神州数码
Digital China

Next Generation Cloud Data Center

Jason Pan
Senior Director and Consultant
SGI and Digital China

Cloud before 1990s

Cloud Today

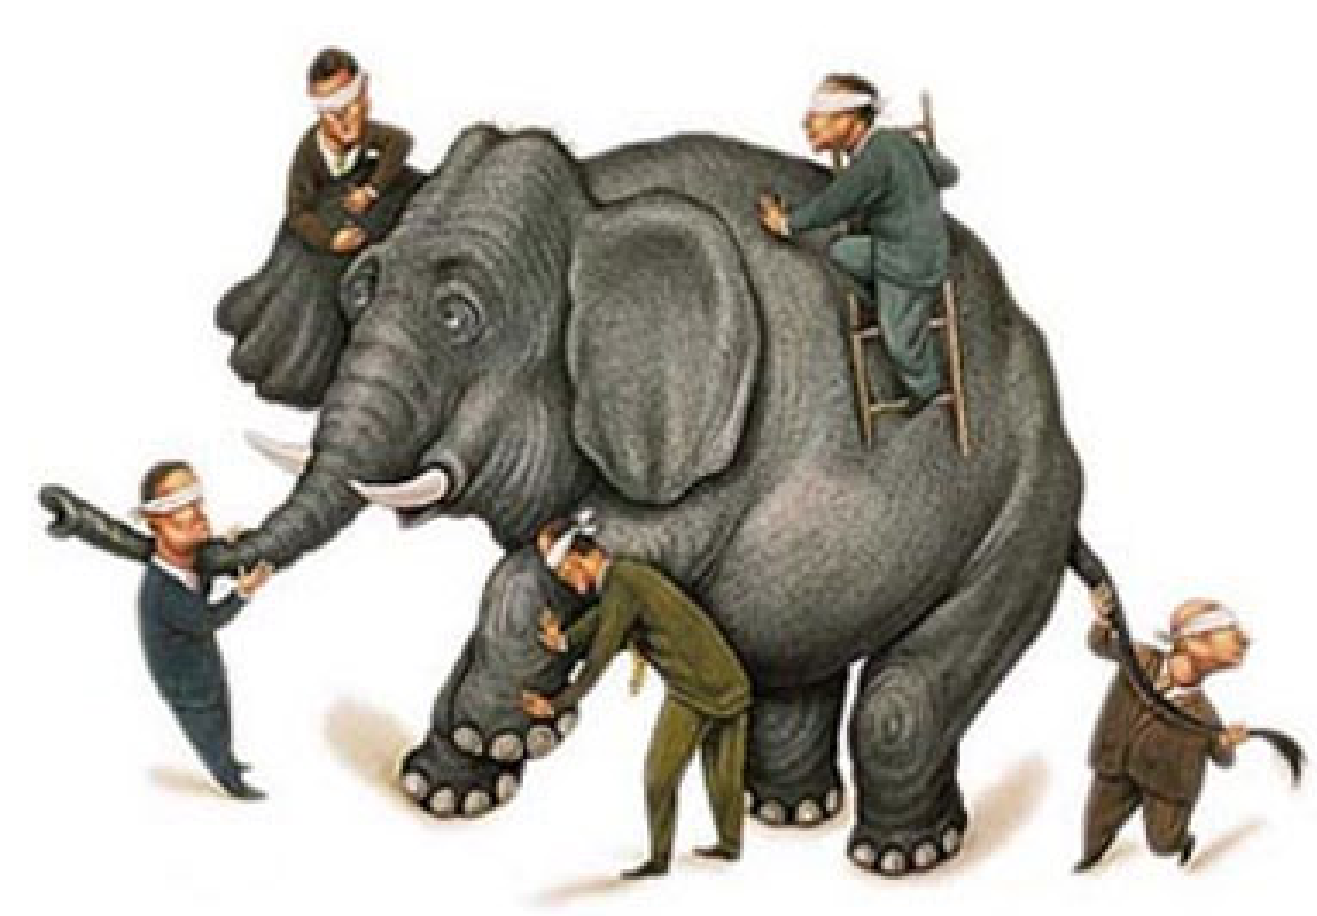

神州数码
Digital China

The road to cloud…

- “The network is the computer” – Sun (1988)
- .dotcom era in 2000 – Outsource IT services and infrastructure
- Web email (Yahoo, Hotmail, Gmail) & YouTube
- Salesforce.com and Amazon Web Services
- Nicholas Carr – The “Big Switch” 2008 and “IT is dead”
- Mobile Apps on hundreds of Million app-capable smart devices

What are we seeing in Cloud?

- Moved from buzzword and hype to real technology
- Move from “When” to “How”
- Agreements on Various Definitions
  - Payment Model
    - Software as a Service (SaaS)
    - Platform as a Service (PaaS)
    - Infrastructure as a Service (IaaS)
  - Delivery Model
    - Public, Private, Hybrid Clouds

神州数码 Digital China

Comparing the Delivery Vehicles

Traditional Delivery	*Cloud Based*
Applications	Software-as-a-Service (SaaS)
Development Tools	Development-as-a-Service (Daas)
Middleware	Platform-as-a-Service (PaaS)
Infrastructure and Hardware	Infrastructure-as-a-Service (IaaS)

Cloud trends in 2011 – Private Clouds

- The move to a hybrid world
  - Large enterprises build out private clouds and leverage public clouds for workloads such as
    - Dev/Test, Archiving, Disaster Recovery, Redundancy, Web-Serving
  - Gartner expects large enterprises to have a dynamic sourcing team for cloudsourcing and management
  - Focus on selecting the right building blocks

Cloud trends in 2011 – Big Data

- Focus on “Big Data” and Databases
  - The old paradigm of RDBMS is no longer valid
  - Cloud-native database
    - “No SQL”
    - Graph Databases
    - Hardoop HBase
- Database.com – Transactional data processing and analysis
- Cloud is great for “Data-as-a-service” combined with Compute and Storage from Infrastructure as a service

神州数码
Digital China

Cloud trends in 2011 – Truly Global

- Cloud allows Emerging markets to levels the playing field
  - New business models and access to markets
  - No previous IT investment or Legacy baggage
- US based cloud providers are going after Asia and EMEA markets

Cloud Platform Design Guidelines

- Modular architecture
  - Compute, Storage, Management, Network
- Homogeneous pool of compute and storage with uniform connectivity
- Compute –
  - No more than 16 “small” machine instances per server
  - Managed/VLAN networking mode to ensure mutli-tenancy network security

Cloud Platform Design Guidelines - Storage

- Single large store for all cloud storage
  - DAS, SAN, NAS
- Block Storage (EBS) vs Application Aware(S3)
- Separate persistent vs non-persistent
- Use SSD and Flash to improve performance
  - Compute DAS for caching VM images
  - Meta-data Store
- Storage based on “cheap good enough storage” (JBODs)

Design Guidelines - Redundancy

- Redundancy moves from hardware into the software and application
- Most public clouds builders lack network, power and disk redundancy
- PDU and Rack Switch failure will take down lots of servers
- Flatness at Edges – Avoid tiers when possible

