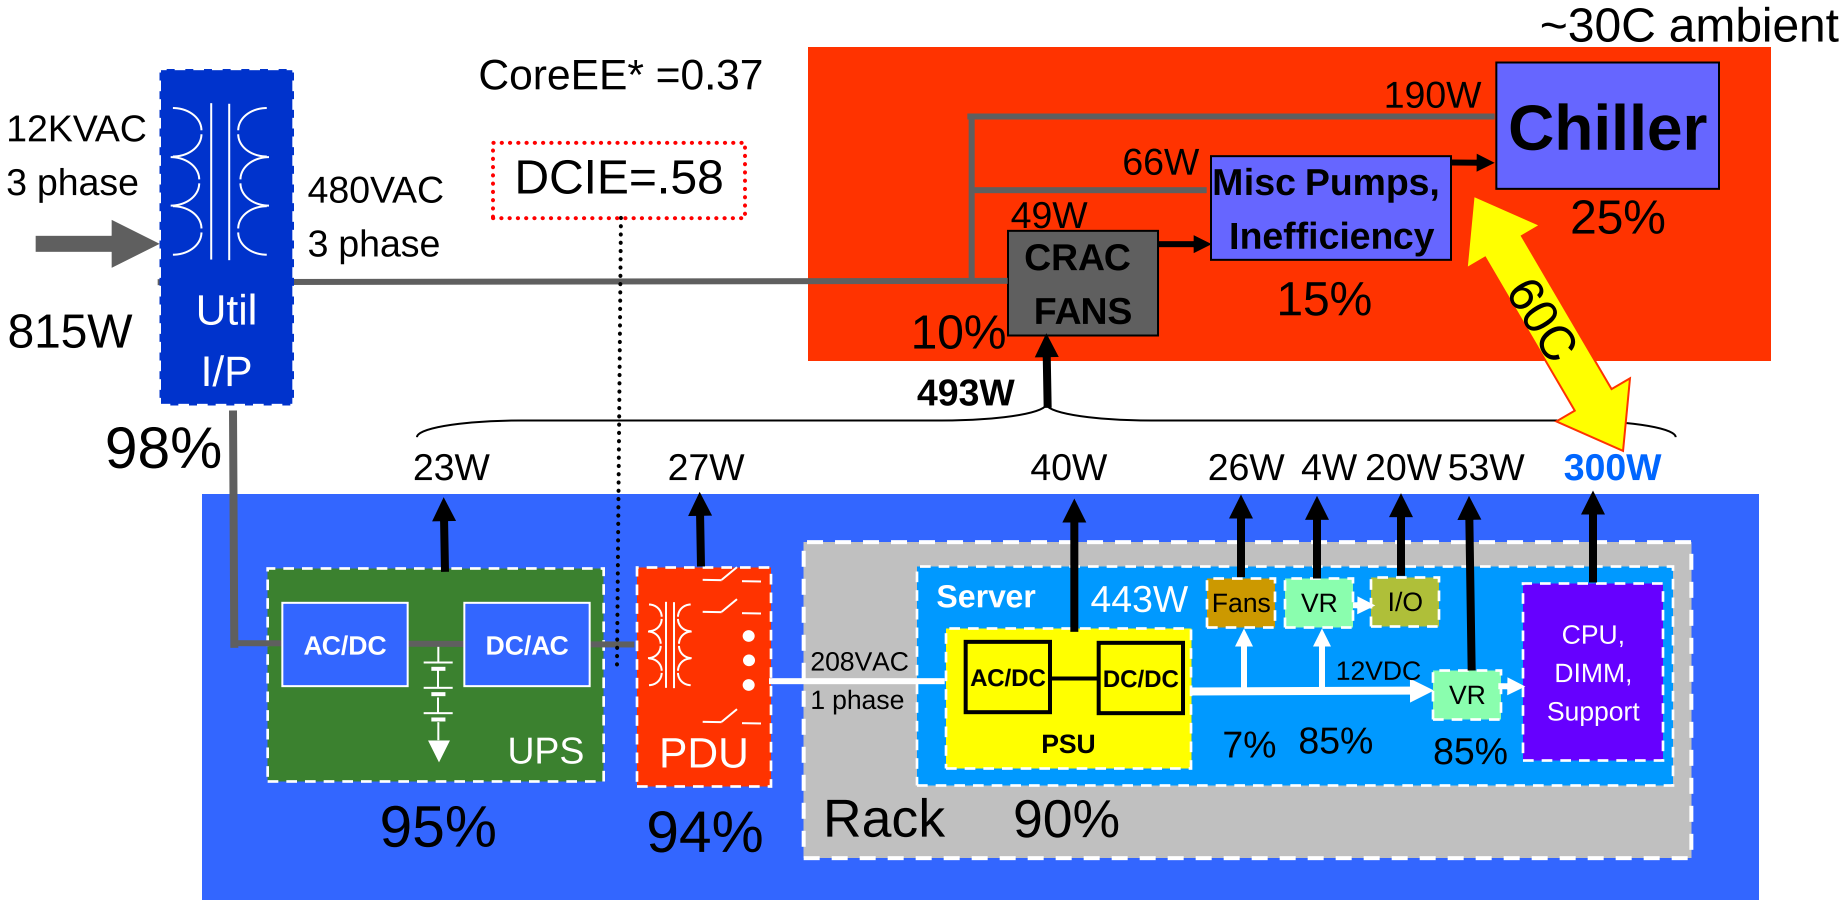

- Legacy air cooled system with water cooled chiller
- Air system’s high thermal resistance forces chiller use

