Developing Technologies in Server Design

Stephen W. Montgomery, Ph.D.
Sr. Systems Engineer

Intel Labs – DuPont, WA

May 7, 2002

Agenda

- Current challenges
- Integrated Rack Cooling and Docking Station (IRCDS)
- Potential applications
- Summary

Current Challenges

- Cooling performance processors in highest density systems
  - Capacity reduced → height constraints
  - Limited redundancy
- High density in data center
  - Potential air recirculation issues
  - Difficult to scale CRAC units
  - Extensive cabling leads to flow blockage
- Serviceability
  - Multiple cable connections
  - Rack must be rear-accessible
    - Greater service time
    - Greater chance of human error

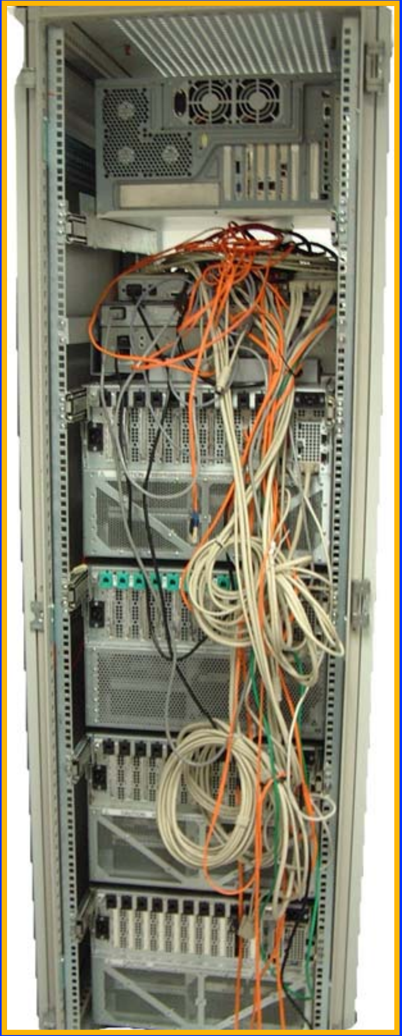

IRCDS

- Integrated Rack Cooling and Docking Station (IRCDS)
  - Electrical → Rack Docking Station (RDS)
  - Thermal → Rack Cooling Station (RCS)
- Goals
  - Front serviced, rapid “swappability” (blade-like)
  - Solve system and data center thermal and infrastructure issues

Concept Overview

- ‘Pluggable’ interfaces
  - Electrical → RDS
  - Thermal → RCS
- Hot swap capable
- Front access only
- Remote thermal management

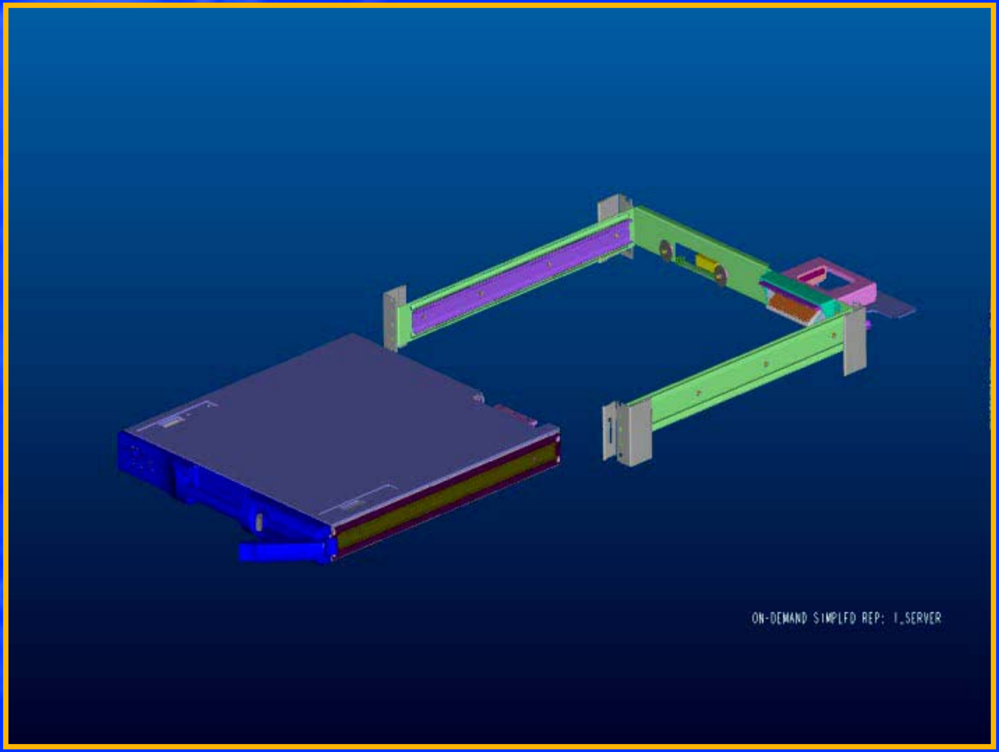

Intel Labs

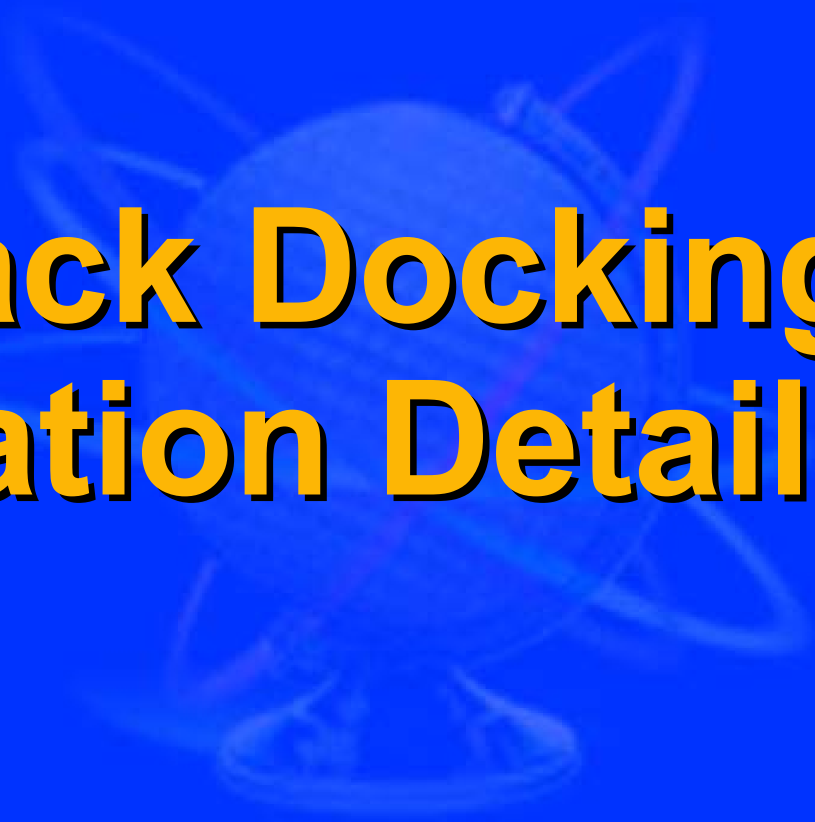

Rack Docking Station Details

Rack Docking Station

- What is it?
- Why?
- Challenges
  - Server engagement
  - EMI containment
  - Signal integrity
- Cost Savings

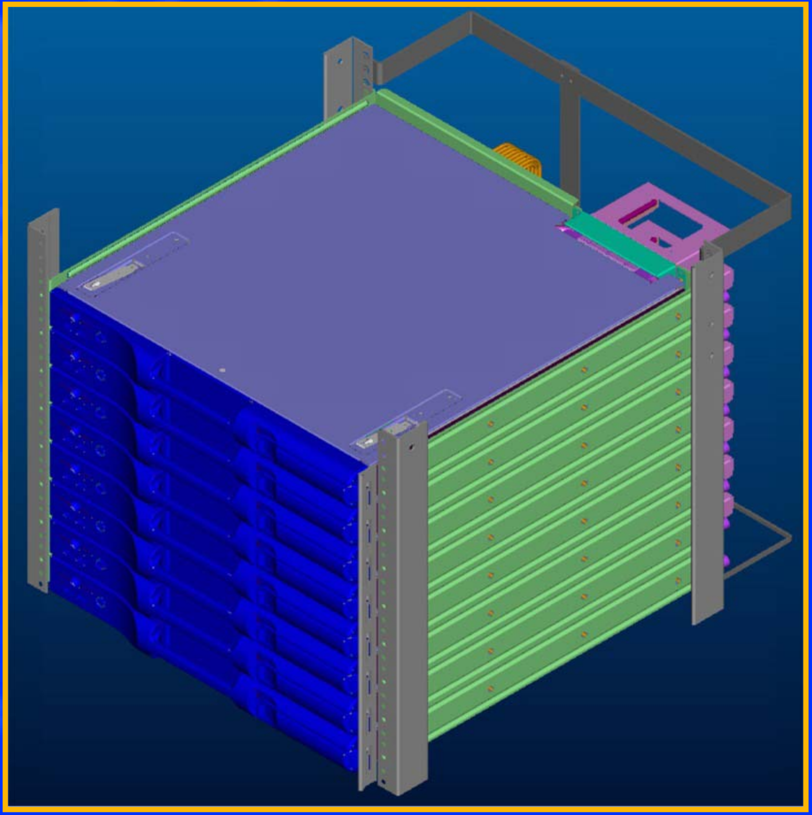

Rack Docking Station (RDS)