神州数码 Digital China

Design Guidelines – Hyperscale

- Simple network routing and configuration
- Higher Density
- Minimize shared physical infrastructure
- Modularity, consistency and simplicity is the key
- Have the ability to add capacity to the datacenter rapidly

Cloud Datacenter Key Design Points

- High x86 server & storage density
  - Designed to scale
- Build To Order flexibility
  - Broad range of configurations
- Eco-Logical™ focus on lowering OpEx:
  - Low power consumption
  - High efficiency AC and DC power supplies
  - Innovative, effective cooling techniques

神州数码
Digital China

Rackable Half-Depth Advantage

Flexibility

Choice of 22U, 36U, 40U, 44U racks & full range of AMD/Intel configurations

Density

Half-depth server design mounted back-to-back for 2x density

Thermal Management & Power Distribution

Unobstructed central air plenum for superior ventilation

DC Power option provides power savings of up to 30%

Passive cooling design for highest reliability

Serviceability

I/O in front and clean cabling reduces maintenance time by up to 75%

Sophisticated remote management capabilities

sgi

神州数码 Digital China

CloudRack Sample Tray

Eco-Logical™ Fanless, Coverless & Power-Supply-less Design

Ease of Use – Blind mate 12V DC Power Connector

Customer Choice – Huge Selection of Intel and AMD Motherboards

Investment Protection – power and cooling moved to cabinet level

High Density: 1.75”H (1U) x 19.5”W x 31”D Tray Form Factor

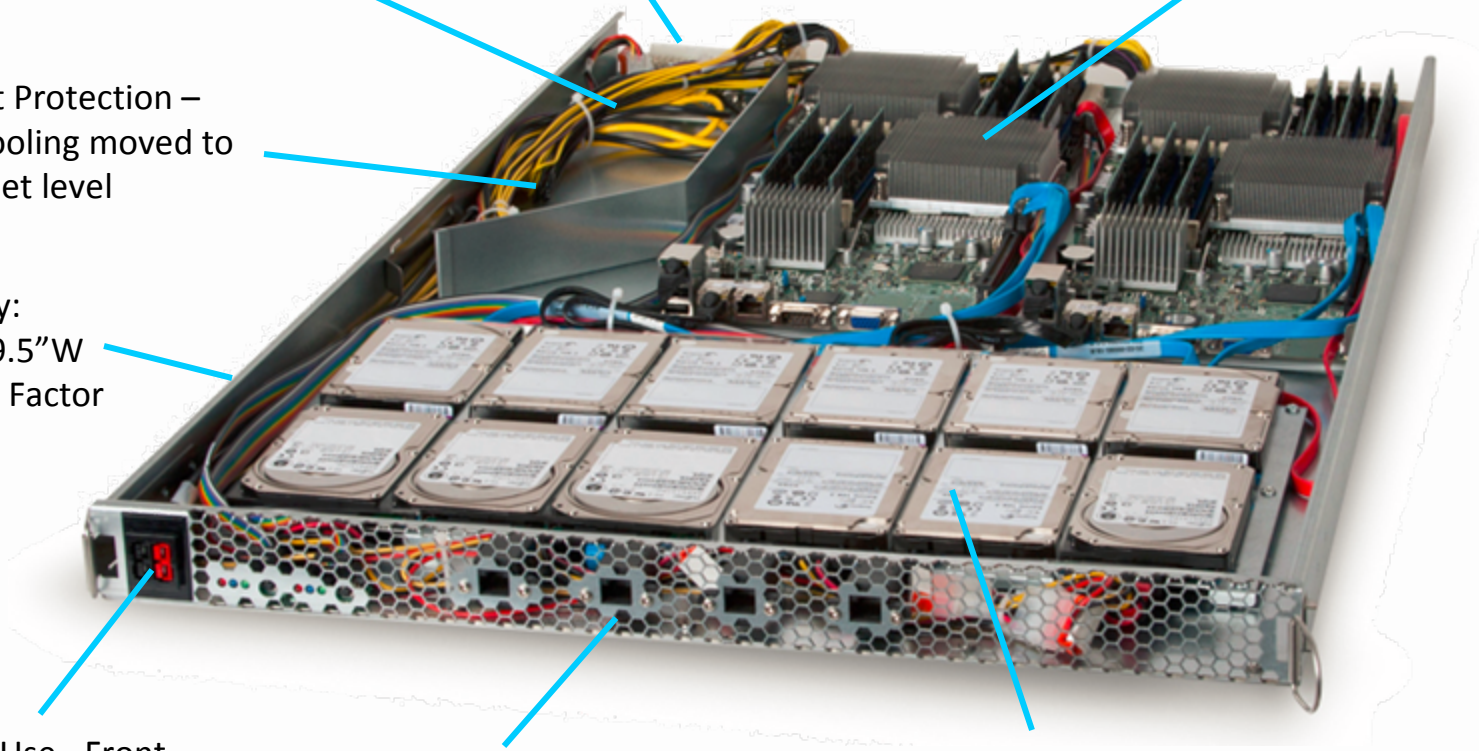

Ease of Use - Front Access Service Loop

Full IPMI 2.0 Remote Management & Serial Redirect

Best in Class Storage Density – Selection of “Quick Release” 2.5” or 3.5” Drives

神州数码 Digital China

CloudRack™ C2 Cabinet (42U)

Serviceability: Super-Easy Top & Front Access to All Key Components & Cabling in Seconds

Significant Reduction of Metallic Piece Parts >> Ecological Design and Lower TCO

Ultra Efficient & Standard 24”Wx 46”D Design >> Best in Class Compute & Storage Density

Customer Choice of 24U & 42U (78.7” Tall) Cabinets

Redundant, Hot Swap 2x3 Fan Arrays >>> High Reliability

Common Cooling Infrastructure >>> High Reliability & Efficiency

12V DC Power Distribution…99% Efficiency to the Trays

Redundant, Hot Swap Rectifiers >> High Reliability

Reconfigurable AC Input (3xL6-30, 2 Phase Connectors Shown)

CloudRack™ C2 Innovations Offer Dramatic Advancement in Efficiency & Reliability

神州数码 Digital China

BTO Tray Examples

Model TR2108-2TY5 Example

32 Intel Xeon 5500 Cores/U

For App/Database

Model TR2106-3SC1 Example

24 Cores/U

For Web/Mid Tier

Model TR2108-F7 Example

Quad Opteron 8300: 16 Cores/U

For Database & Virtualization

Rich Selection of Base Boards Optimize for Wide Range of Applications

Model TR2106-6AT1 Example*

6 Servers & 12 Cores/U

MicroSlice For Web/Cloud Computing

sgi

神州数码 Digital China

* Picture shows CloudRack TR1000-AT1 Tray. TR2106-6AT1 removes the Power Supply & AC Connector in Front