*CPU+DIMM pwr/total power

Where the Power Goes: Refrigerant

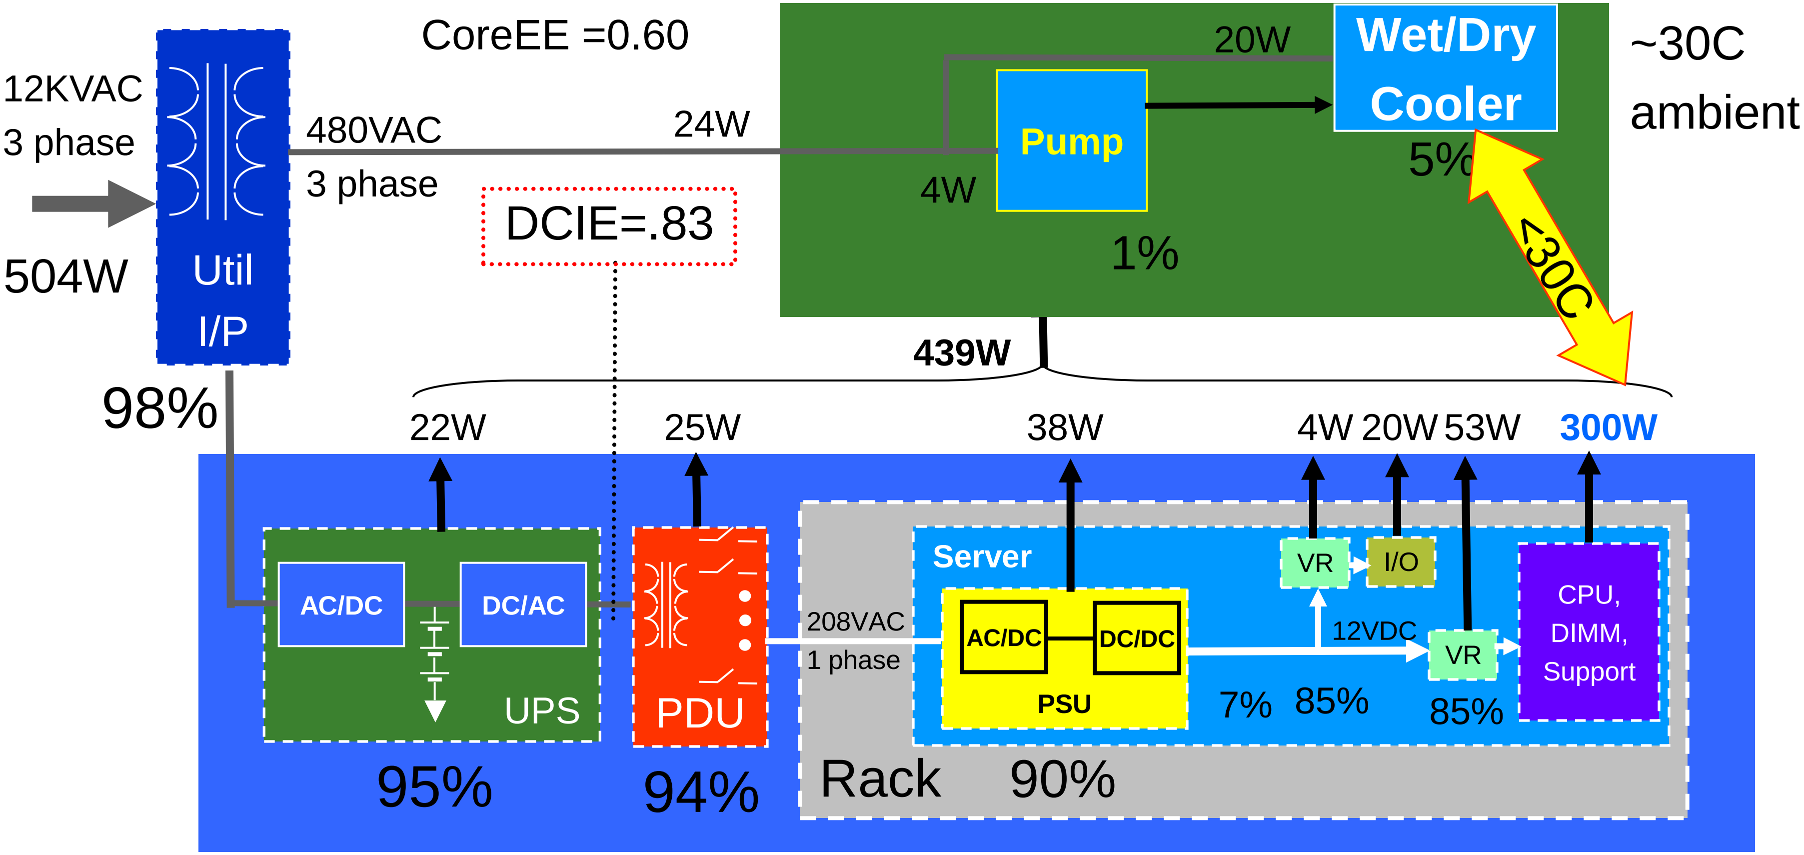

- All fans eliminated
- Lowered thermal resistance enables 95% economizer use or total chiller elimination

Direct Touch Refrigerant

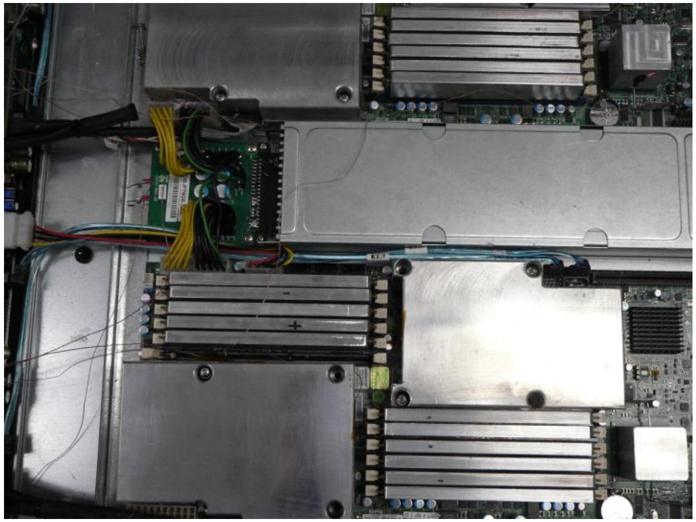

Converted 4 CPU Server Example:

Fans & Heatsinks replaced by heat risers

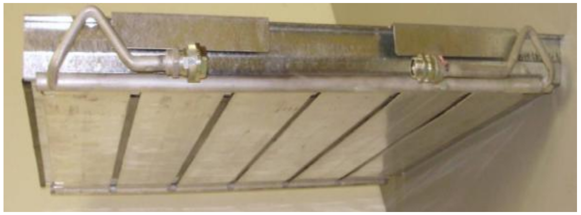

Rack mounted cold plates (1/server) remove heat through specially treated server lid

Liebert XDS

20KW & 40KW options

Chill Off II Results

Direct touch cooling demonstrated most efficient

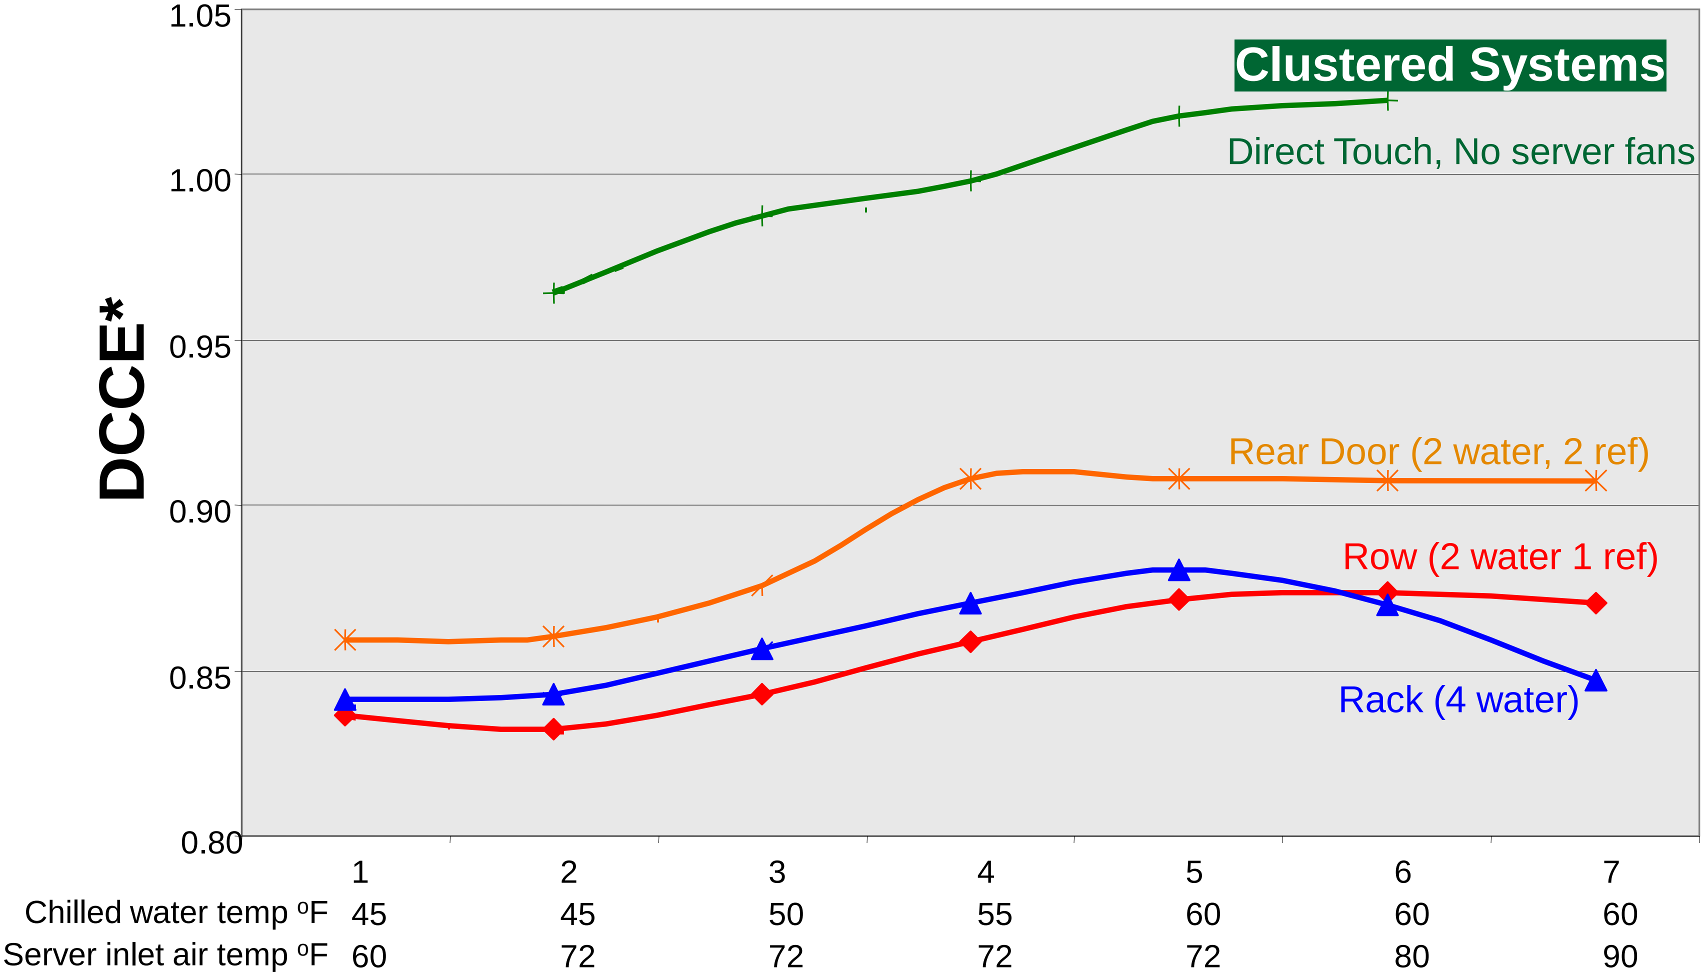

* Data Center Cooling Efficiency

Source: Lawrence Berkeley Labs

