Open. Together.

OCP
REGIONAL
SUMMIT

Immersion optimised Enterprise IT Platforms

Rolf Brink, CEO, Asperitas

Immersion technology

RACK & POWER

Asperitas AIC24-15/21”
Passive immersion technology
Power density:

- Unit power (rack footprint): 43 kW, 60 kW potential (2n: 3-phase 400V/63A)
- Compute density: 43-66 kW/m2@32°C
- Solution footprint: 23-33 kW/m2 @32°C (ASHRAE W3)

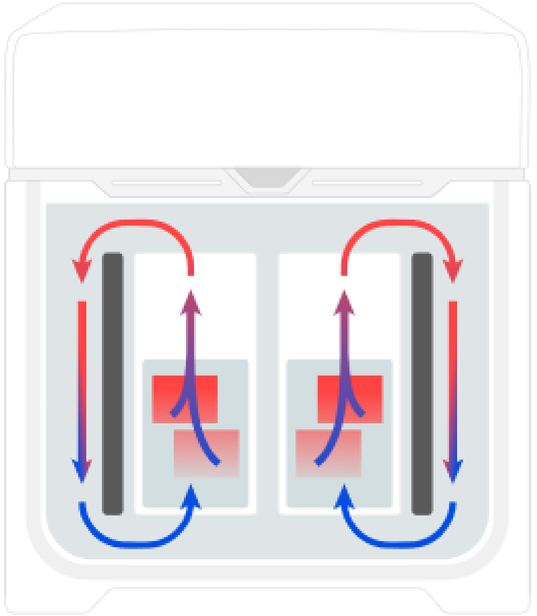

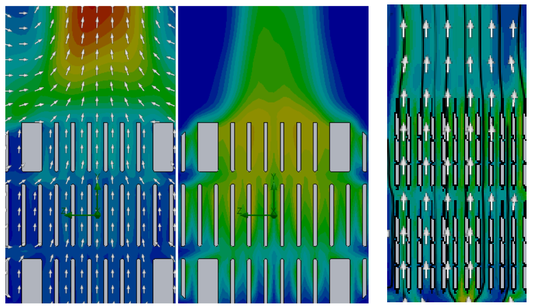

Passive immersion technology:
Common flowrates of 10-20 mm/s

Only *heated* liquid is circulated

Solution integration

Electrical
- 43 kW power delivery
- Rack level switchgear
- Power management

Cooling
- Direct connection to FWS
- Flow control
- Thermal and pressure monitoring

Management
- Monitoring
  (FWS, DECS and IT thermal, Power)
- Alarming
- Logging (including IT thermal information)

Containment
- Dual hull
- Liquid level monitoring
- Cable management (capillary action)
- Thermal (insulated)

Serviceability
- Redundant switching
- Automatic lifting (no manual handling)
- Containment (leak trays)

IT optimisation
- Liquid design servers (based on air design)
- Full material compatibility research
- Full performance research
- Fully tested

High Availability

RACK & POWER

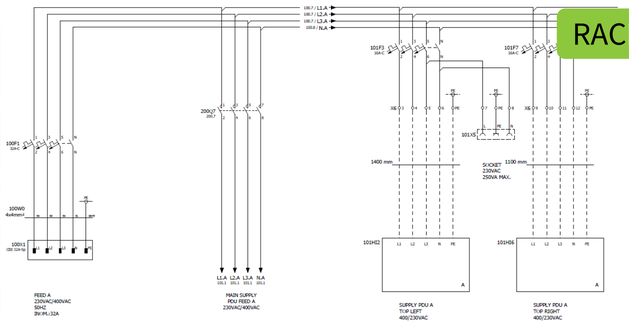

Electrical Selectivity
- Dual power
- Multiple sensors for power safety
- Full selectivity on monitoring/autonomous safety/DCIM

Concurrent maintenance
- Redundant convection drives®
- Independent thermal management
- Exchangeable during operation

Availability optimisation
- High thermal buffer (ride through)
- Autonomous safety (early reporting)
- Integrated cable management (easy servicing)