1. Power XE: Power Distribution Tech.

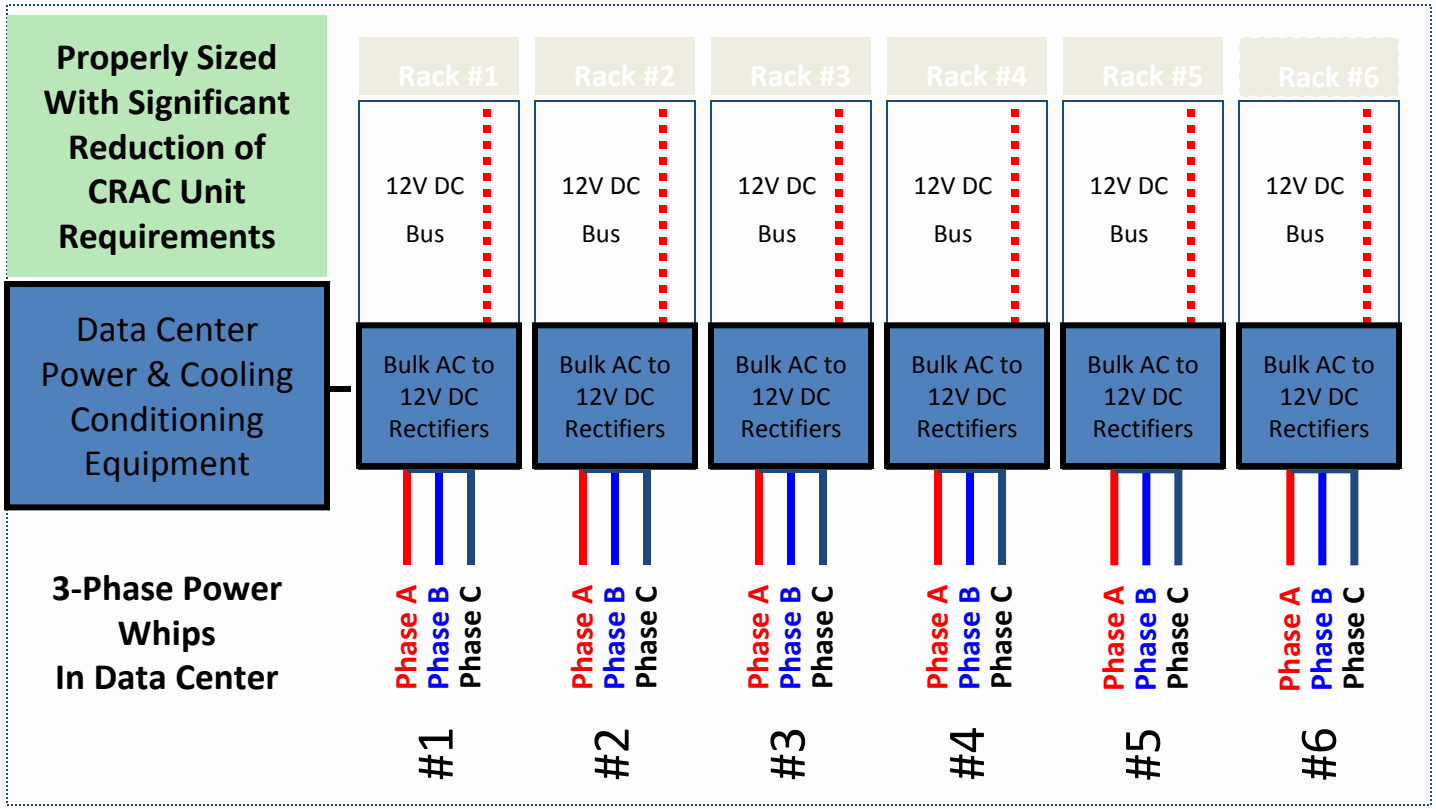

- Power XE virtually eliminates the stranded power problem
  - # of bulk rectifiers in cabinet perfectly divisible by 3 (phases)
  - Near perfect (>95%) phase balance with 99% power factor
  - Load imbalance not appear as “feed imbalance” due to isolation

神州数码 Digital China

Other Power XE Key Advantages

- Significant improves the power dist. reliability & availability
  - AC cabinet with 99% 12V efficiency to the server
  - Hot-pluggable, N+1 or better redundant rectifiers
- Decreases point of failures
  - All server level power supplies eliminated
- Minimizes injection of “power harmonics” to the power mains
  - AC Total Harmonic Distortion below 5% with no single harmonic above 3%
  - More servers on the same circuit
  - 0.99 Power factor
- Incorporates real time rack level power metering

神州数码 Digital China

2. Dramatic Thermal Efficiency

- Each redundant six-fan array is “Autonomic” – monitoring and adjusting fan speed as required
  - Uniform flow pattern throughout cabinet
- Supports elevated temperature environments – up to 40° C (104° F)
  - Improved data center Power Usage Effectiveness (PUE) through reduced use of Computer Room Air Conditioning (CRAC) units
- Improved rack-level cooling power profile
  - < 8% of rack power vs. competitive > 25%
  - 4W per server fan consumption (800W per 228 servers)

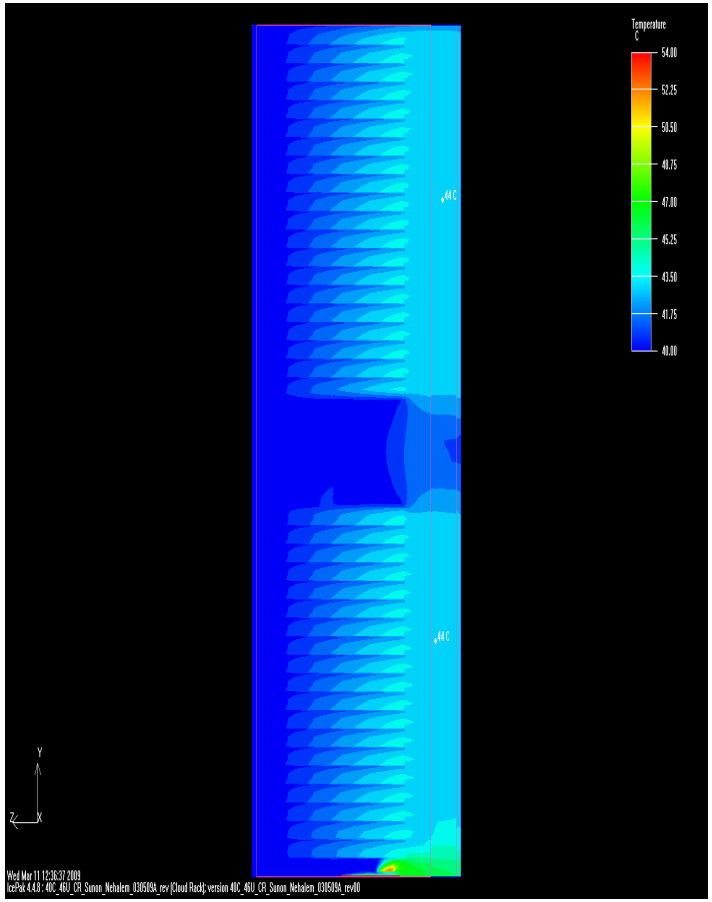

Thermal Simulation At 40° C Ambient

3. Leading Reliability

- Extending server life & availability
  - Server trays drastically simplified (no more power or cooling per server)
  - No moving parts on trays in SSD drive or diskless configurations!
  - Near elimination of Rotational Vibration (RV) transmitted to internal drives vs. competitive systems due to centralized cooling
- N+1 redundant input power and thermally-managed fan arrays
  - Significant improvement in reliability & availability
  - Hot swappable for improved serviceability
  - Redundancy at no incremental cost or loss of efficiency
  - Dramatic reduction in rack-level fan power vs. competitive products