15%-20% Additional Energy Savings

I/O Virtualization, DC power

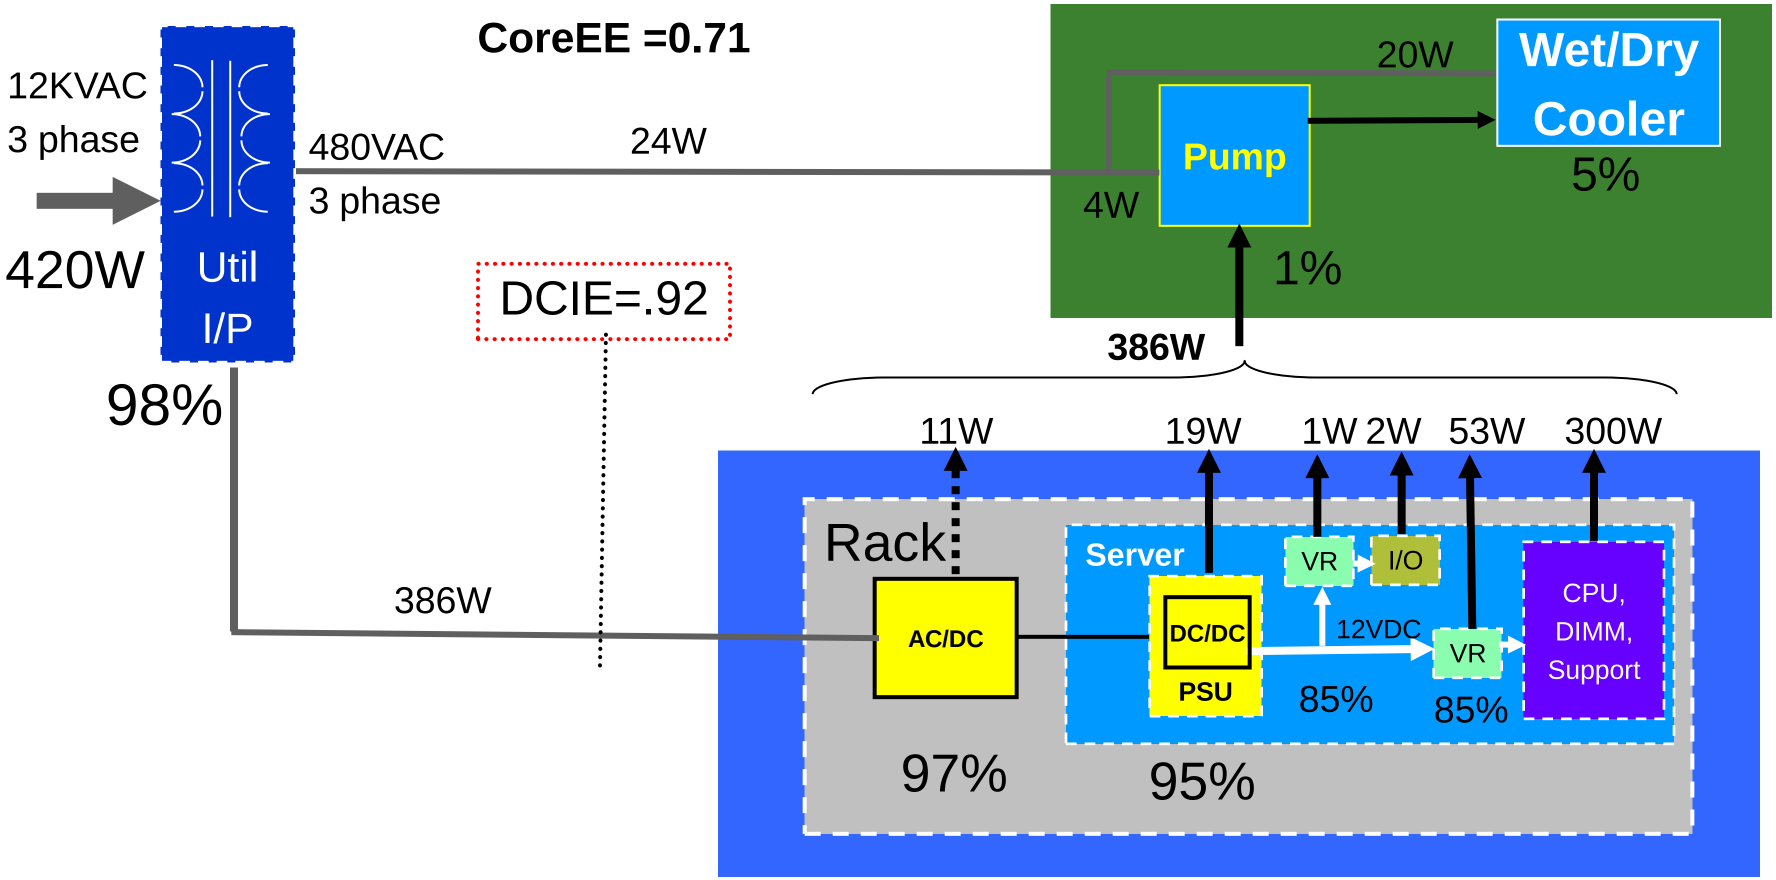

100KW Rack Overview

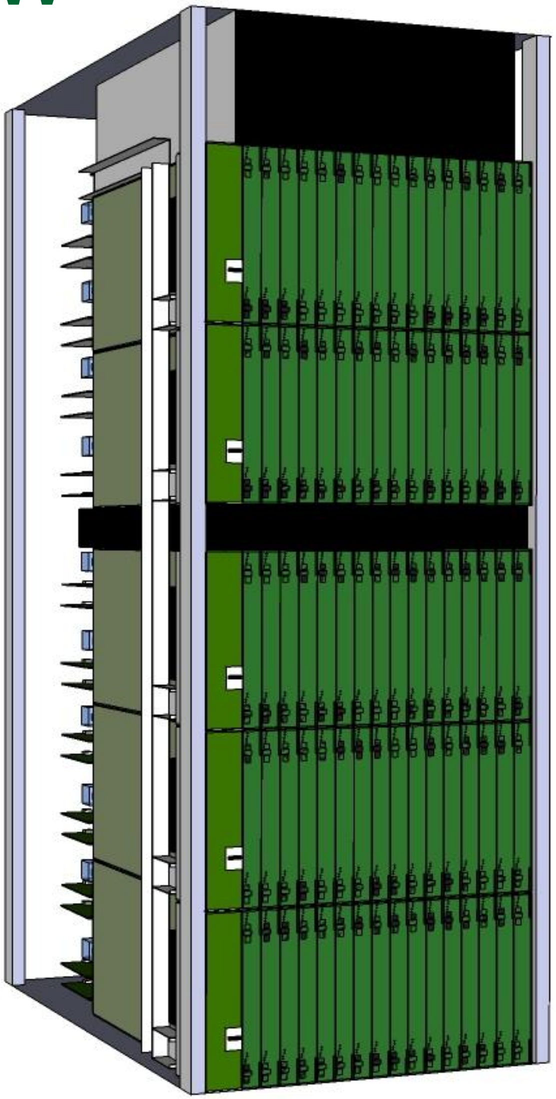

- 5 Chassis
  - 80 blades
- 100KW power & cooling
  - Proprietary non invasive liquid cooling system
  - 480VAC to 380VDC rectifier available
- Open System
- All electronics field replaceable

