Turn your critical data into real-time business insights

HPE Superdome Flex

Jacky Chan

Kinmax
擎昊科技

Can you capitalize on the data explosion?

Structured data **40 petabytes**	Human interaction data **4 petabytes**	Digitization of Analog Reality **40,000 petabytes a day***
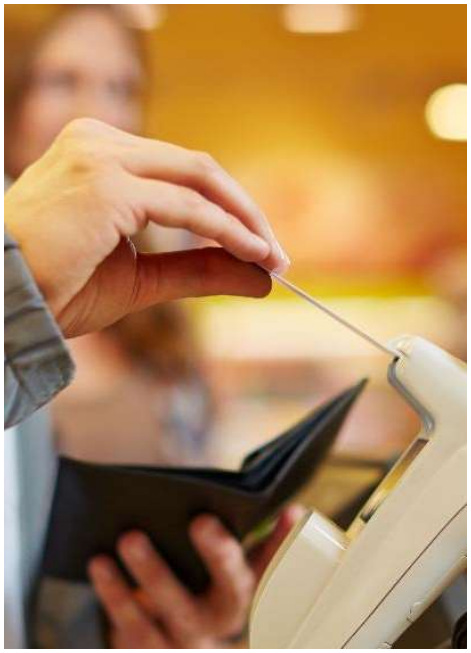		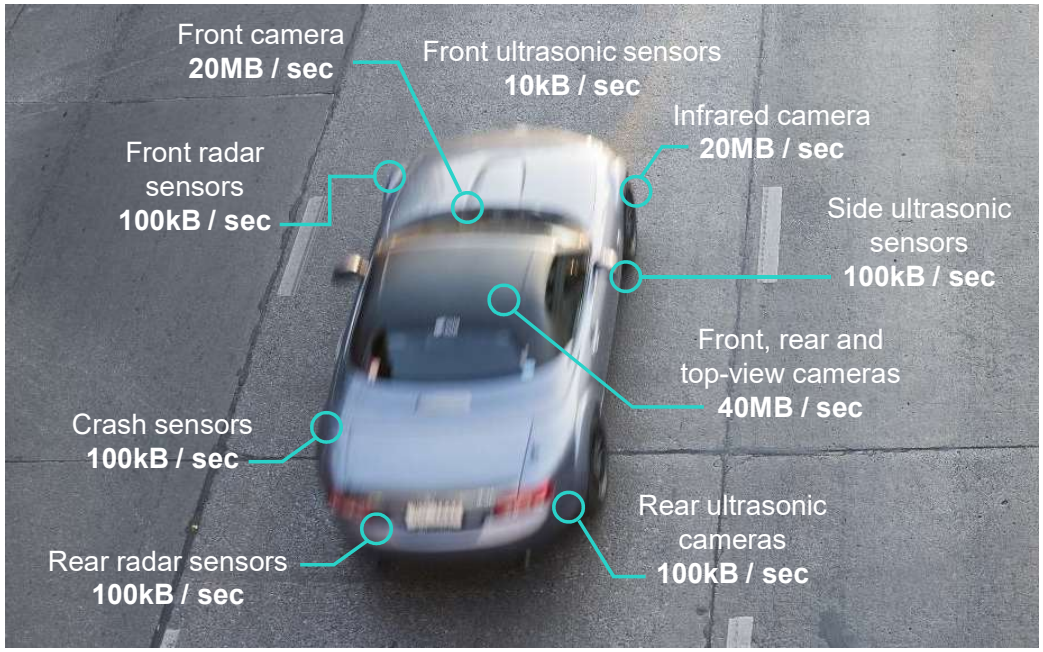
Walmart's transaction database (2017)	Per-day posting to Facebook across 1.1 billion active users (May 2016) 4kB per active user	10m self-driving cars by 2020