神州数码 Digital China

4. Amazing Density

- Highest compute density in tray configs
  - Up to 76 Intel Xeon 5600 dual-processor servers per 42U (78.7”) rack
  - Up to 912 Intel Xeon 5600 cores/cabinet or 24 cores/U
- 24” wide cabinet for space efficient data center deployments
- Massive storage density – up to 24 TB/U
- Broad choice of compute/storage density ratios
  - Fan-less, coverless & now power supply-less server trays
  - ..& broad selection of base boards & storage components

神州数码 Digital China

ICE Cube “Air”


ICE Cube Air
The Expandable Modular Data Center
sgi
ICE Cube® Air

ICE Cube Air is about Integration

The process has already started…with the IT and the Datacenter

Modular Datacenters – Built, Tested, Deployed

Light Weight & Sturdy

Factory Testing

ISO Base frame

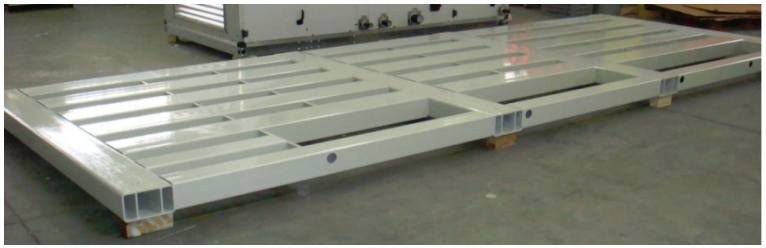

神州数码
Digital China

EXPANDABLE Unit Supporting Air, EVAP & DX Cooling

Low Voltage (208-240V) or High (380 – 480V) Head unit expandable with Tails as required

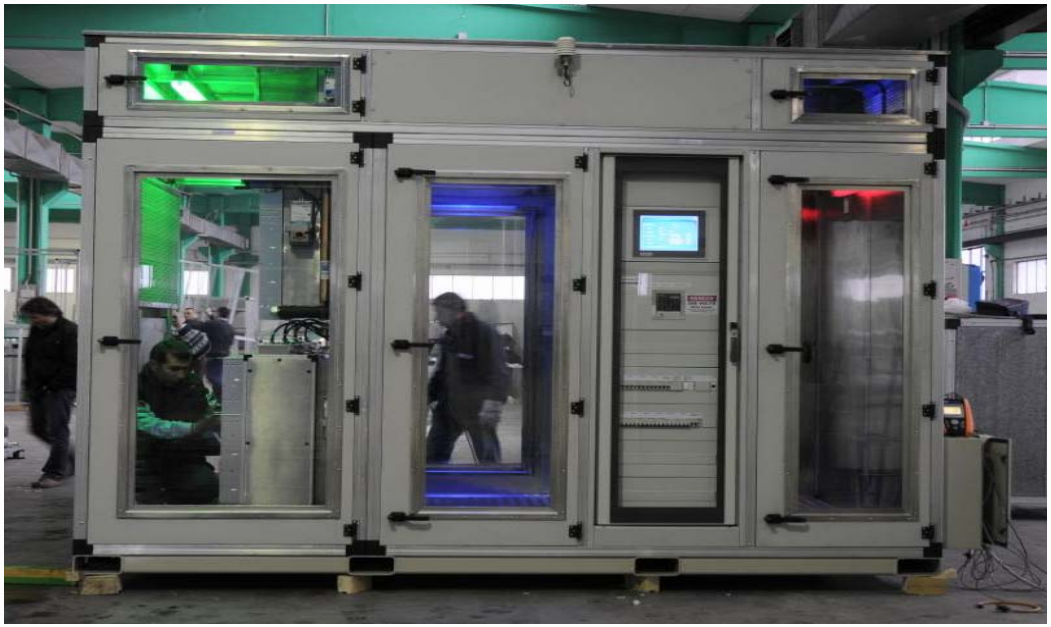

ICE Cube® Air – The Expandable Modular Datacenter
Low Voltage (208-240V)

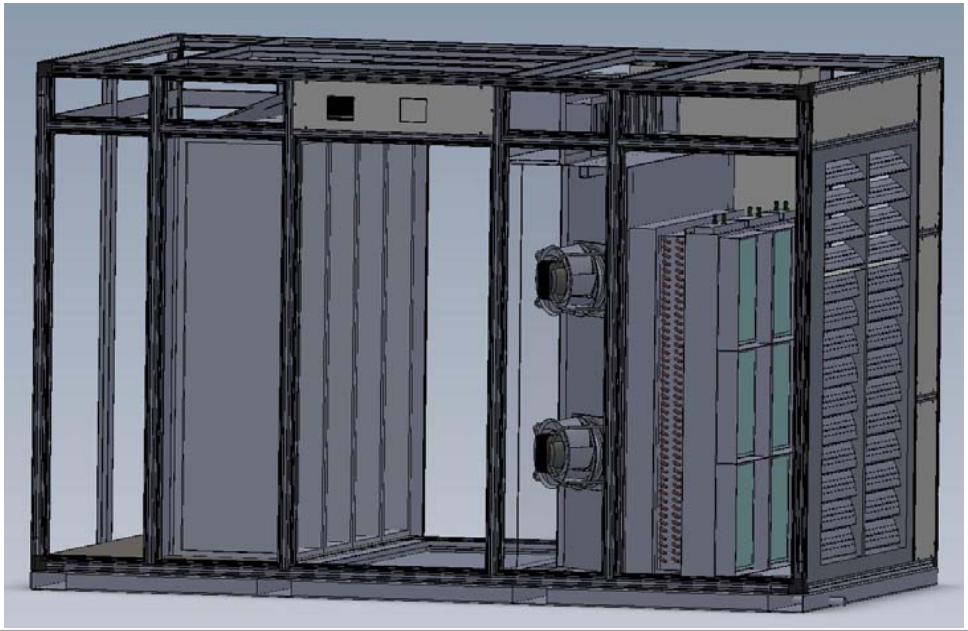

Implement a New Design Philosophy and Practice
ICE Cube® Air SMALL Head Unit Side View