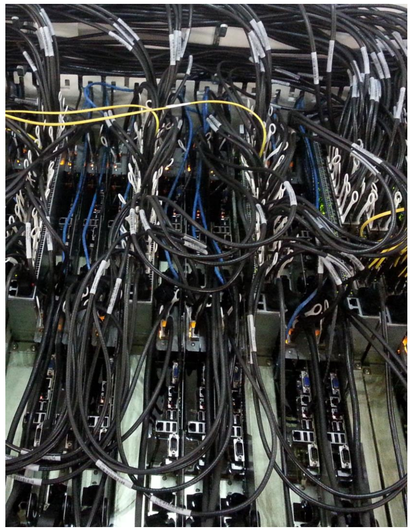

Optimisation

Thermal optimization and high safety

RACK & POWER

	Sensors (reporting)	Control	Auto safety
Power	Amps, Voltage, VA, W (Feed & per outlet)	On/off	On/off
FWS	Temperature (in/out) Flow & Pressure	Control and safety Valves	Open/close Non-return
Dielectric	Temperature (in/out) Volume Quality	Control	

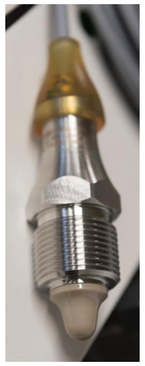

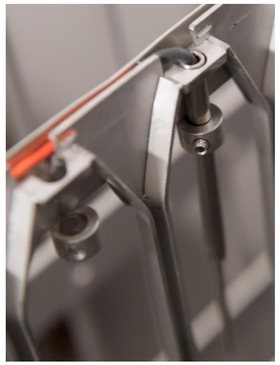

New industry developments

RACK & POWER

ASHRAE focus on liquid datacenters

New fundamental concepts being defined

- Temperature chaining/energy cascade
- Dielectric ride through
- Facility ride through
- Availability measures and metrics

Examples being put into practice

- CoolDC, Bytesnet, Asperitas Lageweide project (DCF presentation)

Ride through example

RACK & POWER

Full analysis per request, contributed within ASHRAE TC9.09

Start:

- 10+ kW
- 42°C cooling
- Ridethrough 1h13m

#	System	Chipsets	Tolerance	Condition
1	GPU-12	12x AMD GPU 1x broadwell	36°C	Stress load
2	GPU-8	12x NVIDIA GPU 1x broadwell	36°C	Stress load
1	Intel CPU	4x Intel Skylake	38°C	45% load
1	AMD GPU	2x AMD EPYC 2x Nvidia Titan V	44°C	Stress load
1	Broadwell	2x Intel Broadwell	46°C	Stress load
1	HDC AMD	4x AMD Epyc	46°C	Stress load
1	AMD CPU	2x AMD Epyc	48°C	Stress load
1	Management	2x Intel Broadwell	55°C	Idle
1	CFD	2x Intel Broadwell 1x NVIDIA GPU	50°C	Idle
10	Demo Single	2x Broadwell	55°C	Stress load
4	Demo Twin	4x Broadwell	55°C	Stress load