RDS is…

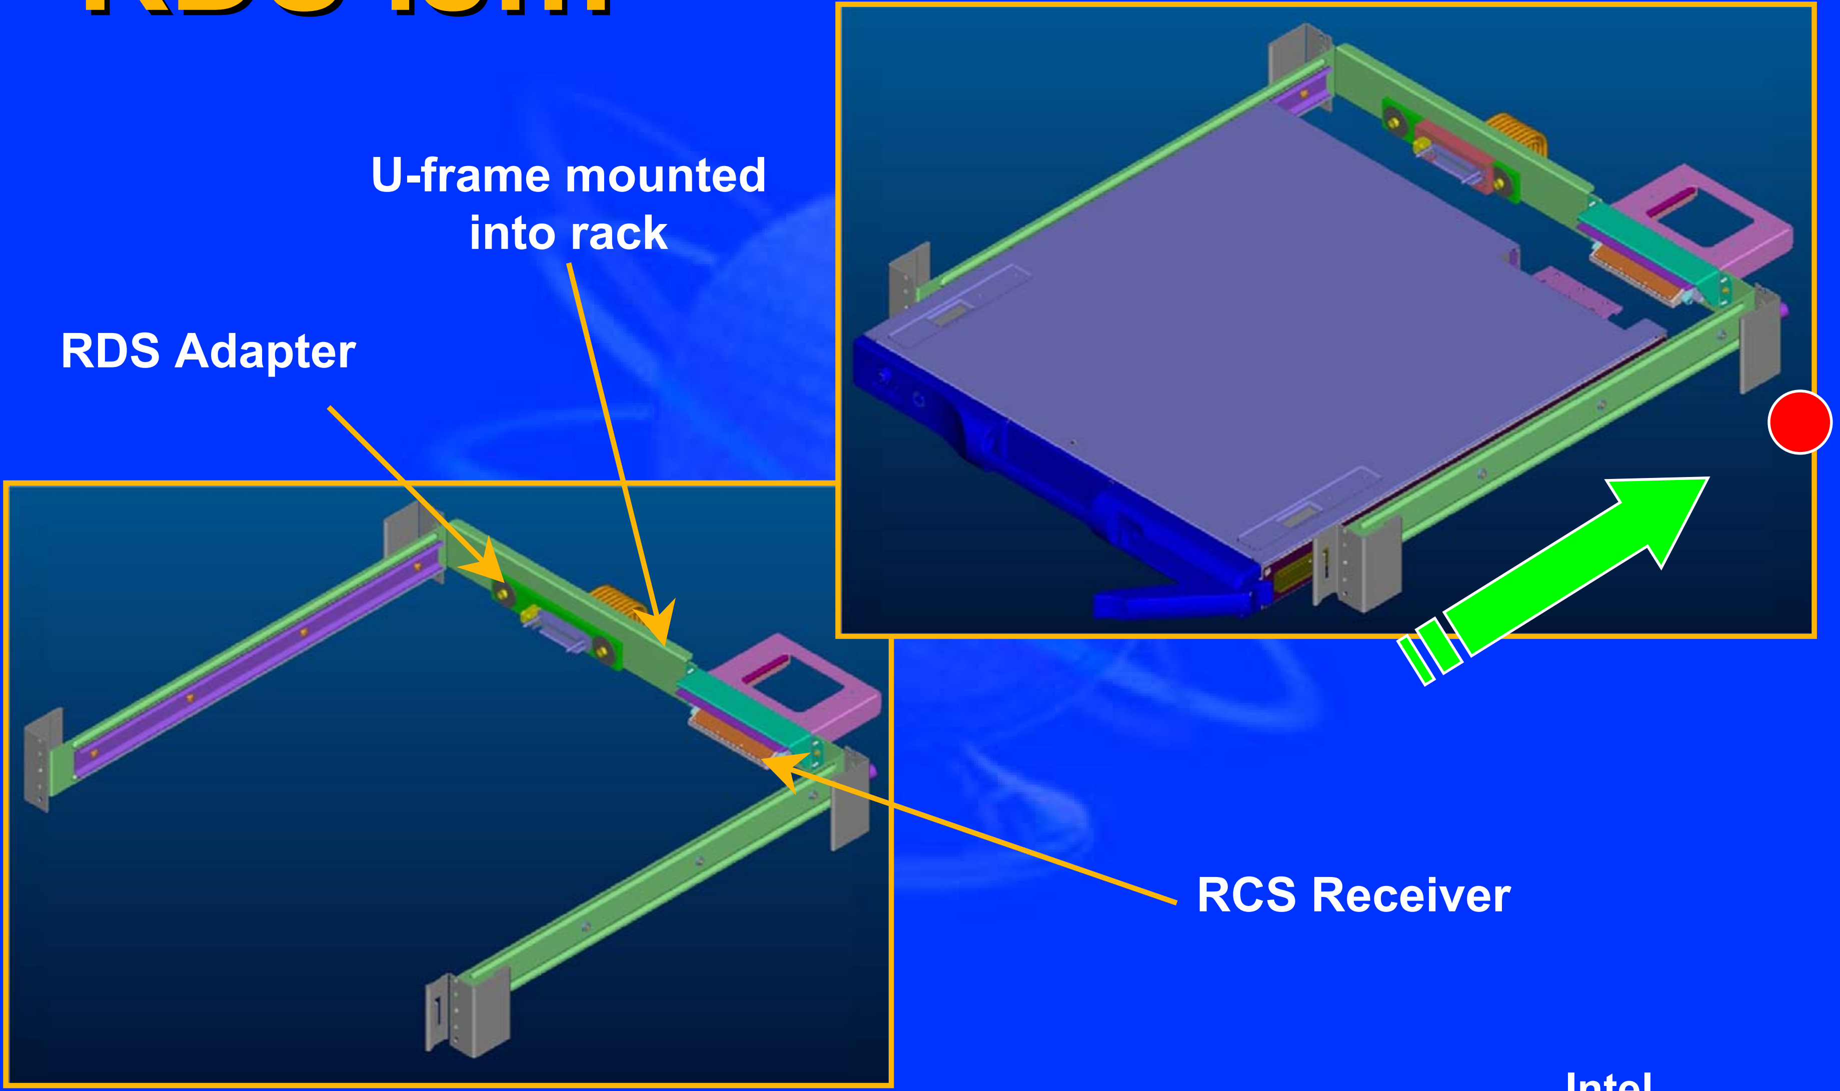

U-frame mounted into rack

RDS Adapter

RCS Receiver

Why RDS?

- Solution to “rack-level” problems
  - 👍 Cable management
    - Rapid deployment of rack components
    - One-time cable assembly
  - 👍 Facilitates rapid, error-free swapping
    - Eliminates multiple cable connections
    - Server accessed via front only
    - Enables full-performance servers, w/ serviceability of blades