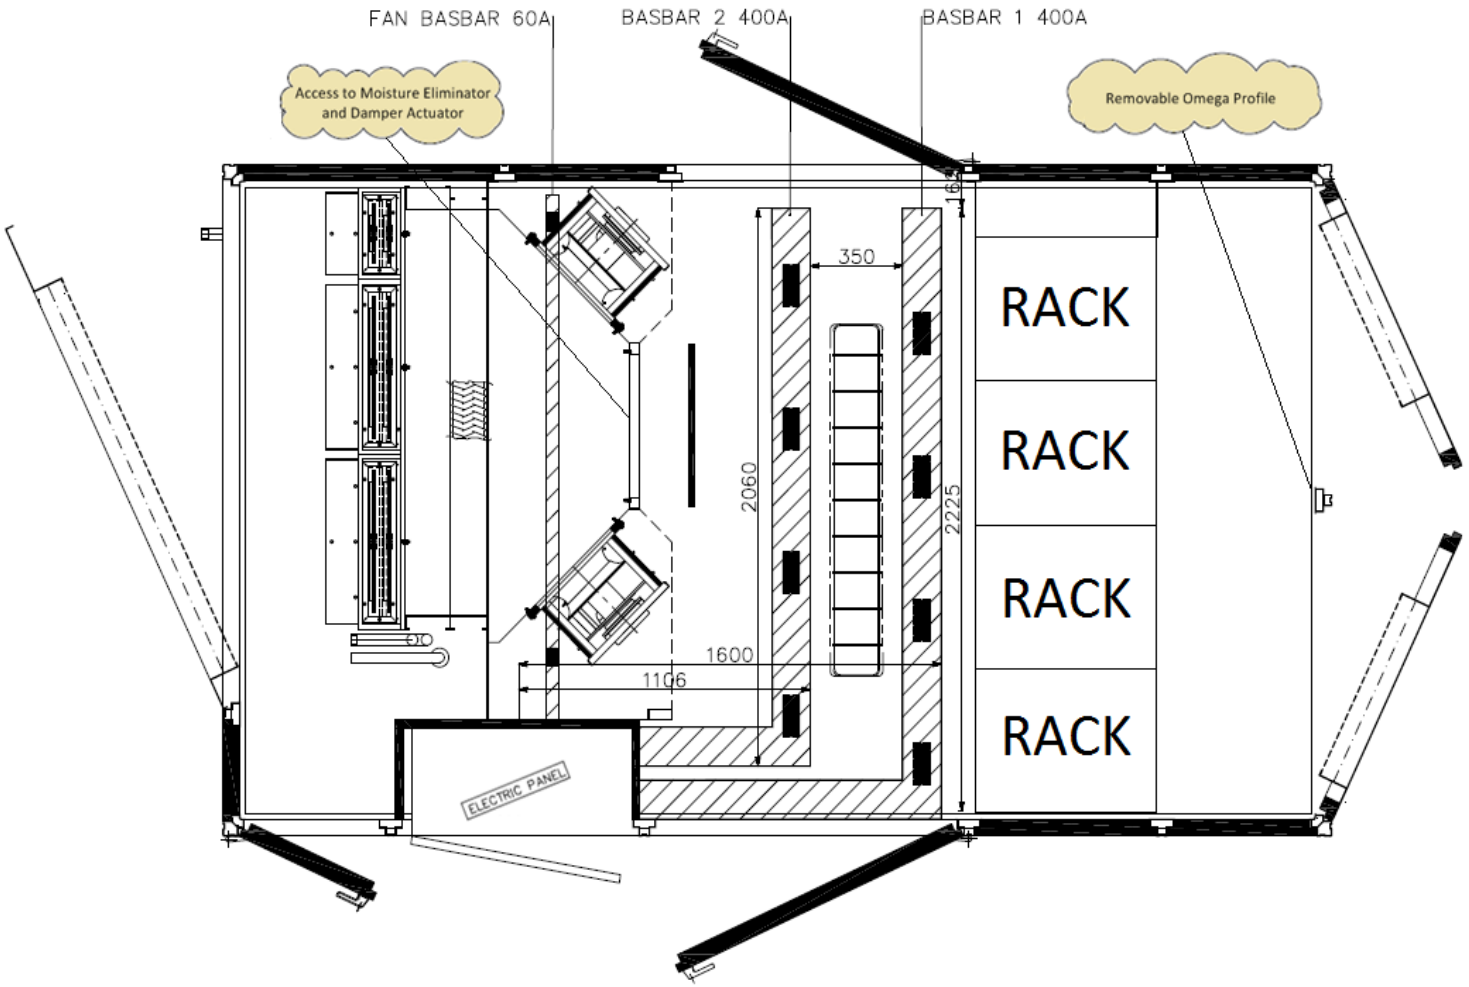

ICE Cube® Air SMALL Head Unit with 3 Expansion Modules (16 Racks)