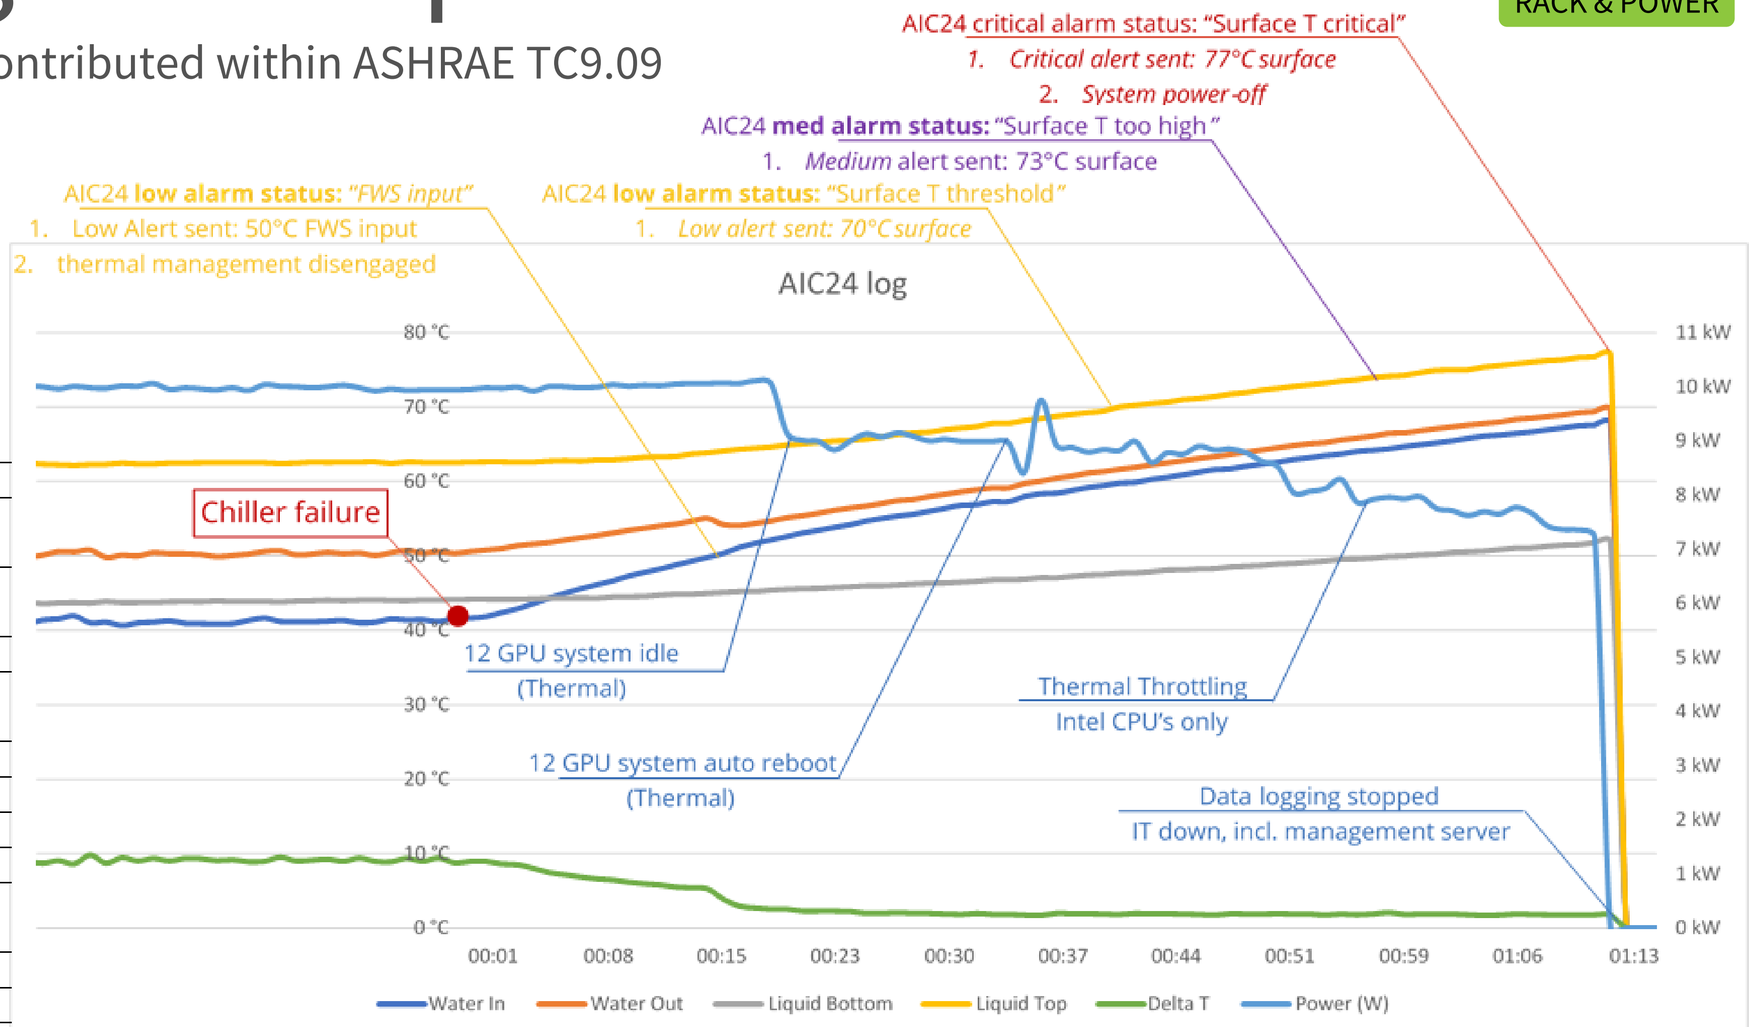

IT/Platform design

The platform availability, then IT are the highest value within immersion technology.
Not all IT components can be immersed safely!

No assumptions **regarding IT compatibility or safety**

- Material compatibility is a serious concern
- Material compatibility can be dealt with
- Material compatibility must be researched

Thermal optimisation *will* increase IT availability and lifetime

- Thermal design is relevant for liquid flow (liquid is different from air!)
- Reduced failures (material, thermal and electrical)
- Increased thermal stability/reduced thermal shock
- Designing IT gear for immersion adds new possibilities (space optimisation)

Example: EPYC 7351: Stable at 318W@3,7 GHz /40°C+ cooling

Optimised chassis development

RACK & POWER

Universal chassis

- 15” or 21”
- Guidance rails
- Optimised for liquid flow
- Dry interfaces
- Flexible interfaces