Server Engagement

- Assembly tolerances ~0.08”
- RDS adapter board moves in X-Y, not Z

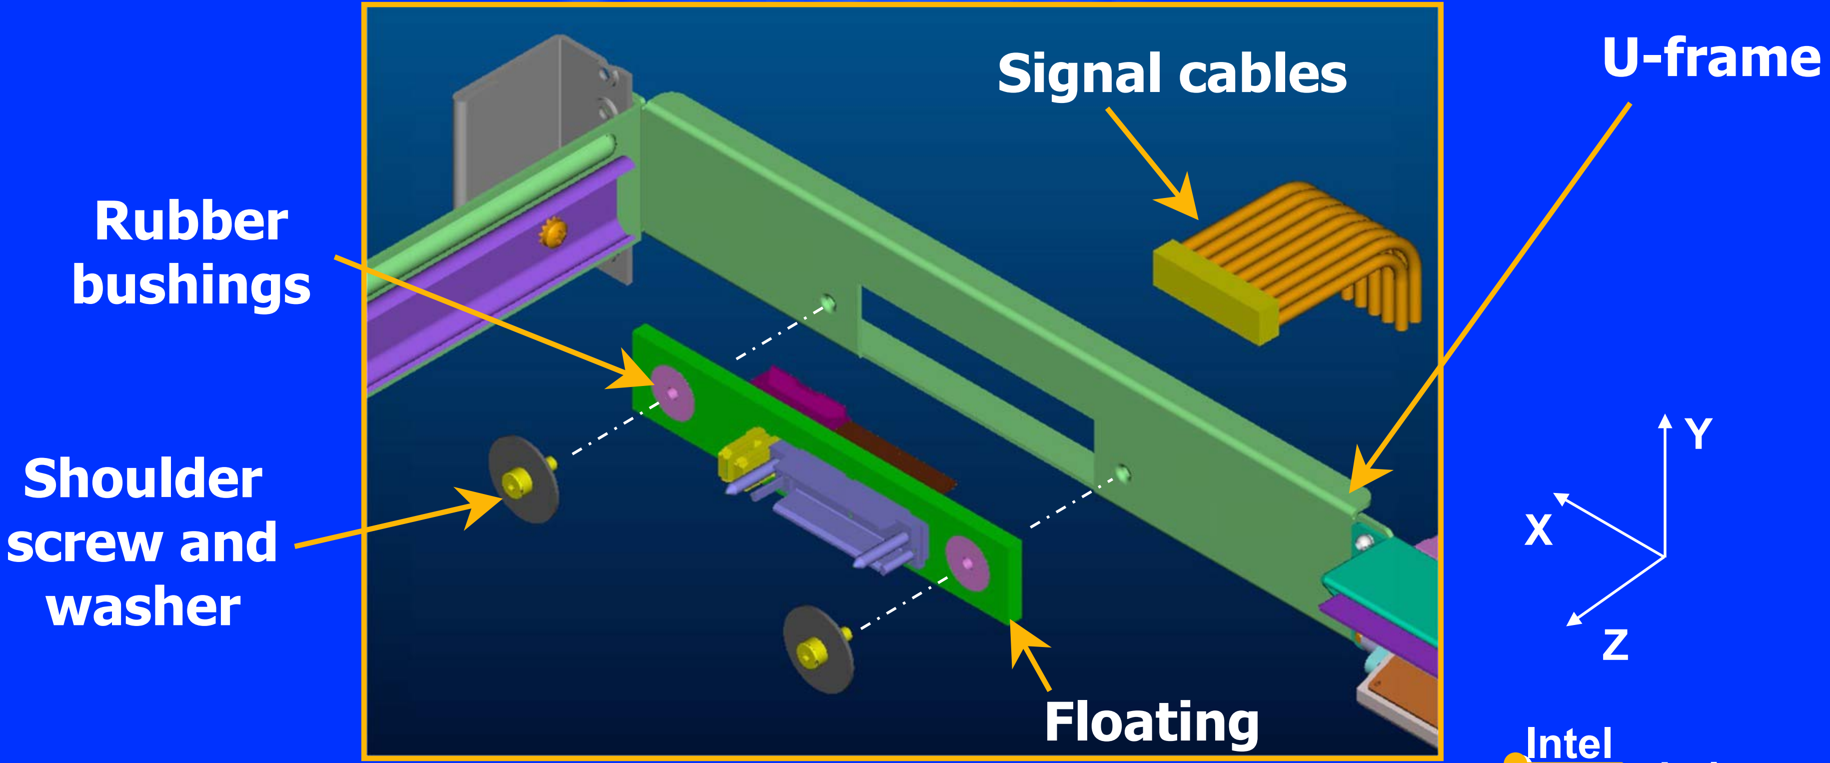

Server Engagement

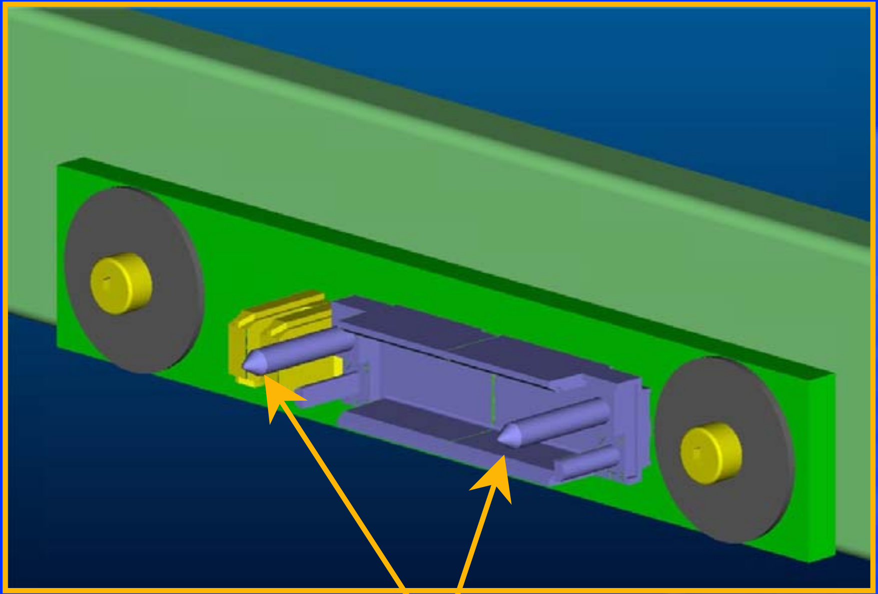

- Gathering capability = ± .102”
- Sufficient to overcome assembly tolerances of ± .08”

Alignment pins

Oversized chamfered holes

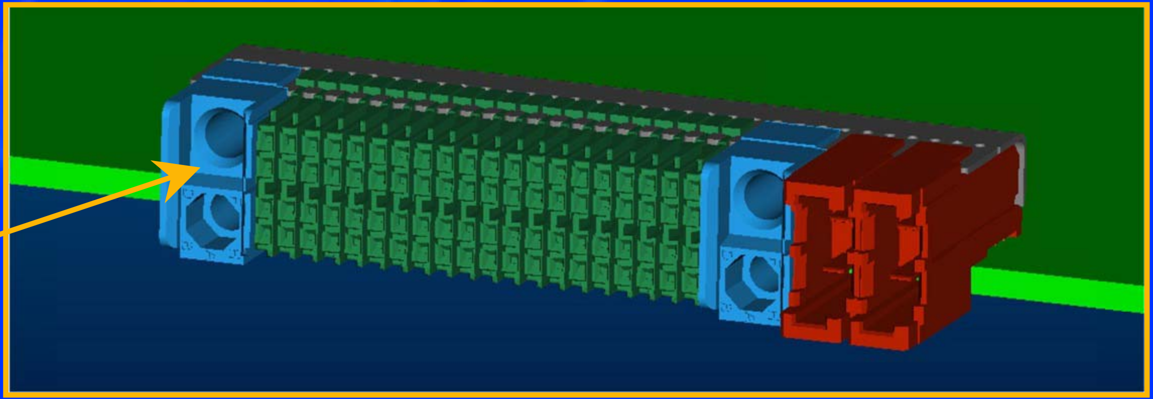

Server Engagement

- Combined insertion force ~ 20 lbs
- Implement mechanical advantage to ease assembly

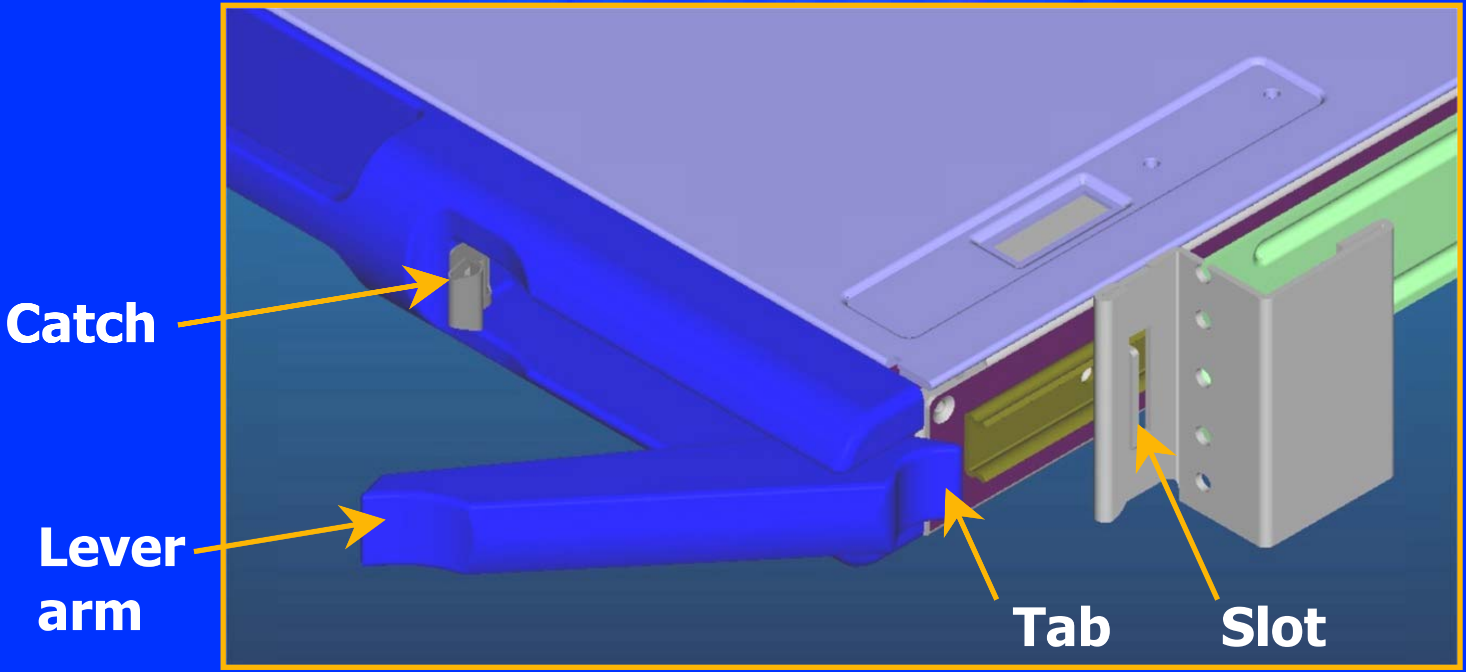

EMI Containment

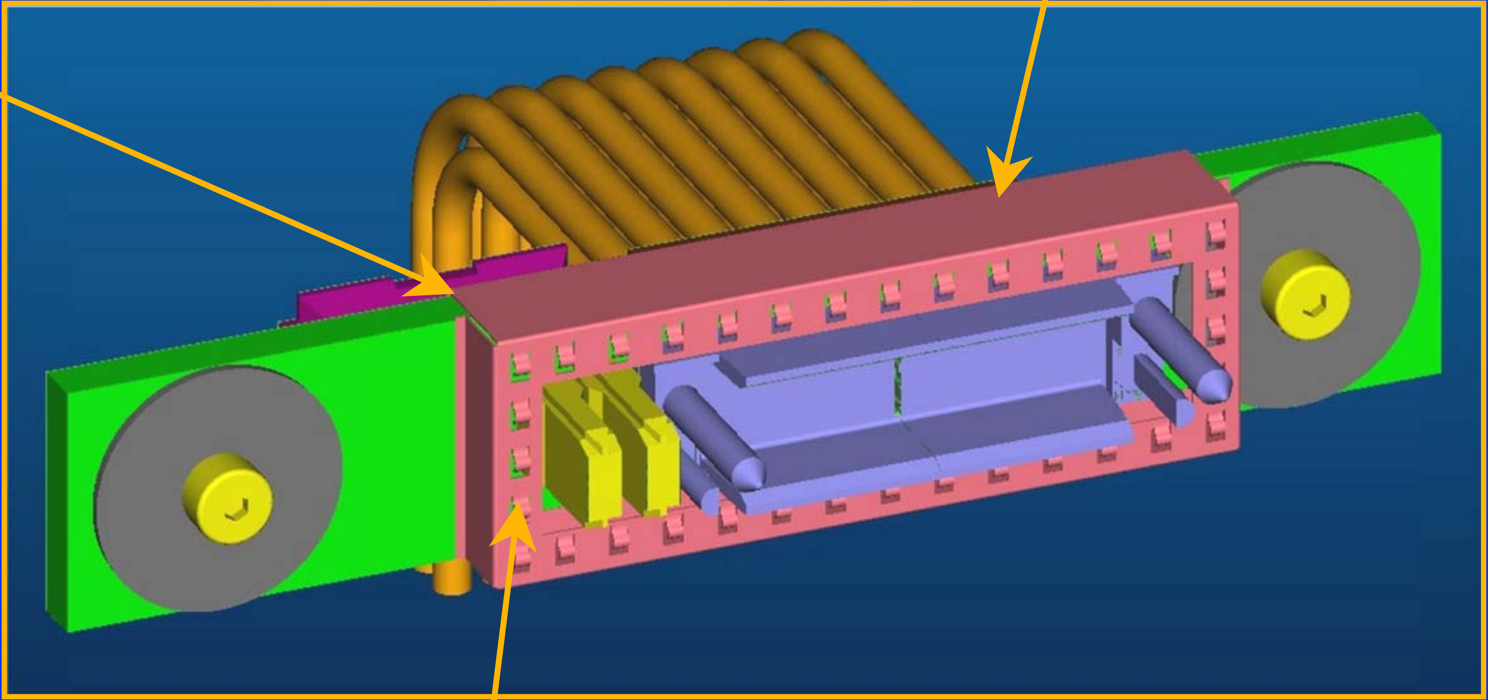

Grounding and shielding frame

Attached to adapter board ground plane

Spring fingers (contact chassis sheet metal)

Signal Integrity

- I/O impact
  - InfiniBand* minimum eye
    - V_{RSense} = 175mV
    - T_{Reye} = 140ps
- Margin exists at 5Gb/s and 2.5m cables