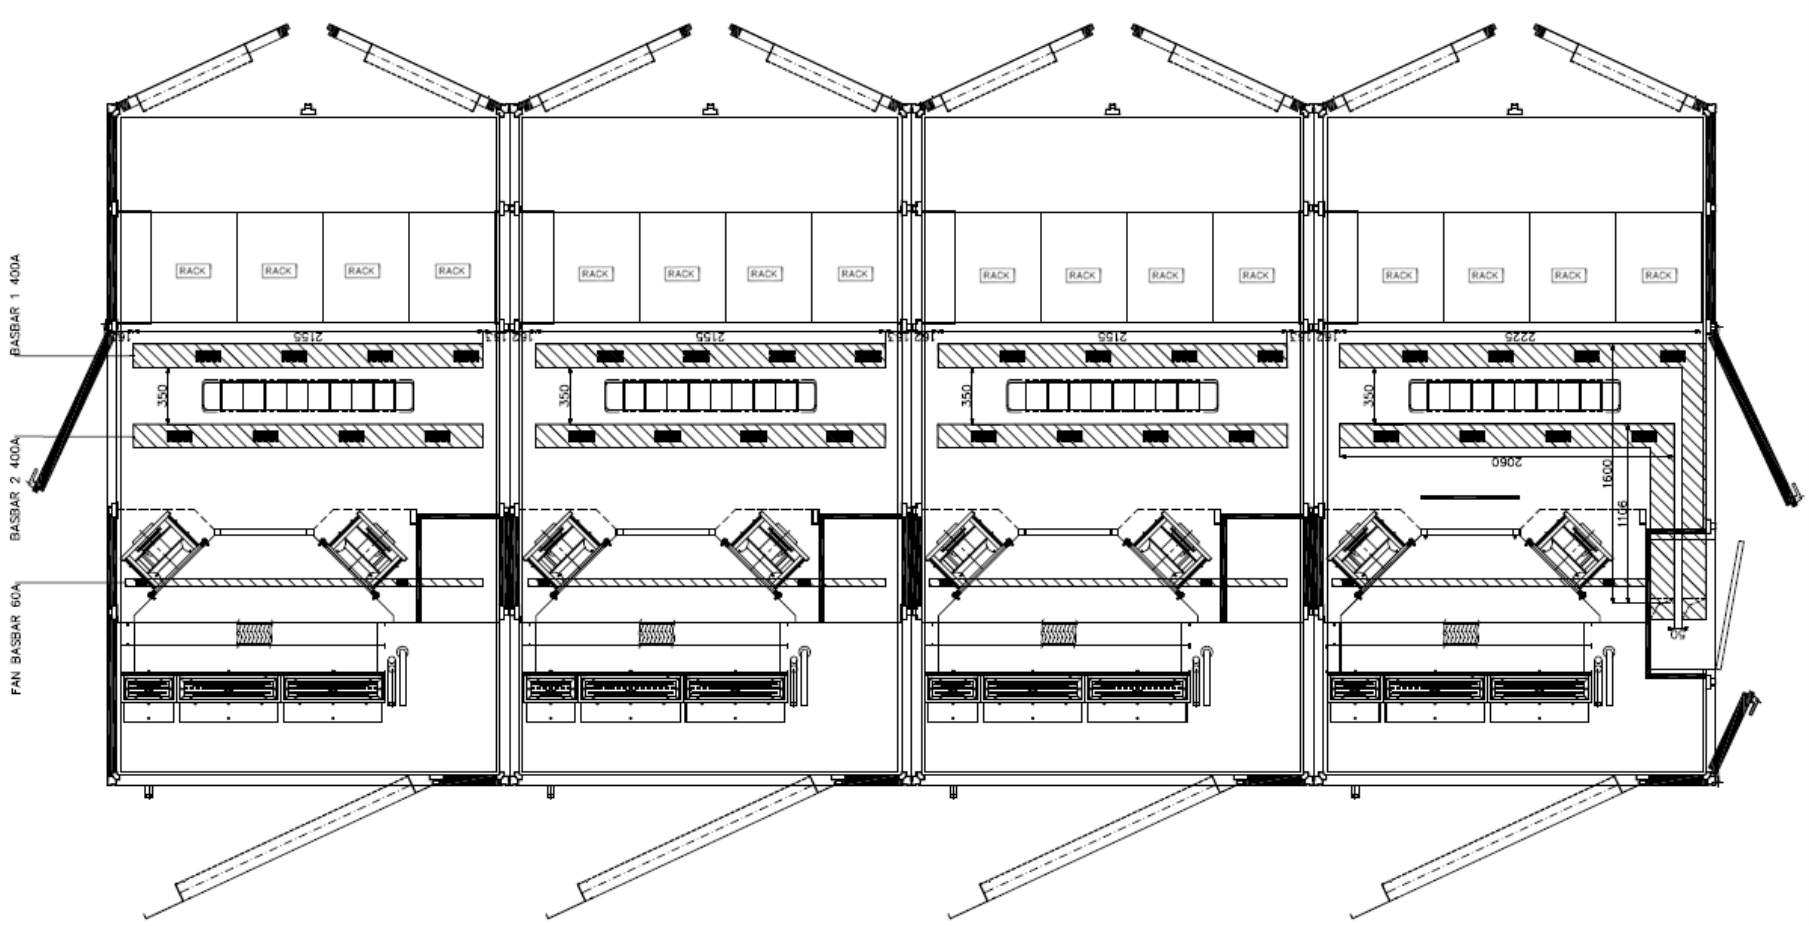

Common Hot and Cold Isles

Air-Cooled Container Deployment Area

- A Single Air-Cooled Deployment Area (DA)
  - 6 x 20 ft. segments (back-to-back 2 x 3) per DA
    - Assume equivalent of 64 roll-in racks per DA
      - May require use of built-in racks or DA length adjustments