Architecture-Chassis

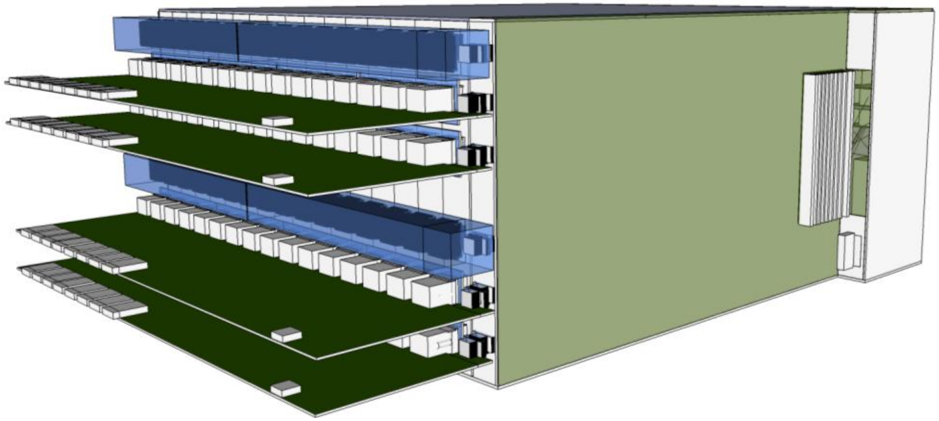

- 16 blade capacity
  - 1-2 circuit boards in each
  - 2 cold plates per blade
- 4 removable activeplanes for switching
- 20 KW power distribution unit

Cooling System

- 2 rows of 16 cold plates for motherboard cooling
  - >600W per plate capability
- 4 plates for activeplane cooling
  - 500W per plate