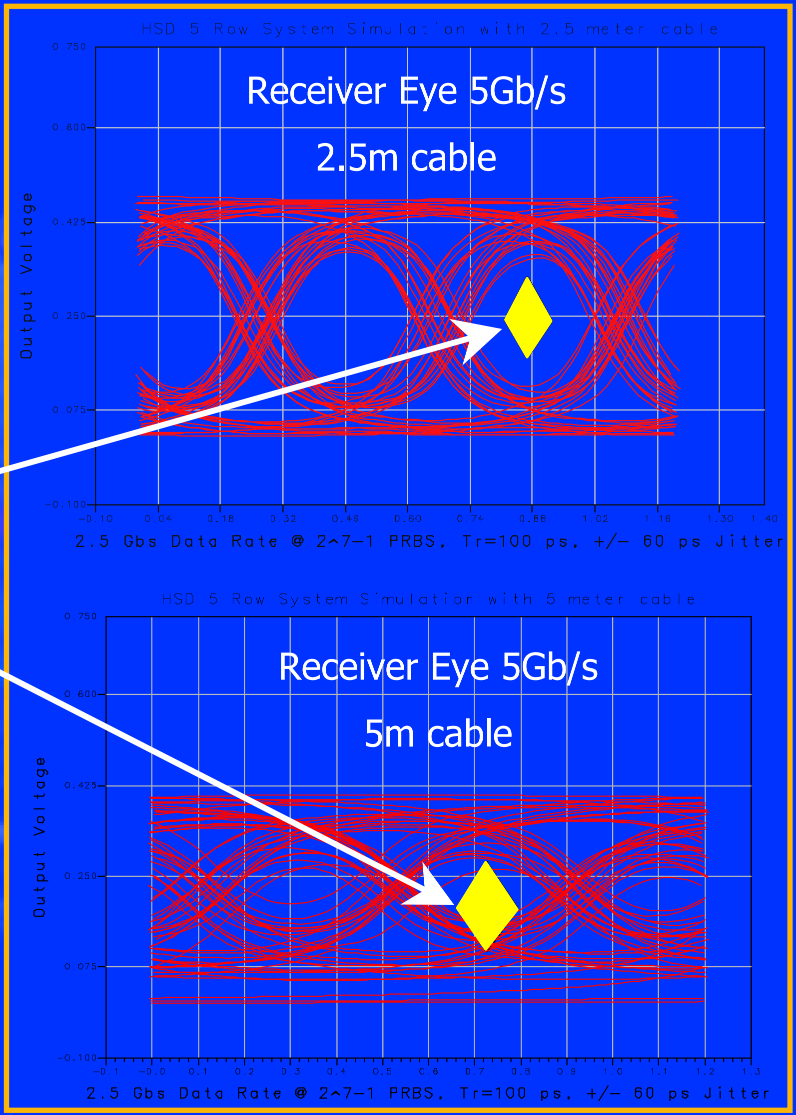

Analysis courtesy of Molex Inc.

Intel Labs

Cost Savings

- Increased serviceability
- Reduced integration time
- Standardized BOM for interconnect
- Leads to…

Reduced total cost of ownership

Rack Cooling Station Details

Rack Cooling Station

- What is it?
- Why RCS?
- Goals
- Operation
- Components
- Analysis
- Applications

RCS is…

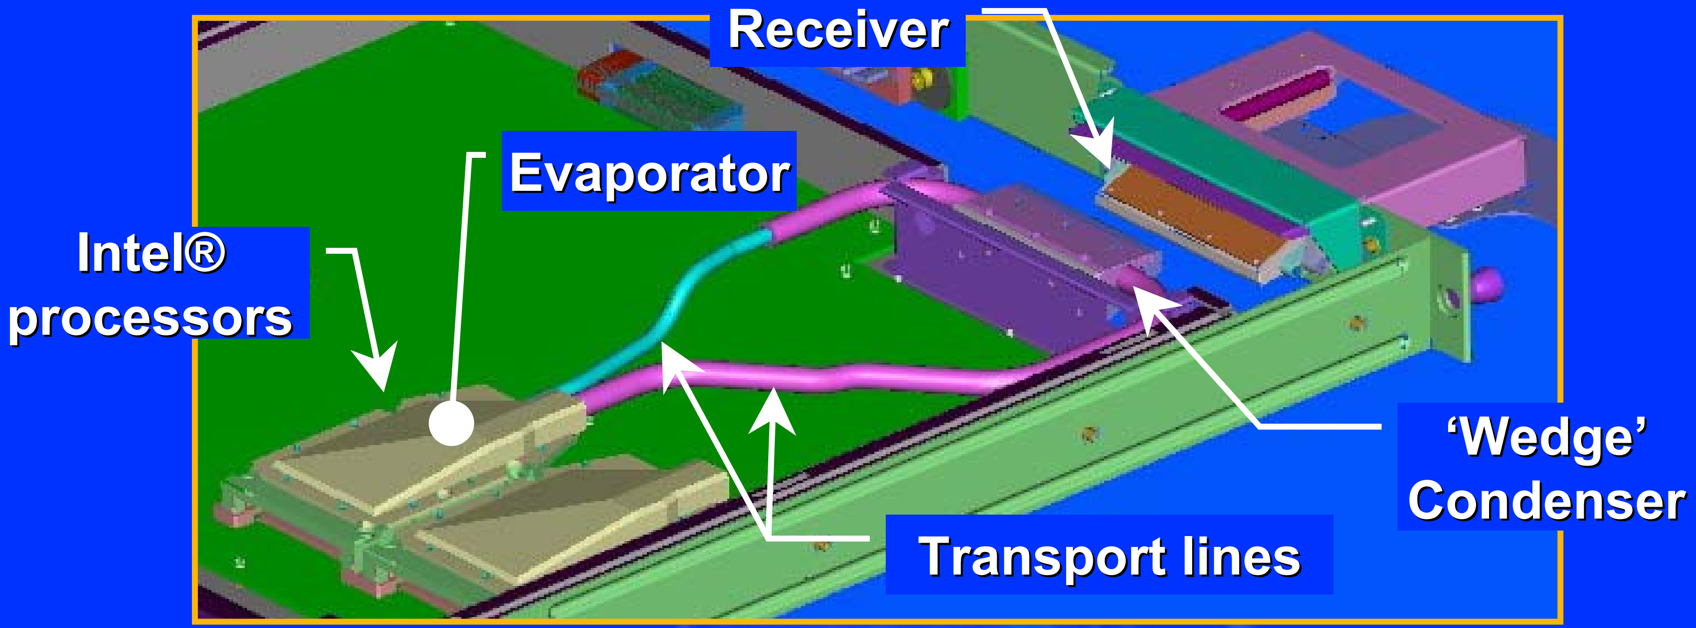

- Remote cooling via a 'pluggable' interface
  - Liquid/Vapor phase-change technology
  - Sinks heat to system on or out of rack
  - Supports high heat loads

Why RCS?

- Current chassis small and cramped
  - Limited room for fans or heat sinks
- RCS moves heat out of rack
  - Passive cooling inside/active outside
- Potential use beyond current roadmap
- Helps optimize board design
  - Can place processors for max electrical benefit

Intel Labs

Goals

- Adaptable to increased densities
- Low thermal resistance
- Passive cooling only inside chassis
- Low insertion force (≈18 N)
    - Complement RDS
- Allow tolerance ‘takeup’

Operational schematic

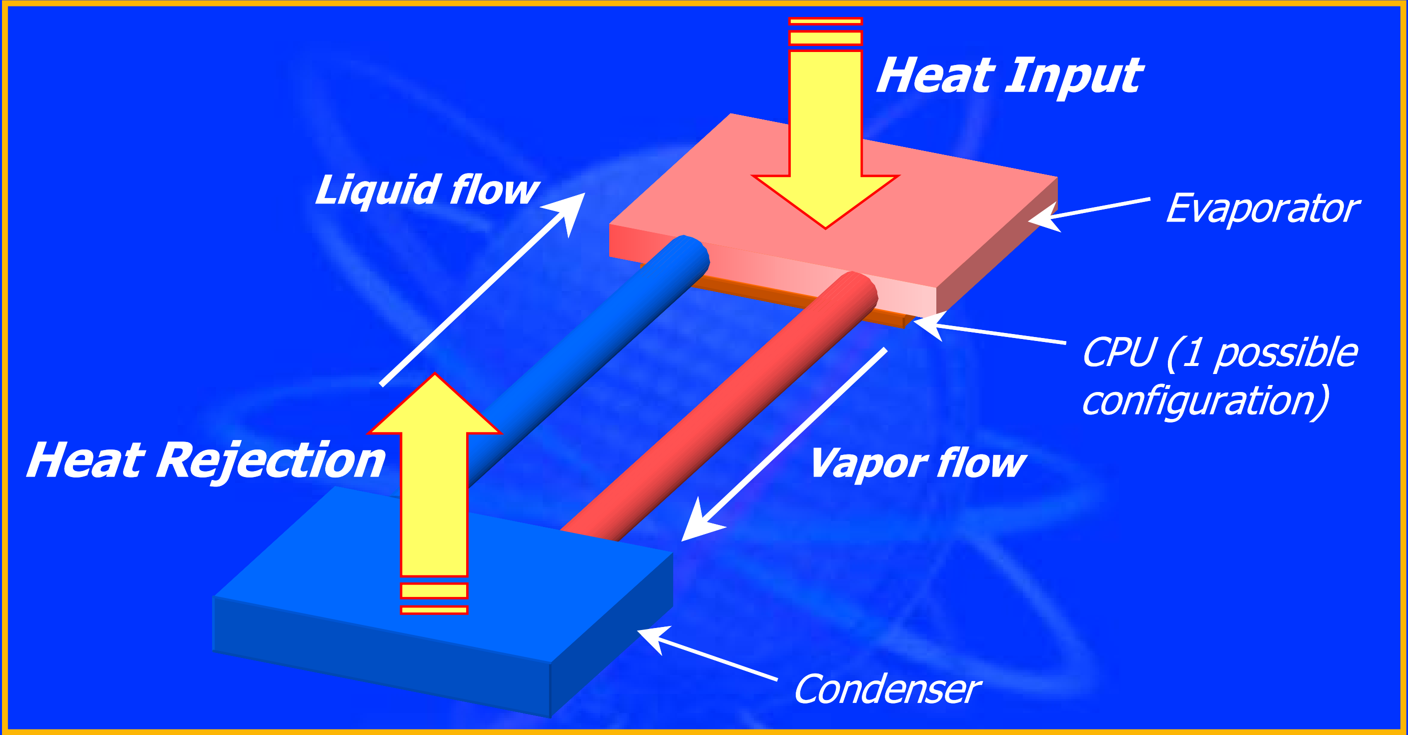

- Very reliable – in use in avionics, satellites
- H_2O = very high latent heat of vaporization
- Long-distance heat transport, w/low resistance