Exhaust Blowers Above Hot Aisle

DA approximately 30 ft W x 63 ft L

Size of a Tennis Court

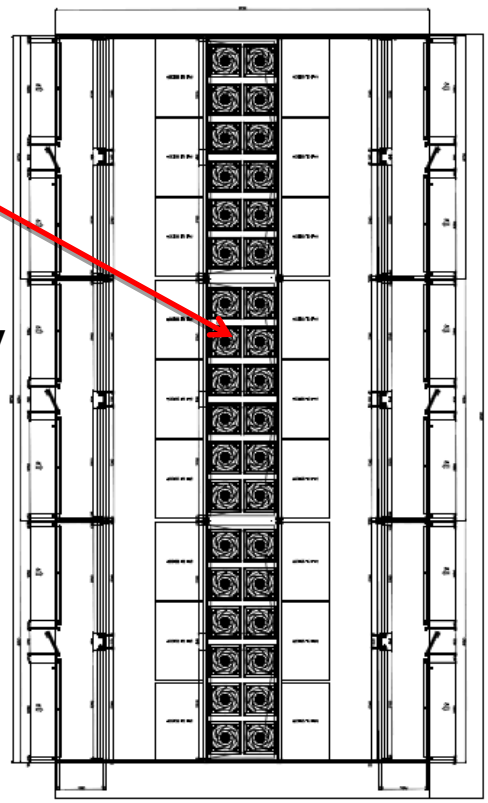

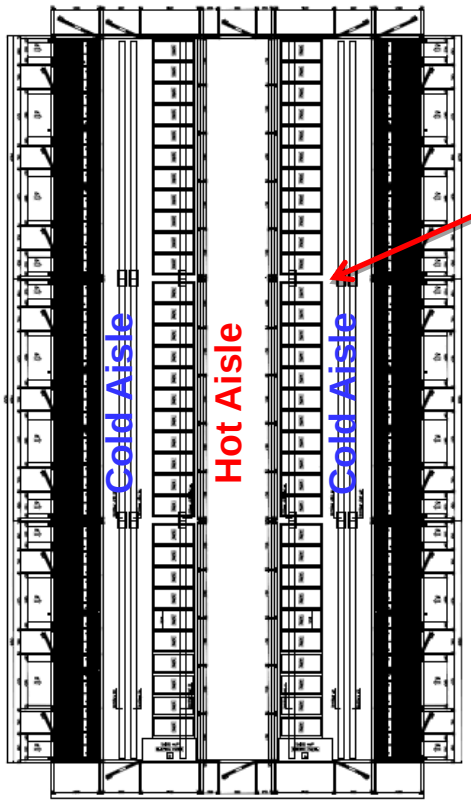

64 Racks per DA

256 Racks in 4 DA (12800RU)

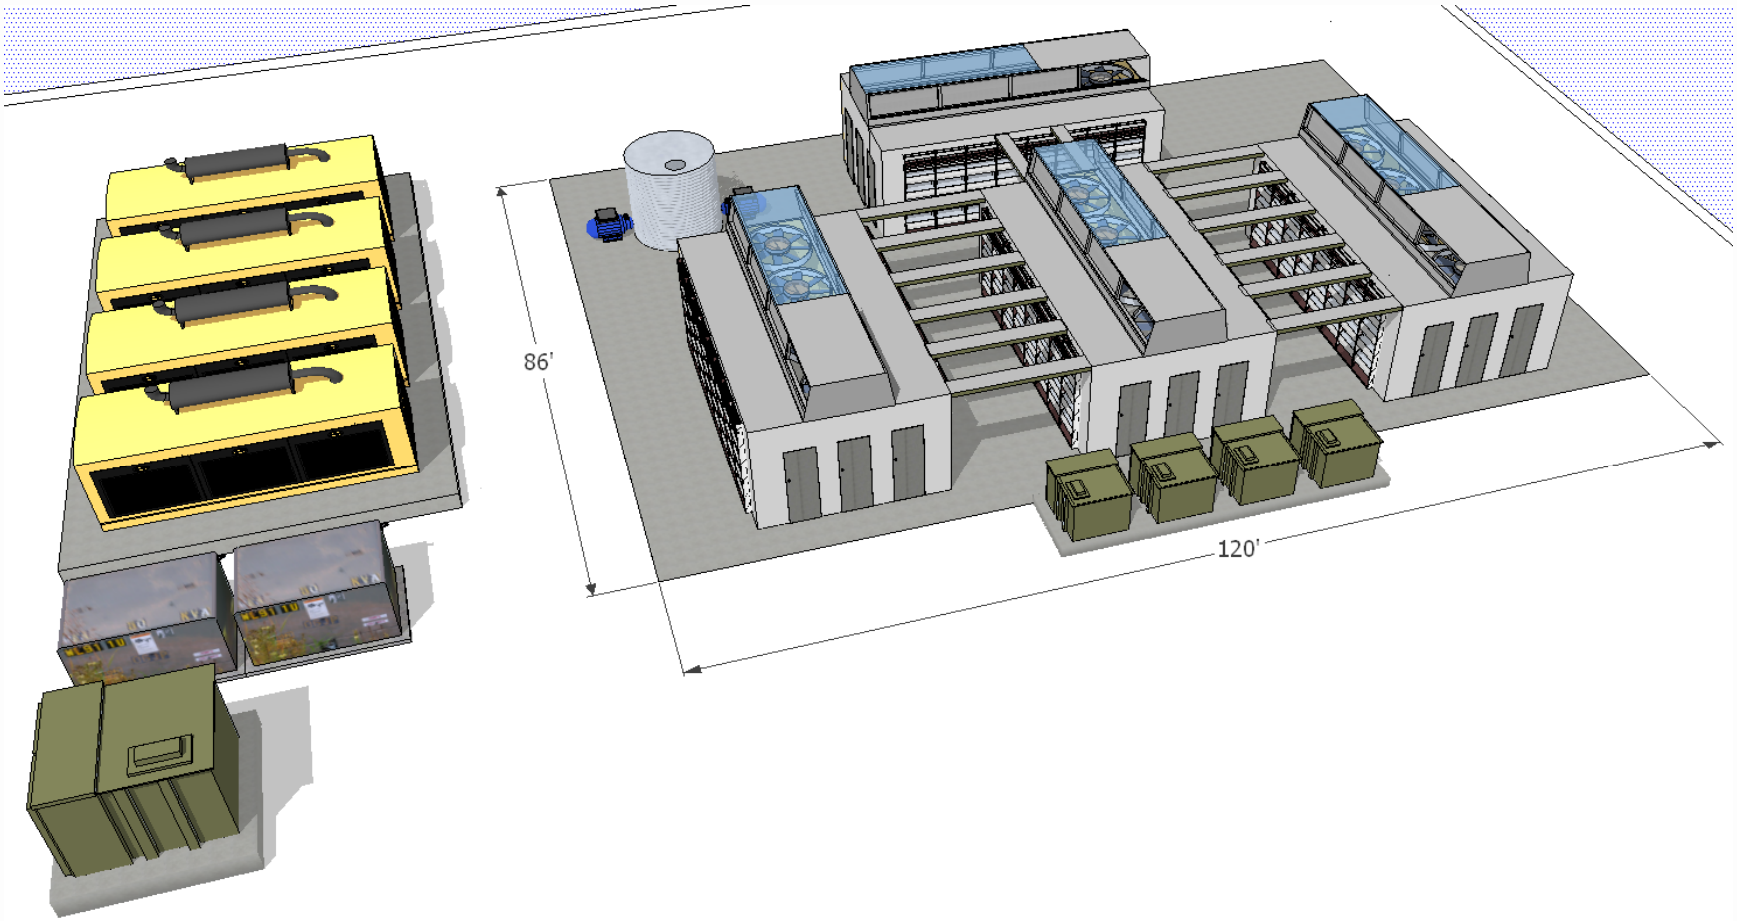

神州数码
Digital China

Cooling the Systems

- ICE Cube has cold water supply and return lines
- Fans draw air through radiators between each rack
- Air is cooled before passing through the servers
- Allows for higher water loop temp and reduced air handler power usage