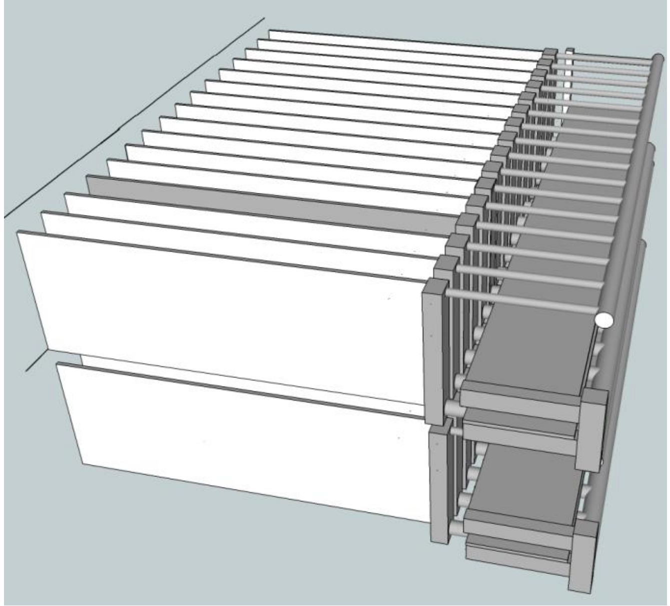

Power Train

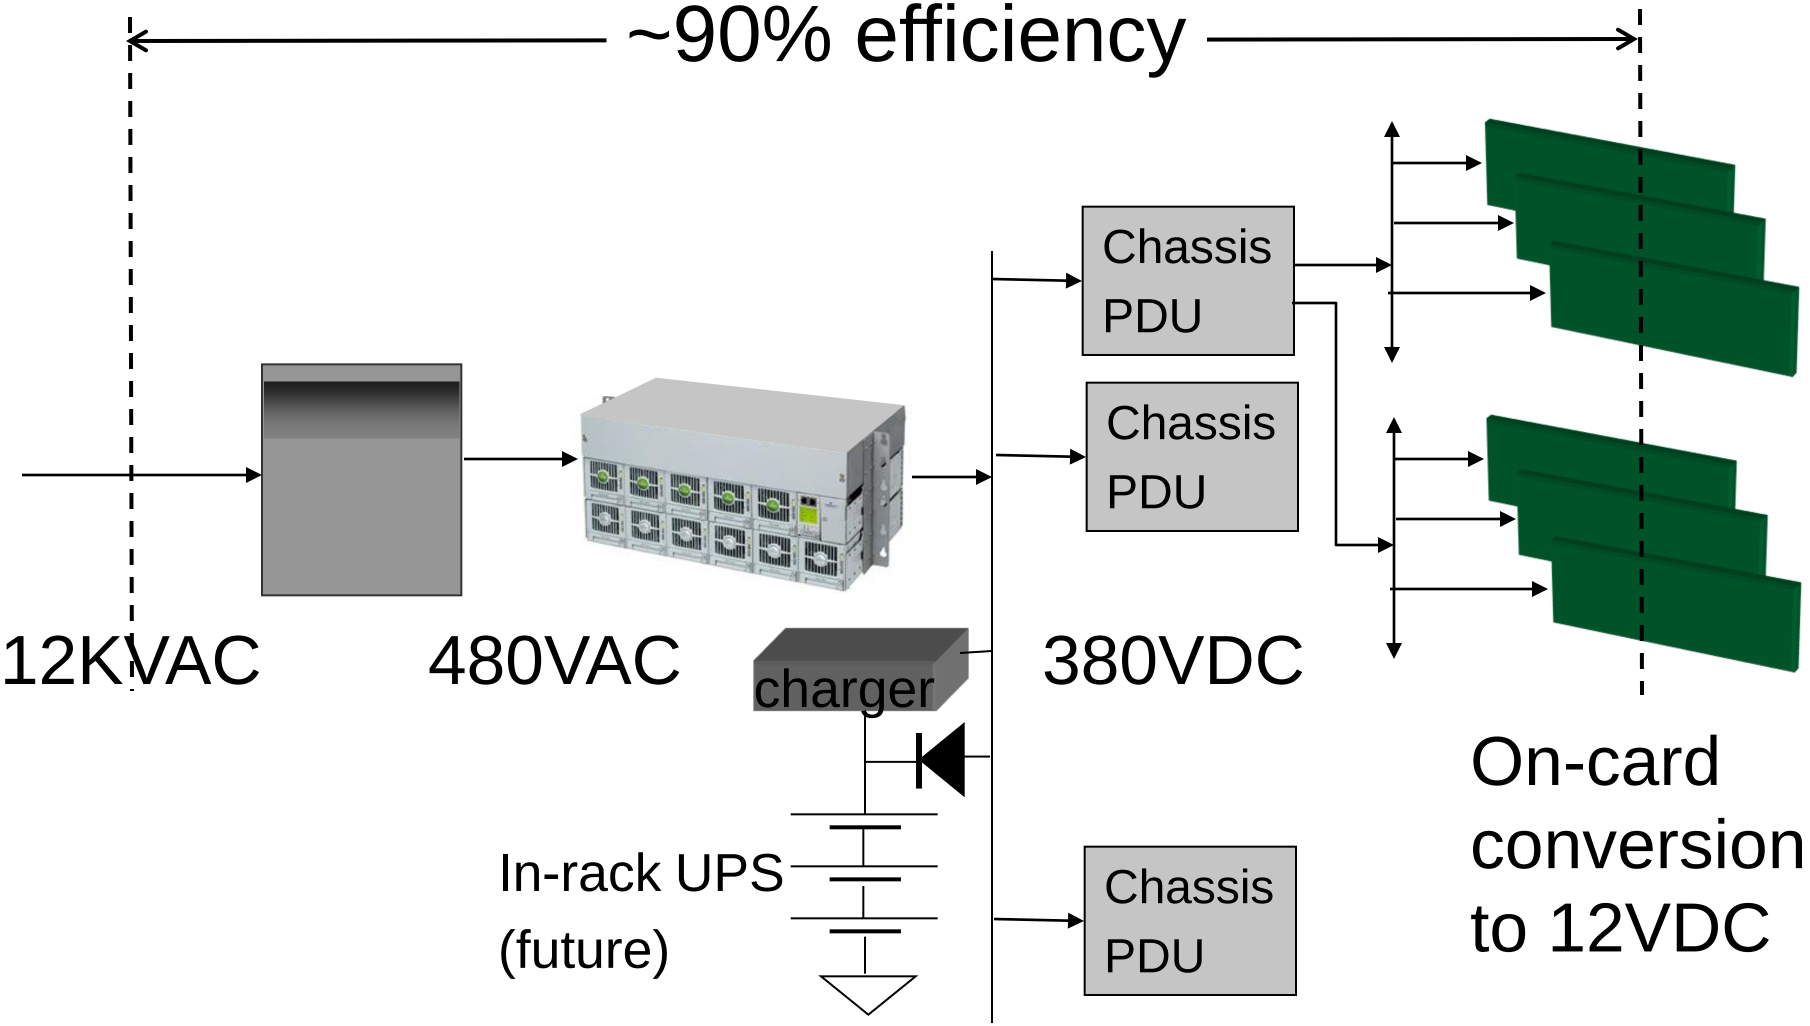
~90% efficiency
Chassis PDU
Chassis PDU
12KVAC
480VAC
charger
380VDC
In-rack UPS (future)
Chassis PDU
On-card conversion to 12VDC

System Intercommunication Architecture:

I/O virtualization

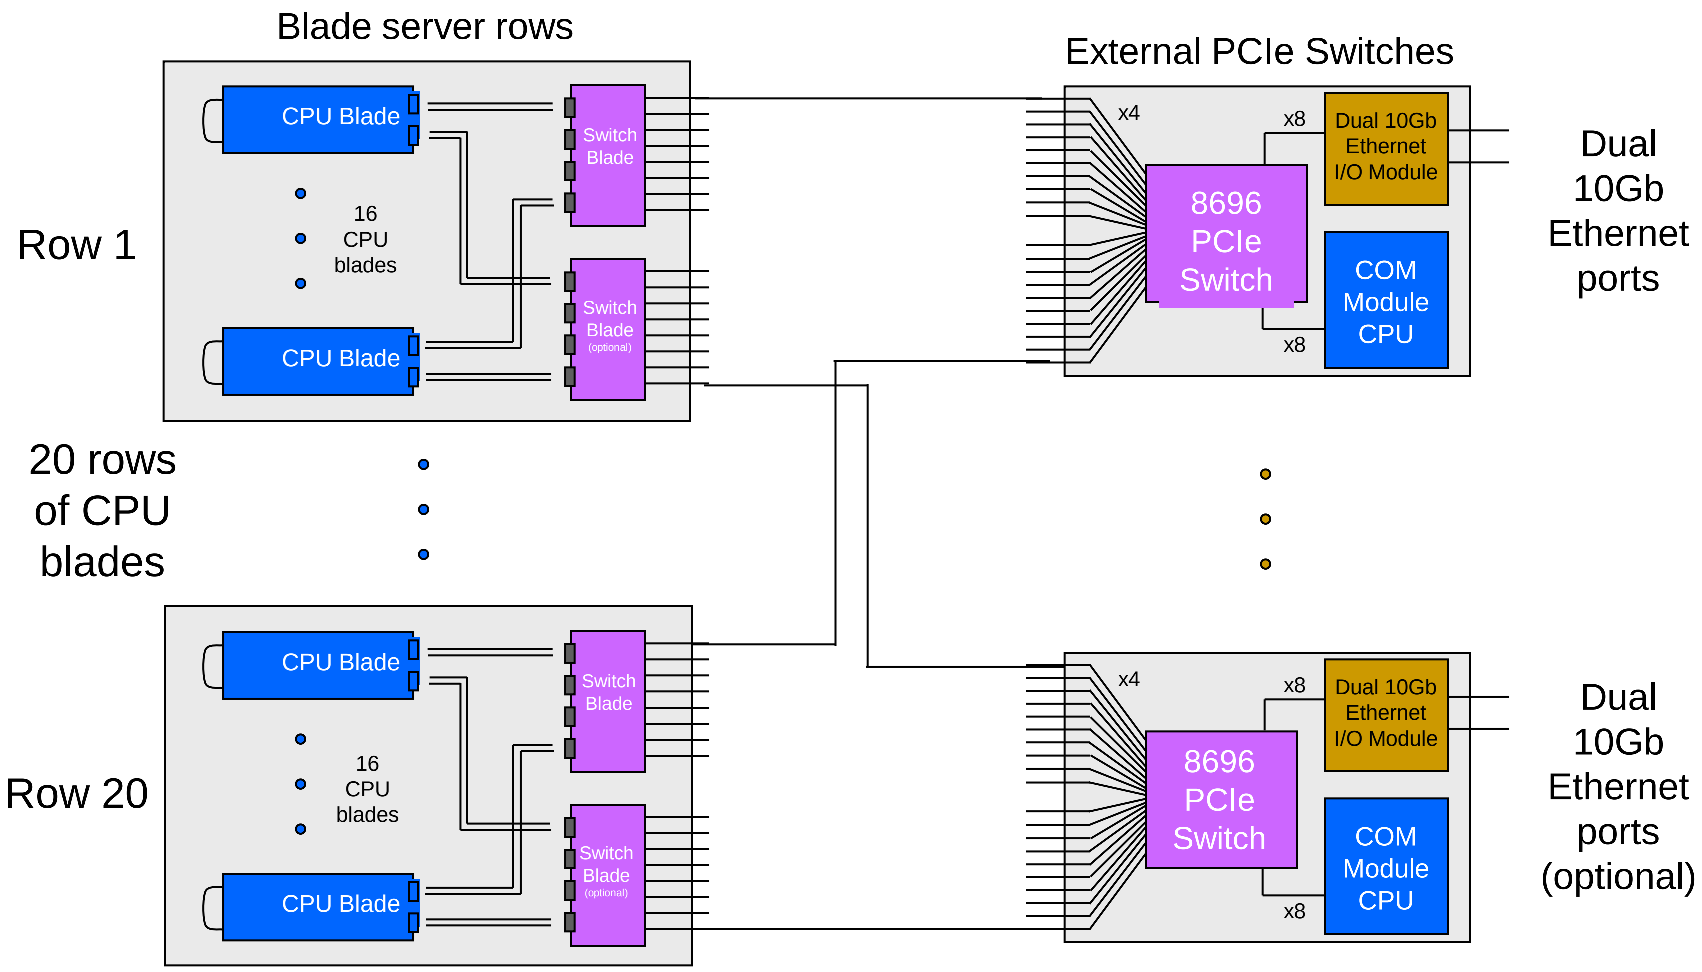

Application Example – 20’ Container

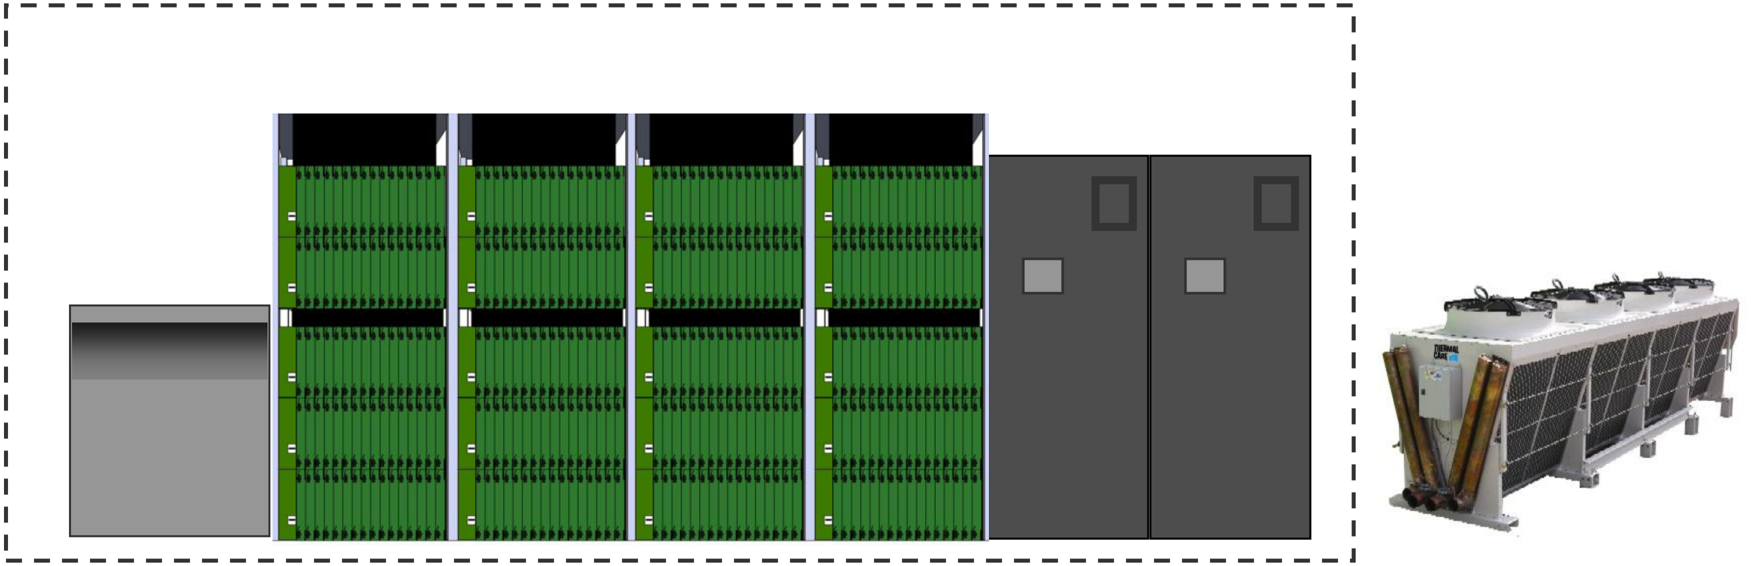

12KV-480VAC
400KVA Xfmr

CSys Racks

2 x Liebert XDP

Adiabatic cooler+
chiller (optional*)

- Capacity: 640 Servers
- 400KW power and cooling
- Infrastructure cost ~$1,000 per server

*Very hot climates

Potential Cost Savings

3.2MW Data Center

	Next-gen	**HP Ultra-modular**	**Clustered**	**% reduction**
Site and facility development	$58,000,000	$26,000,000	$ 9,207,360*	84%/65%
Total annual energy cost	$3,053,106	$2,637,677	$ 2,136,518 **	30%/19%
Total annual water cost	$67,440	$7,418	$ 6,454	90%/13%

*Estimated price for 8 container based 400KW modules

** Based upon new blade server design

Status

- PCB layouts complete, prototypes mid/end April