Components

- Modified Capillary Pumped Loop (CPL)
  - No moving parts
  - Existing mounting hardware
  - Sintered copper wick
  - Wick only in evaporator module
  - 3.175 cm head pressure available

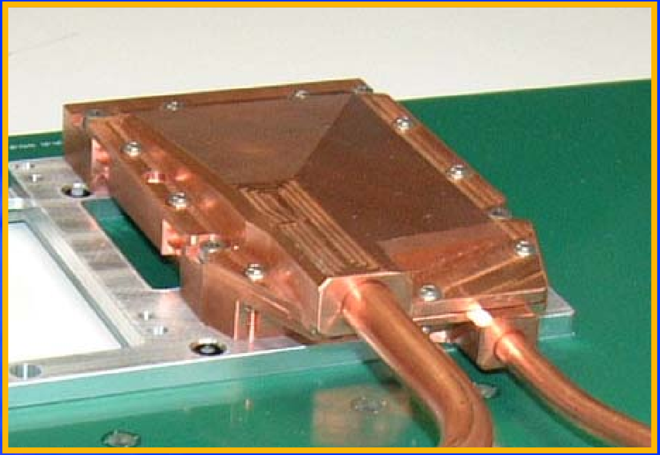

Evaporator module

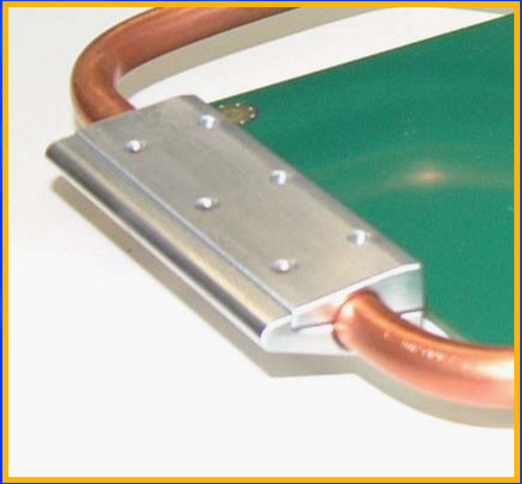

Wedge condenser

Intel Labs

Components

- Receiver module
  - Metal-to-metal contact w/wedge
  - Thermal interface material optional
  - Spring loaded
  - Takes up mechanical tolerances
  - Is mated to external cooling system

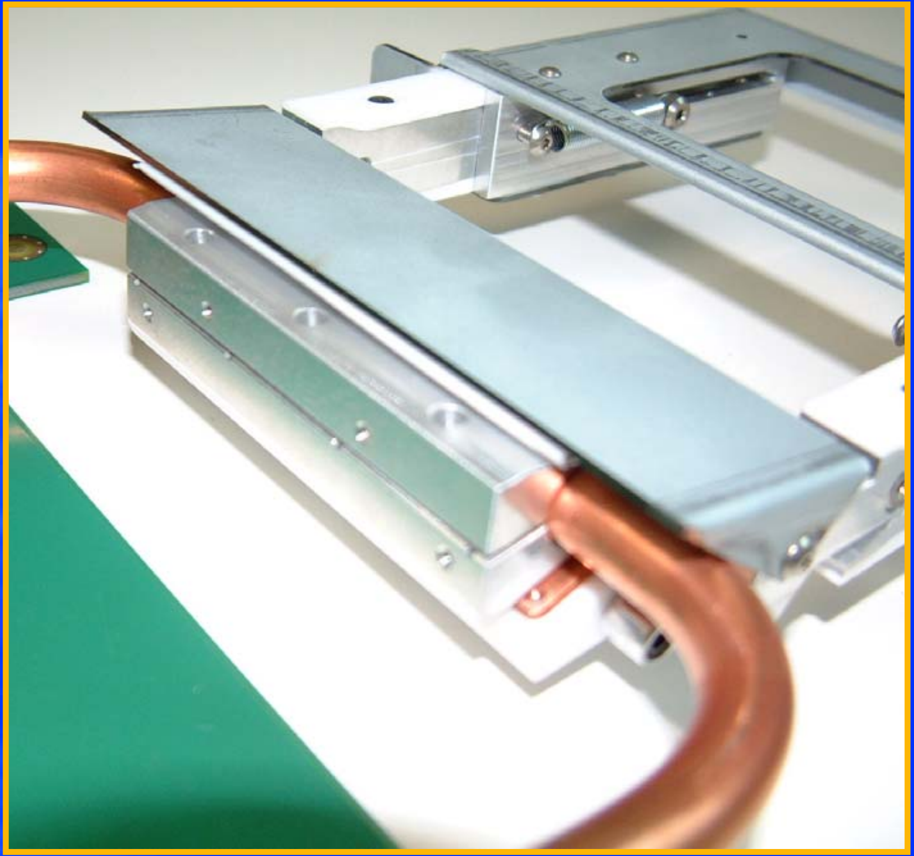

Spring-loaded receiver, shown with Al wedge (Cu wedge optional)

Thermal resistance analysis

- Thermal circuit in loop
- Sum resistances
  - Total:
    - 0.16˚C/W
  - Measured:
    - 0.17˚C/W

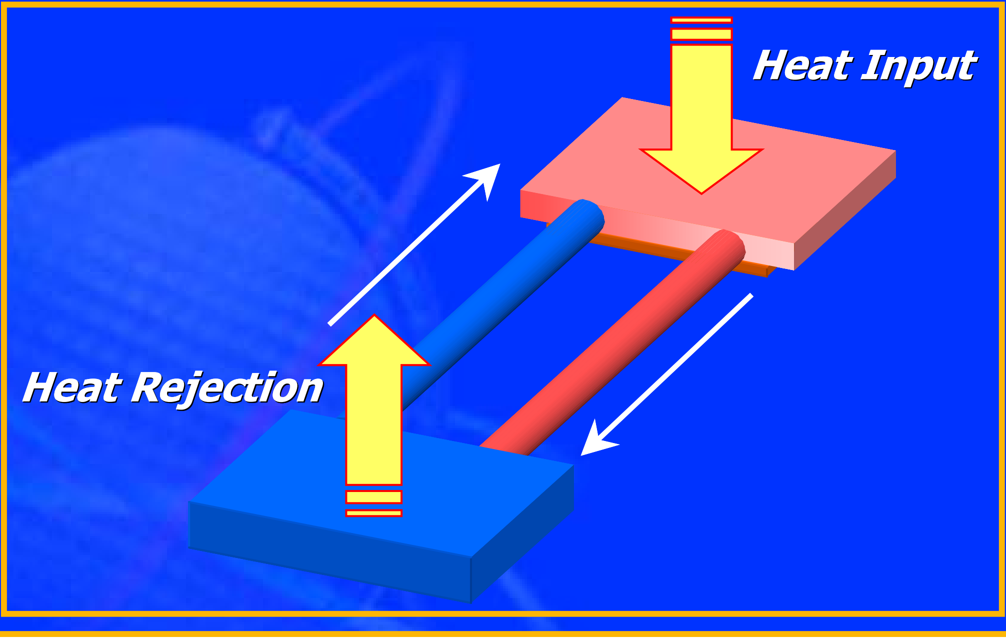

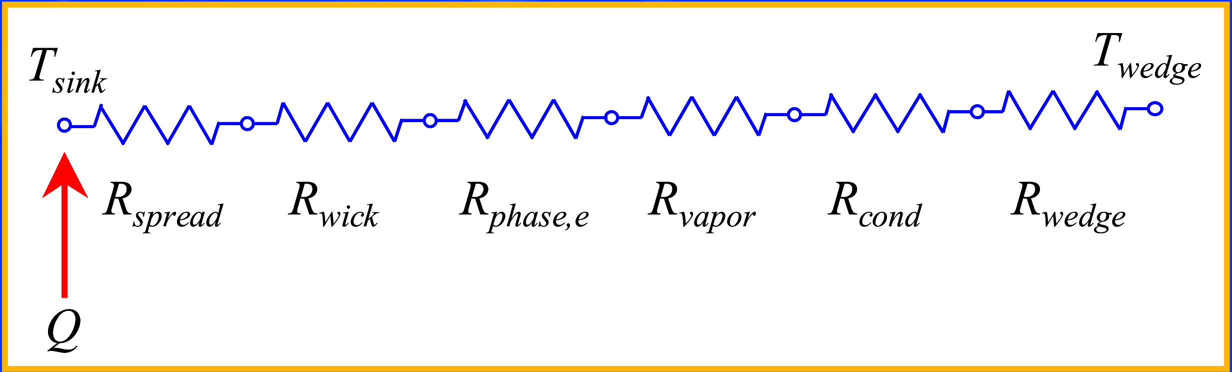

Metal-to-Metal Contact

- Thermal resistance at junction
- Correlation of Mikic [2]:

$$\theta_{contact} = \frac{\sigma}{m}\frac{h_c}{k_s} = 1.54\left(\frac{\sqrt{2}P}{mE'}\right)^{0.94}$$