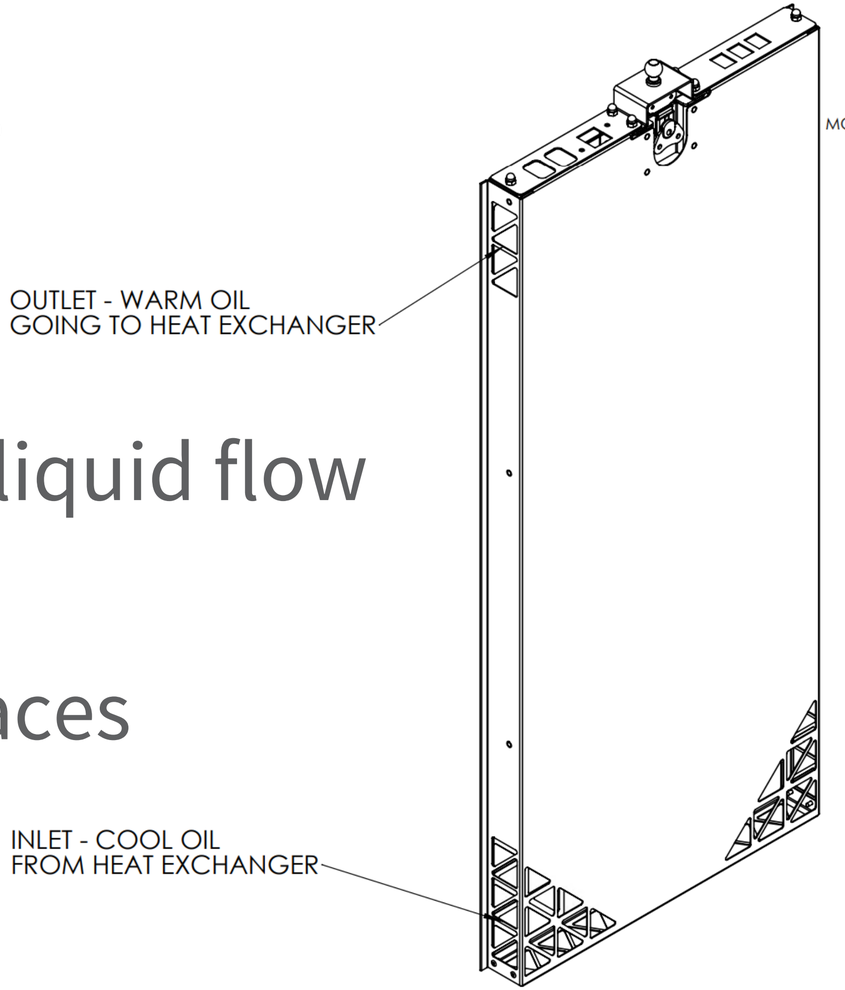

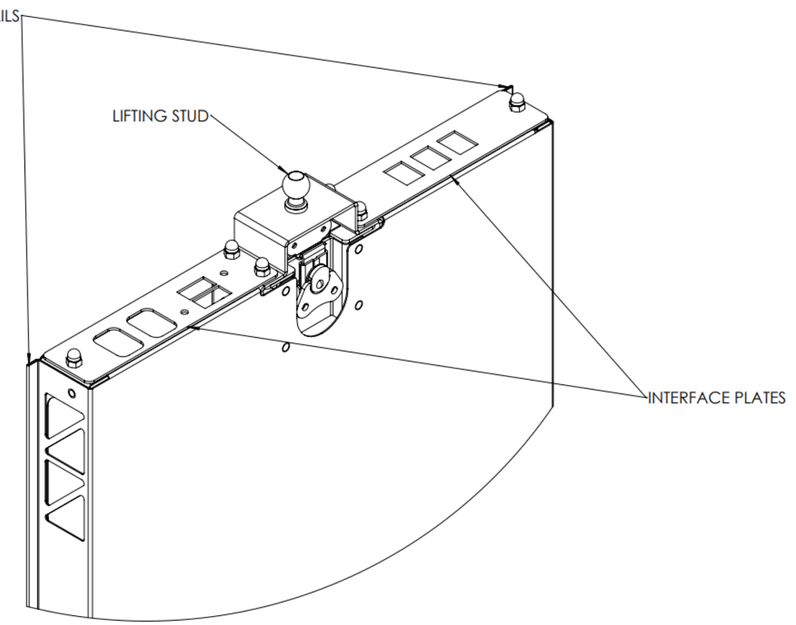

Platform design - Validation

RACK & POWER

1 - Platform requirements:
- CPU, GPU or other workload types
- Storage, network or other interface requirements
- Power density/thermal requirements

2 - IT selection
- Mainboards, CPU/GPU, network, storage etc.
- Brand/partner selection

3 - Concept mechanical design
- Existing assembly vs optimised (depending on requirements)
- 15”/21”, U/OU, component lay-out, accessibility PSU’s (previous OCP sessions)

4 - Thermal analysis for power density
- Ease of flow
- Flowrate
- Temperature layering
- Component temperatures

Platform design - One-off evaluation

RACK & POWER

1 - Engineering samples from manufacturers based on optimised chassis
If not assembled in optimised chassis:

- Transfer components to optimised chassis
- Build prototype system with components only
- Evaluate material compatibility

2 - 48 hour max load test

- Analysis power benchmark, thermal performance
- Evaluate material compatibility

3 - Product documentation

- Material analysis report
- Test report
- Thermal analysis (corrected for actual measurements)
- Build specification

Platform design - duration

RACK & POWER

1 - 10 week duration test (simulate extensive lifetime usage)
- High temperature fluctuations
- Power testing
- Multiple servicing
- Multiple full (dis)assemblies
- Functional testing

2 - Product documentation
- Material analysis report
- Test report
- Thermal stability and tolerance
- Build specification