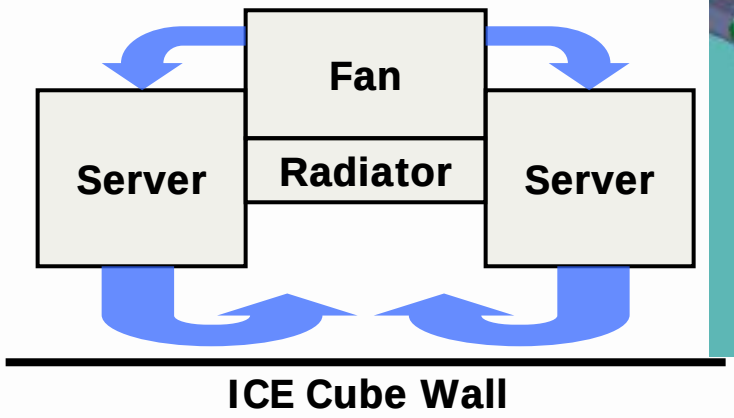

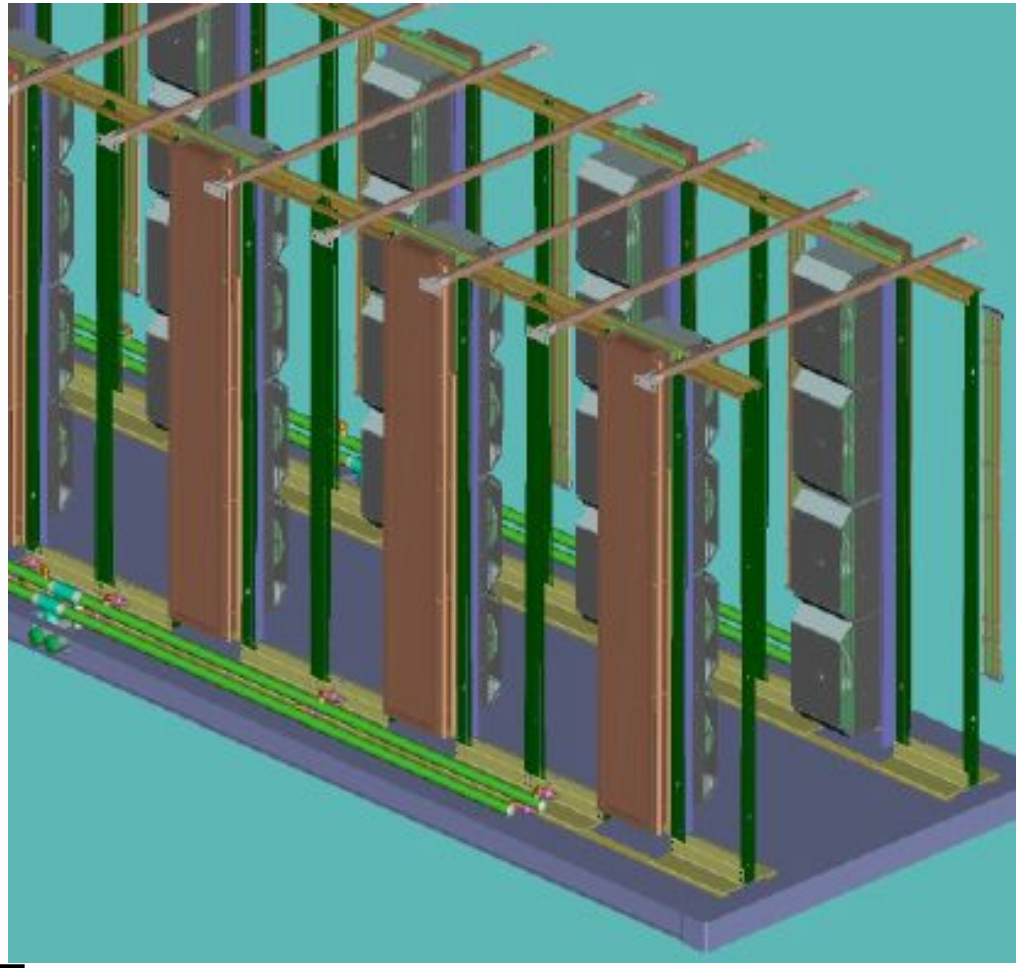

ICE Cube ISO Container

STANDARD SHIPPING
Rackable Systems' ICE Cube puts a custom-configured data center in a standard 20- or 40-foot shipping container, providing companies with an expandable, fully portable computing solution.

Thinking outside the box…with a Cube!

KEY COMPONENTS

SERVERS/STORAGE
A 40-foot container can hold 28 server racks with up to 1,400 servers or storage systems. This translates to 11,200 processing cores (using Quad-Core Intel Xeon processors) or 7.2 petabytes of storage.

COOLING COLUMN
Impeller fans circulate air through center, eliminating need for individual fans at the server level.

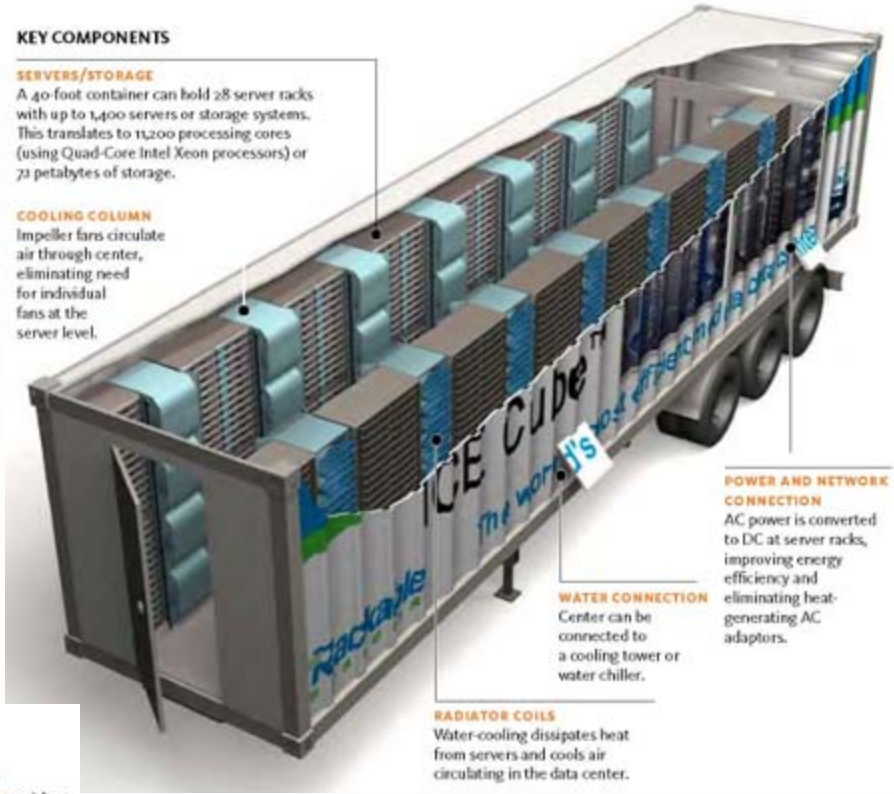

POWER AND NETWORK CONNECTION
AC power is converted to DC at server racks, improving energy efficiency and eliminating heat-generating AC adaptors.

WATER CONNECTION
Center can be connected to a cooling tower or water chiller.

RADIATOR COILS
Water-cooling dissipates heat from servers and cools air circulating in the data center.

BREATHING ROOM

CENTRAL AISLE
36-inch aisle provides easy access to systems.

VESTIBULE
Large area provides space for IT administrators to work or store gear.

神州数码 Digital China

ICE Cube Customers

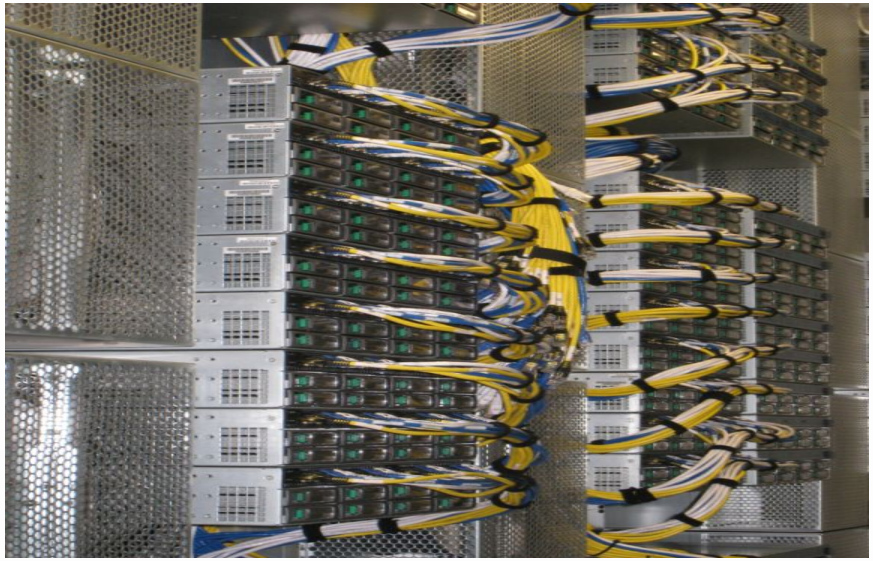

Modular Data Center Evolution

2005 2006 2007 2008 2009 2010

Raised Floor Low KW per Rack

High Density Racks – No Raised Floor

ISO Containerized Datacenters

Modular Data Centers. Container is one Component

Deployment Scale Unit

Server

Capacity

Rack

Density and Deployment

Containers

Scalability

MDC's

Time to Market

Lower TCO

More info at

www.digitalchina.com

www.sgi.com