- Insertion force requirements met

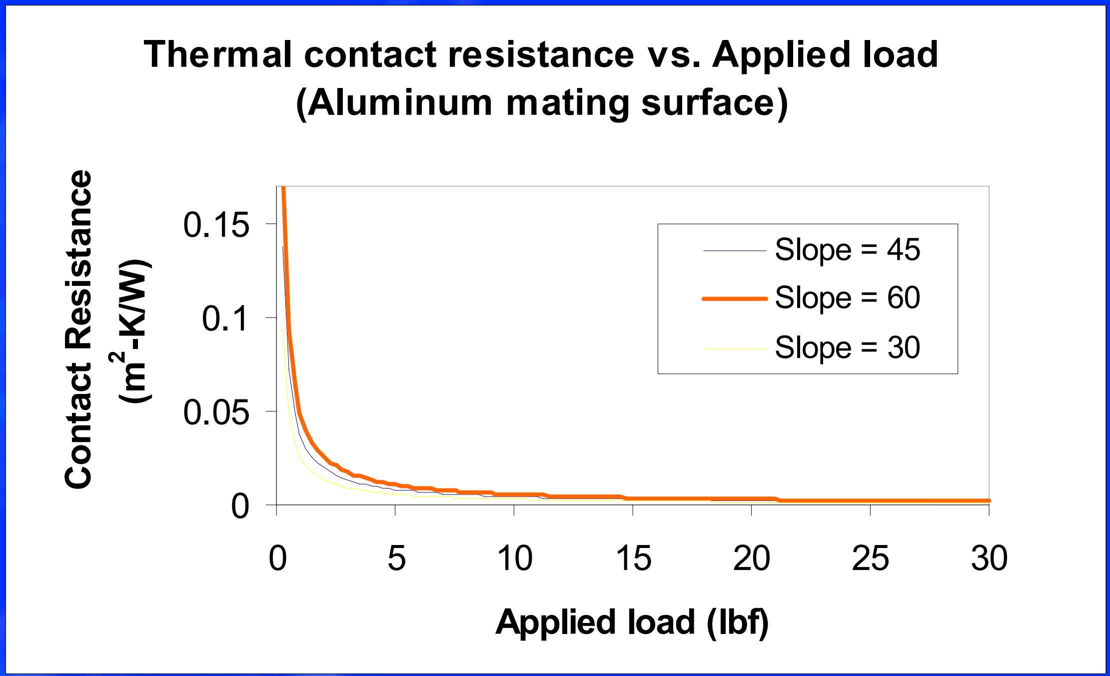

$$\theta_{contact} = f\left(k_i, E_i, \sigma_i, m_i, H_i, S_{yield}, P_{applied}\right)$$

Application – concepts

‘Hurricane Ridge’ – 2P 1U server

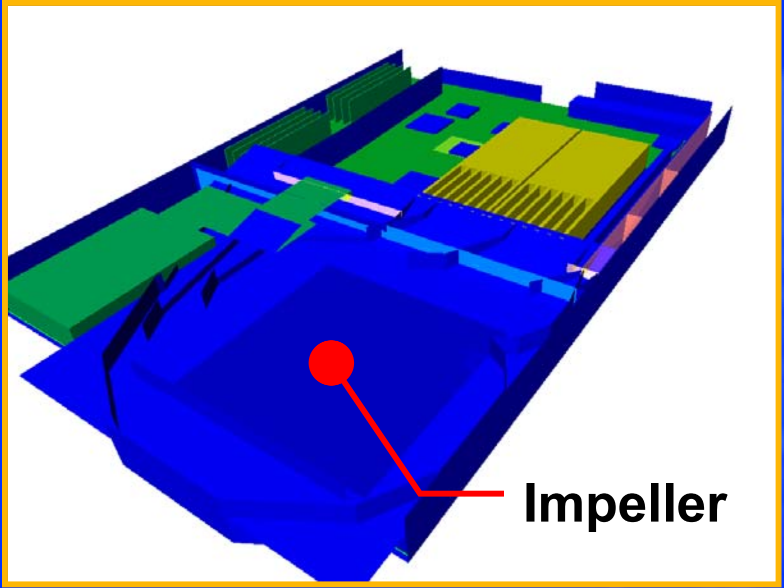

Current setup

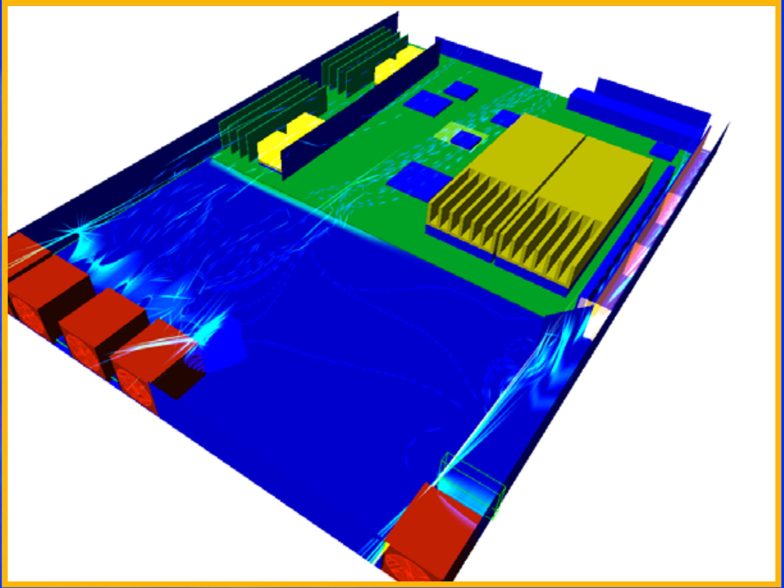

With IRCDS modifications

- Save space – 1006 cm^2 regained
- Acoustical implications – less bulk airflow, less noise

Application – concepts

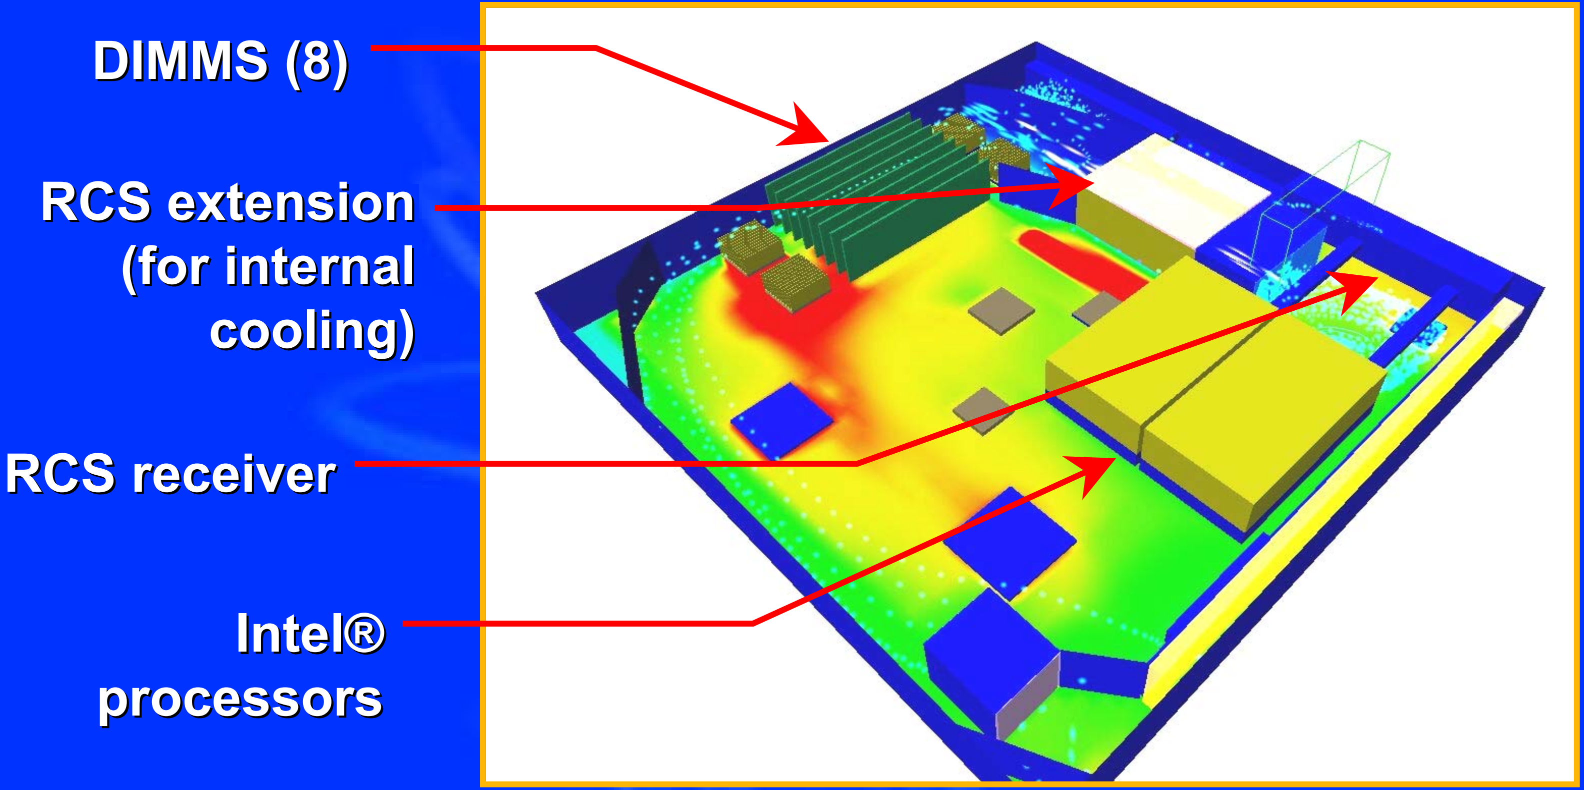

- No pass-through of air
- Fully sealed system

Server Length = 43cm

Concluding Remarks

Pros/Cons

- Pros
  - Serviceability
  - Decreased human error
  - Increased thermal flexibility
  - Potential solution to DC thermal issues
  - TCO decreased
- Cons
  - Added infrastructure
  - “Grounds-up” design
  - Decreased cable length

Paradigm Shift for Data Centers

Summary

- The IRCDS concept is…
  - Electrical and thermal docking interface
- IRCDS represents…
  - Potential new directions for enterprise technology

Increased serviceability and design flexibility

Contact info

- Will Berry
  - Intel Corp., DuPont, WA
  - will.berry@intel.com
- Steve Montgomery
  - Intel Corp., DuPont, WA
  - stephen.w.montgomery@intel.com

Backup

Acronyms

2P = Two processors
CPL = Capillary Pumped Loop
CRAC = Computer Room Air Conditioning
DC = Data Center
EMI = Electromagnetic Interference
HSD = High Signal Density
ICMB = Inter-Chassis Management Bus
I/O = Input/Output
IRCDS = Integrated Rack Cooling and Docking Station
RCS = Rack Cooling Station
RDS = Rack Docking Station
TCO = Total Cost of Ownership

Nomenclature

α = Accommodation coefficient, (See Ref. 1) [-]
μ_ℓ = Liquid viscosity, [N-s/m^2]
ν = Poisson's ratio, [-]
ρ = Density, [kg/m^3]
σ = Surface roughness, [m]
A = Area, [m^2]
D_r = Hydraulic diameter, [m]
E = Modulus of elasticity, [Pa]
E' = Reduced elastic modulus, [Pa]
g = Gravitational constant, 9.81 m/s^2
h_c = Contact conductance, [W/m^2-K]
h_{fg} = Latent heat of vaporization, [kJ/kg]
k = Thermal conductivity, [W/m-K]

Nomenclature

m = Surface asperity angle, [rad]
M = Molecular weight, [kg/kmol]
P = Pressure, [Pa]
Q = Processor power, [W]
R = Universal gas constant, 8.314 kJ/kmol-K
R = Thermal resistance, [C/W]
t_{wick} = Wick thickness, [m]
T_v = Vapor temperature, [C]
v_{fg} = Specific volume of vaporization, [m^3/kg]

Subscripts

cond = Condensation
ph,e = Phase change, evaporation
pipe = Refers to pipe
s = Refers to surface
sat = Saturation value
sink = Location, at evaporator
v = Refers to vapor
vap = Vaporization
wedge = Location, at wedge condenser
wick = Refers to wick

References

- [1] Faghri, A., 1995, *Heat Pipe Science and Technology*, Taylor and Francis, New York.
- [2] Rohsenow, W.M., *et al.*, Eds., 1998, *Handbook of Heat Transfer, 3rd Ed.*, McGraw-Hill, New York.