- Mechanical design 90% complete, assemblies deliveries start mid April
- Cold plate prototypes tested OK
- Chassis test facility designed and built
- Chassis test starts early June
- First rack completes in August

Summary

- Mechanical/cooling framework for multiple applications
- Open system
  - Mix and match switches and blades
    - PCIe, IB, Ethernet, etc.
    - CPU, GPU, Storage, Router
- All electronics field replaceable
  - Forward legacy capability
- ~$20M/PetaFLOPS all included
  - Facilities, hardware etc.
- No special facilities required
  - Direct touch cooling

Hot Water Cooling

Low device to sink (outside air) thermal resistance

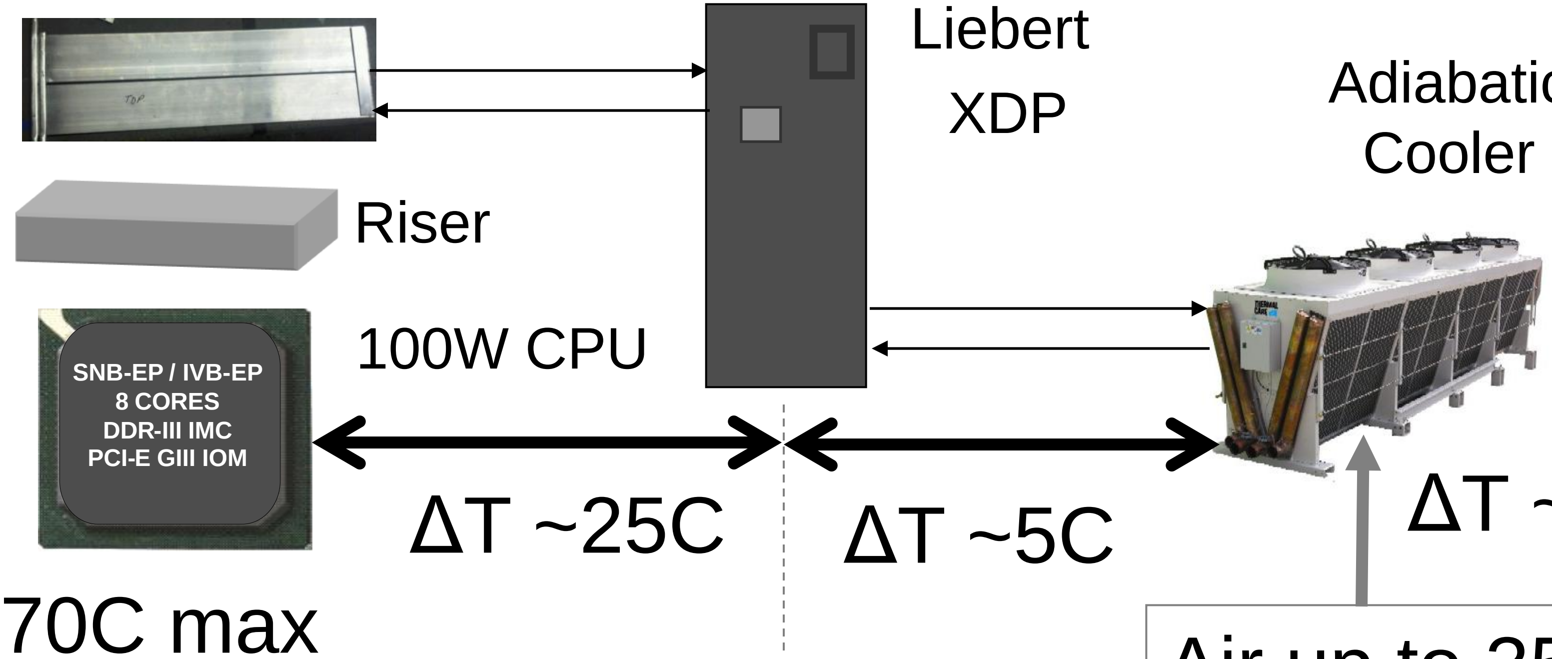

Cold Plate

Riser

100W CPU

Liebert

XDP

Adiabatic
Cooler

ΔT ~25C

ΔT ~5C

ΔT ~5C

70C max

Air up to 35C
(95F) Wet Bulb

- True PUE 1.10