Platform developments based on OCP

RACK & POWER

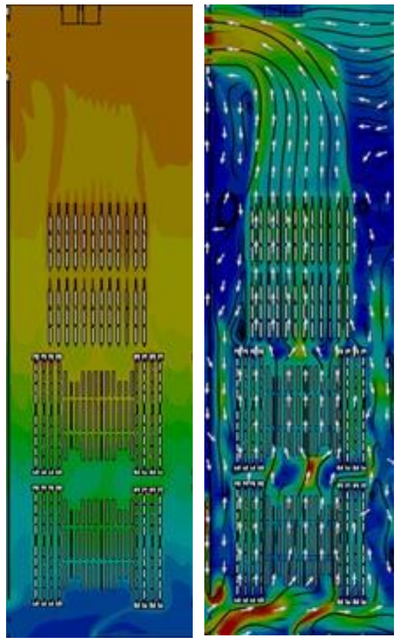

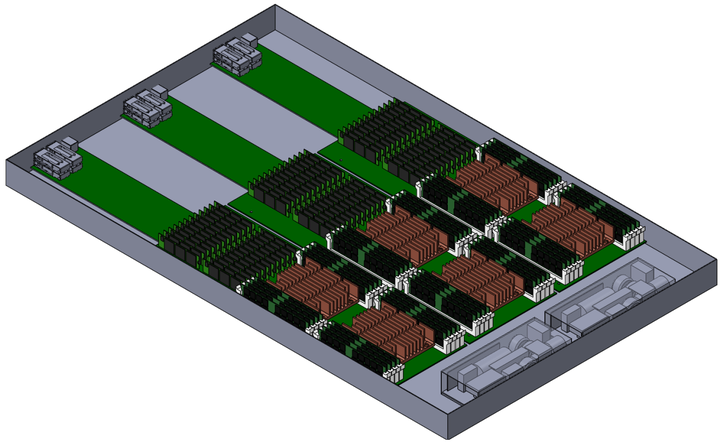

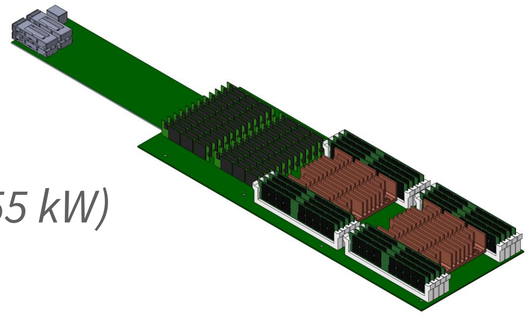

HDC (High Density CPU)
- 96x Intel Cascade lake, 48 nodes
- 96x AMD ROME, 48 nodes
- Upcoming: 144x Cooper lake/ROME, 72 nodes

GPU
- Current: 288x AMD GPU, 24x Intel/AMD CPU, 24 nodes
- Current: 192x NVIDIA GPU, 24x Intel/AMD CPU, 24 nodes

AI
- Current: 96 NVIDIA GPU, 48x Intel CPU, 24 nodes
- *Upcoming: 192x NVIDIA GPU, 48 nodes*

Storage
- Current: 192x NVm internal storage (8 per node)
- *Design concepts: **1728x/2304x NVm (27/35 PB)**, 72 node storage system (40/55 kW)*

Climate independent high temperature cooling: ASHRAE W5+

Liquid optimisation projects

RACK & POWER

Immersion ready optical solutions
- 10GTek (Q)SFP extender optimisation (development partnership)
- Leoni immersible QSFP optical cabling (development partnership)

Immersion liquids
- Shell immersion cooling fluids (development partnership)
  Optimisation for all immersion types (convection, mechanical and enclosed systems)

IT equipment
- Wiwynn, SuperMicro, Dell, Penguin: Server optimization (development partnerships/sponsored)
- {undisclosed}: Storage optimization (development partnerships open/being established)
- Intel and AMD: CPU evaluation and thermal optimization (Material contributions/sponsored)

Call to Action

RACK & POWER

- Contribute to ACS Immersion
- Use/Co-develop optimized IT platforms for immersion
- More information: Asperitas.com/Rolf.Brink@Asperitas.com

ACS Immersion requirements:
https://www.opencompute.org/documents/ocp-acs-immersion-requirements-specification-1-pdf

ACS Immersion schedule:

Each 3rd Tuesday of the month

https://www.opencompute.org/wiki/Rack_%26_Power/Advanced_Cooling_Solutions

Mailing list: http://lists.opencompute.org/mailman/listinfo/opencompute-acsimmersion

OCP REGIONAL SUMMIT

Open. Together.

OCP Regional Summit

26–27, September, 2019