Example Configuration

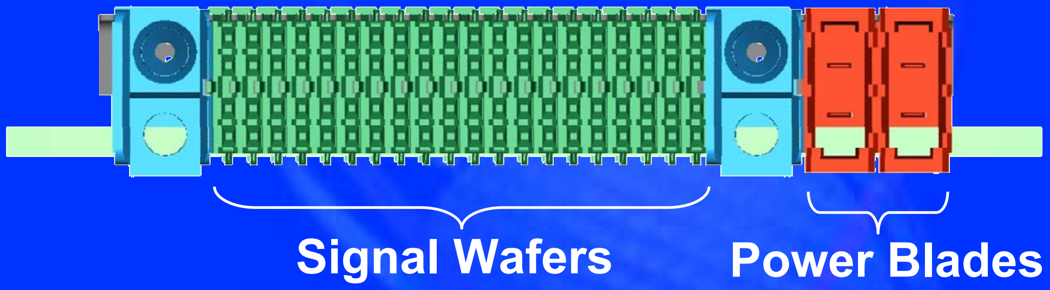

RDS connector

Server management ICMB, RS232

(2) Gigabit Ethernet channels

(2) InfiniBand* 4X channels

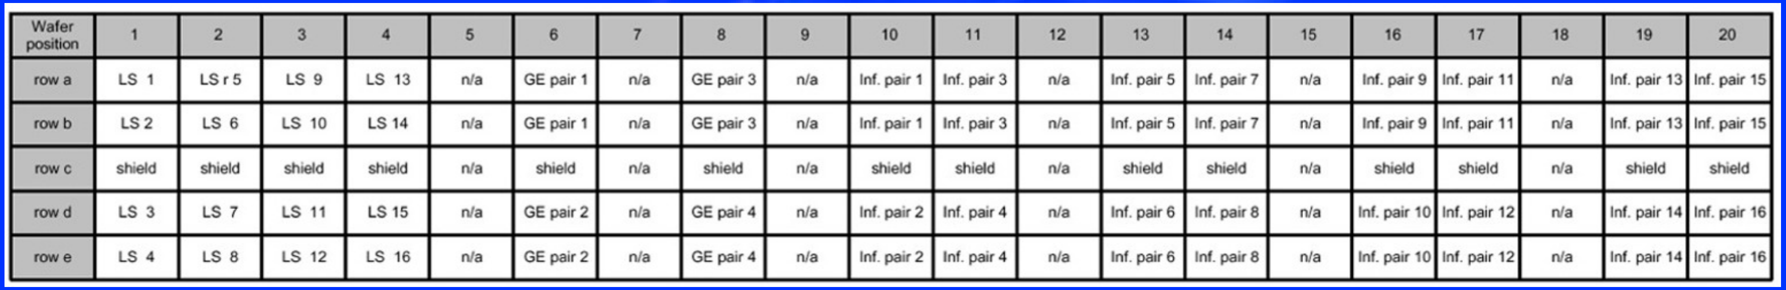

Wafer position	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
row a	LS 1	LS r 5	LS 9	LS 13	n/a	GE pair 1	n/a	GE pair 3	n/a	Inf. pair 1	Inf. pair 3	n/a	Inf. pair 5	Inf. pair 7	n/a	Inf. pair 9	Inf. pair 11	n/a	Inf. pair 13	Inf. pair 15
row b	LS 2	LS 6	LS 10	LS 14	n/a	GE pair 1	n/a	GE pair 3	n/a	Inf. pair 1	Inf. pair 3	n/a	Inf. pair 5	Inf. pair 7	n/a	Inf. pair 9	Inf. pair 11	n/a	Inf. pair 13	Inf. pair 15
row c	shield	shield	shield	shield	n/a	shield	n/a	shield	n/a	shield	shield	n/a	shield	shield	n/a	shield	shield	n/a	shield	shield
row d	LS 3	LS 7	LS 11	LS 15	n/a	GE pair 2	n/a	GE pair 4	n/a	Inf. pair 2	Inf. pair 4	n/a	Inf. pair 6	Inf. pair 8	n/a	Inf. pair 10	Inf. pair 12	n/a	Inf. pair 14	Inf. pair 16
row e	LS 4	LS 8	LS 12	LS 16	n/a	GE pair 2	n/a	GE pair 4	n/a	Inf. pair 2	Inf. pair 4	n/a	Inf. pair 6	Inf. pair 8	n/a	Inf. pair 10	Inf. pair 12	n/a	Inf. pair 14	Inf. pair 16

Signal Integrity Impact

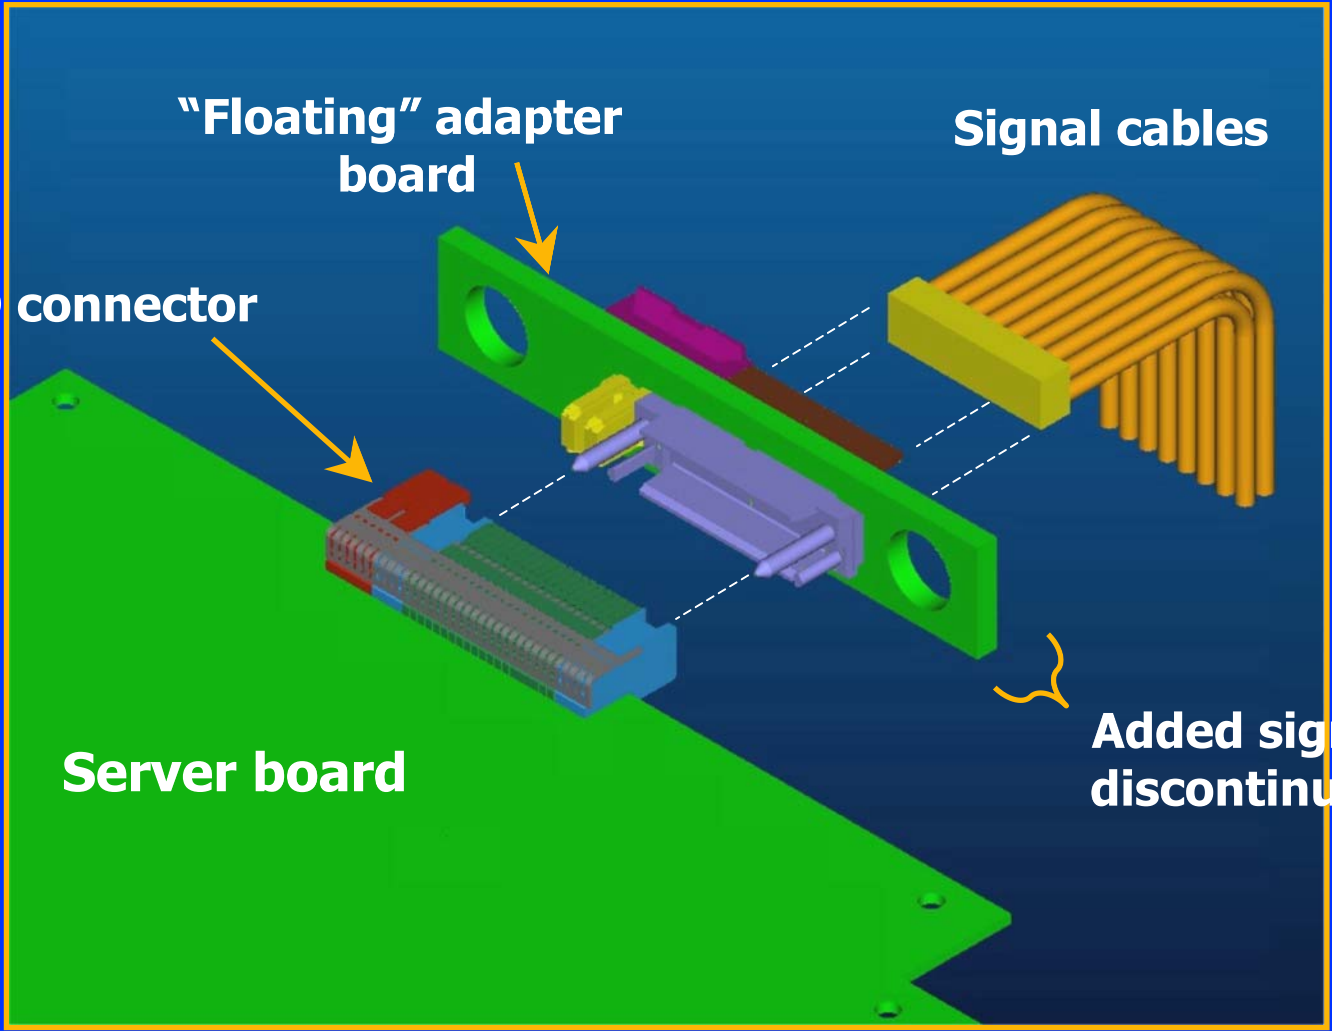

Analysis

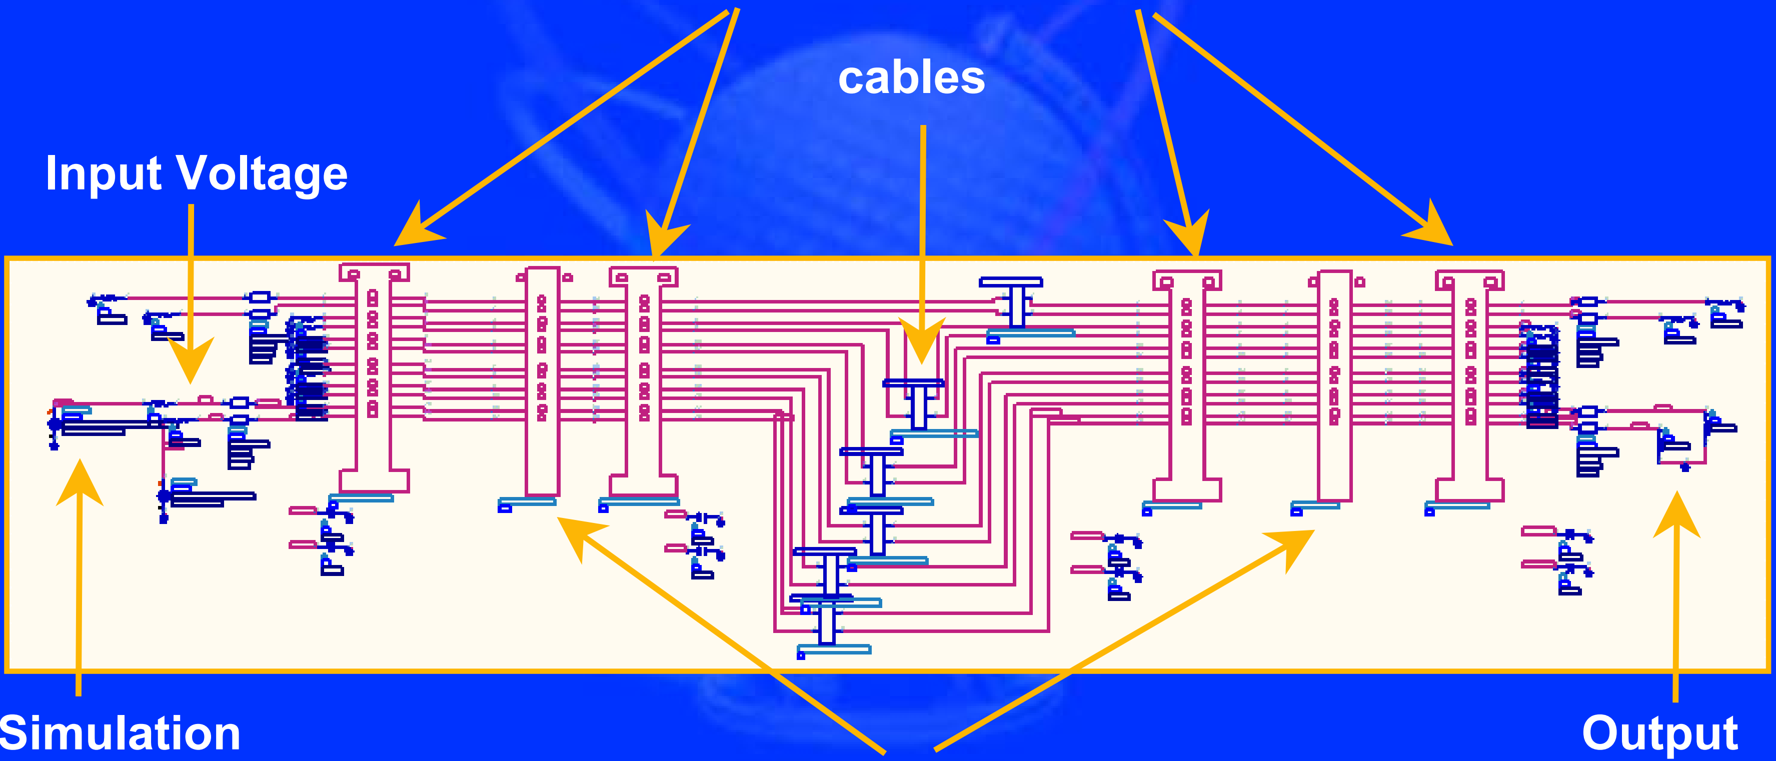

HSD 5 Row Connectors

cables

Input Voltage

Simulation source

Adaptor board

Output Voltage

Analysis courtesy of Molex Inc.

Rack Architecture

- Servers and switches in one rack
- Worst-case scenario (server to switch)
  - length = 2.4m

2.5m Provides Adequate Length

Rear view

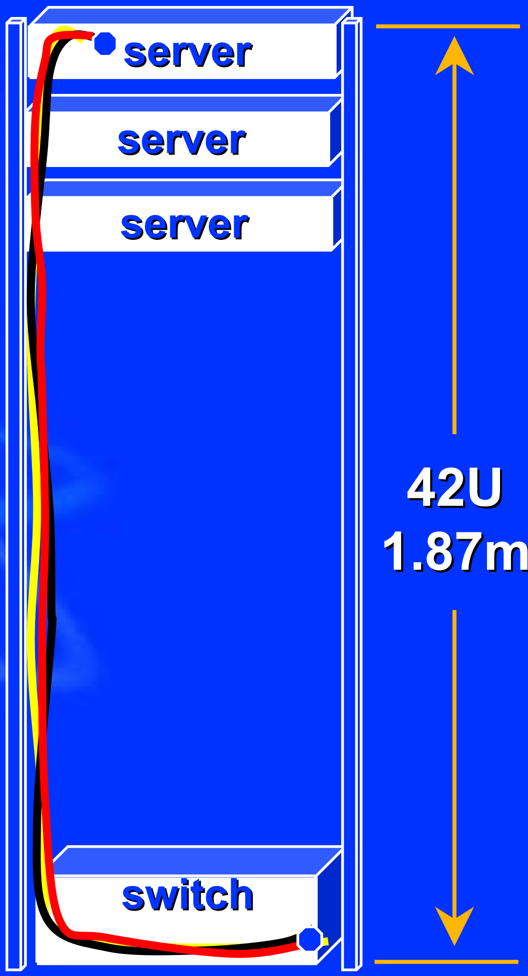

Receiver Details

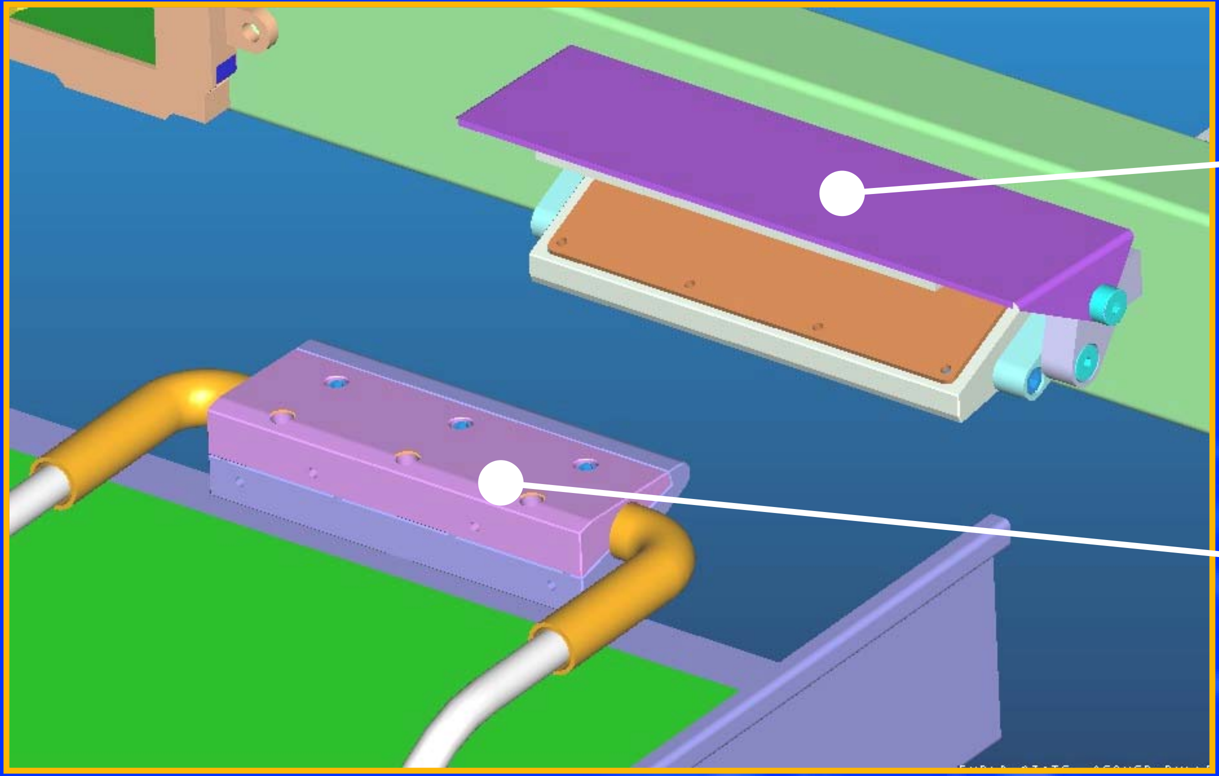

- **Heat transfer accomplished through contact of wedge and chiller plates**

Analysis of CPL

- The resistances are [1]:

$$R_{wick} = \frac{t_{wick}}{k_{wick} A_{wick}} = 0.0024$$

$$R_{ph,e} = \left[\left(\frac{2\alpha}{2-\alpha} \right) \left(\frac{h_{fg}^{2}}{T_{v} \upsilon_{fg}} \right) \left(\frac{\overline{M}_{v}}{2\pi \overline{R} T_{v}} \right)^{1/2} \left(1 - \frac{P_{v} \upsilon_{fg}}{2 h_{fg}} \right) \right]^{-1} A_{wick}^{-1} = 9.5 * 10^{-7}$$

$$R_{vap} = \frac{T_{v} \Delta P_{vapor}}{\rho_{v} h_{fg} Q} = 4 * 10^{-7}$$

$$R_{cond} = \left[\frac{g \rho_{\bullet} \left(\rho_{\bullet} - \rho_{v} k_{\bullet}^{3} h_{fg}^{'} \right)}{\mu_{\bullet} \left(T_{sat} - T_{s} \right) D_{r}} \right]^{-1/4} A_{pipe}^{-1} = 0.111$$

$$R_{wedge} \approx 0.04$$

- Summing the resistances,

Theoretical R_{total} = 0.16°C/W

Benchmark: 0.25+°C/W