* Driver assistance systems only

Kinmax 擎昊科技

Actionable insights from all of your data

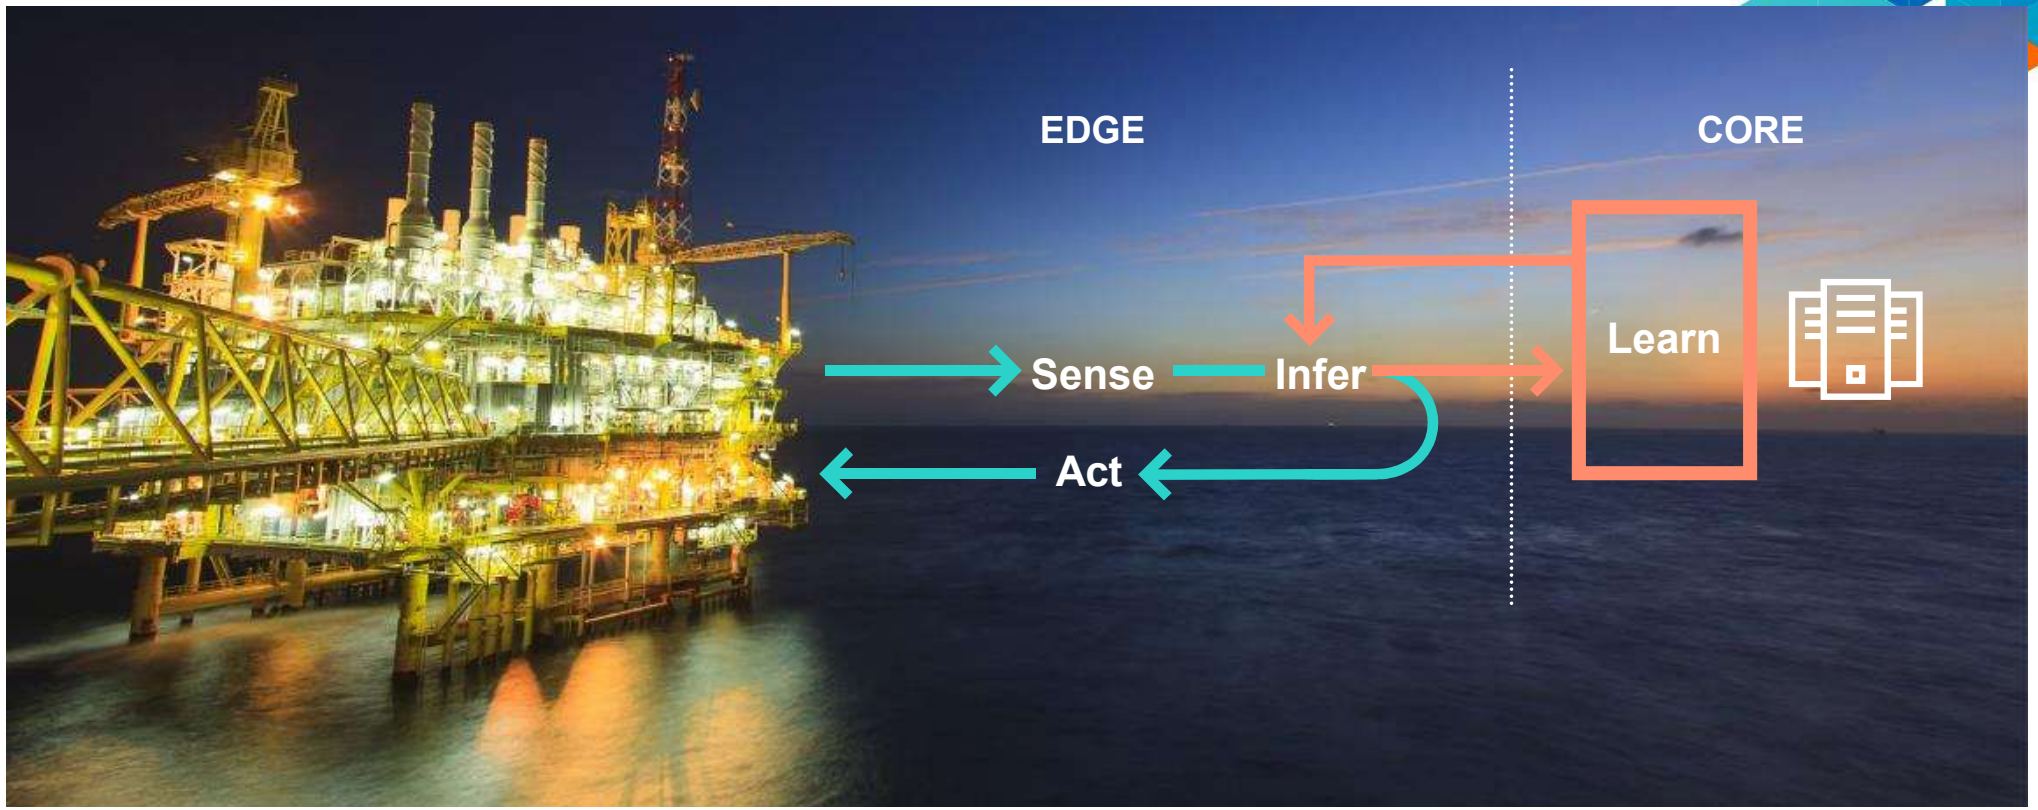

Make **time-sensitive decisions** at the edge of your network

Create a powerful digital core central learning engine

Protect and manage your data from edge to core

Kinmax 擎昊科技

Taking action in real-time demands a powerful digital core

Kinmax
擎昊科技

Harness data in real-time to transform your business

Converging transactions and analytics at the core

Accurate and instantaneous service personalization

Real-time portfolio risk estimation

Hyper-individualized patient treatment

Real-time demand/supply matching

Immediate response to security threats

Kinmax 擎昊科技

Transforming the digital core to run at the speed of memory

HPE is the leader in Memory-Driven Computing

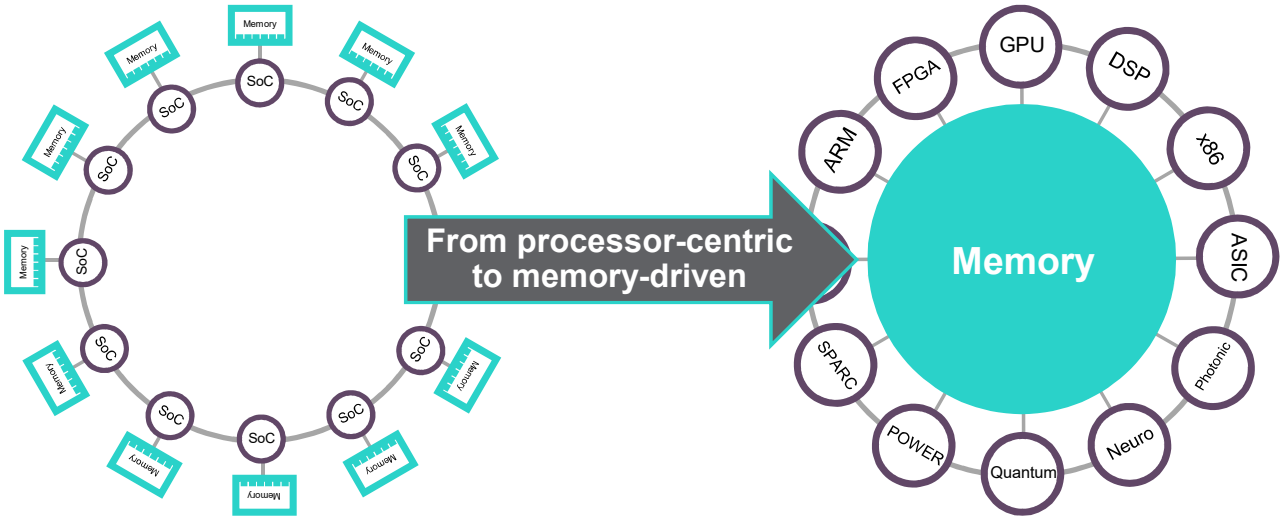

Leading database vendors transform to in-memory software

SAP HANA

ORACLE 12c DATABASE

Microsoft SQL Server 2016

Massive data with unpredictable growth, conventional systems can´t keep up

Kinmax 擎昊科技

Advancing the real-time enterprise journey

Empowering the data-driven enterprise...

SAP HANA in-memory

Project “Kraken”
Co-innovation with SAP

Extreme scalability for transactions & analytics

Superdome X

High performance Data analytics

Acquisition sgi

Platform for the next decade of in-memory computing

NEW

HPE Superdome Flex

Real-time Enterprise

Memory-Driven Computing advancements

.... Accelerating real-time business insights

Kinmax 擎昊科技

HPE Superdome Flex

Turn critical data into real-time business insights

Turn data into actionable insights in real time

- Unparalleled scale 4-32 sockets, 768GB-48TB memory
- Highly expandable for growth; ultra-fast fabric

Keep pace with evolving business demands

- Unique modular 4-socket building block, 45% lower cost at 4s entry point
- Open management for hybrid IT consumption

Safeguard mission-critical workloads

- Proven Superdome RAS with 99.999% single-system availability
- Mission-critical expertise with HPE Pointnext services

Designed with Memory-Driven Computing principles

Kinmax 擎昊科技

World’s most scalable and modular in-memory computing platform

HPE Superdome Flex

Scales up to 8 chassis and 32 sockets as a single system in a single rack

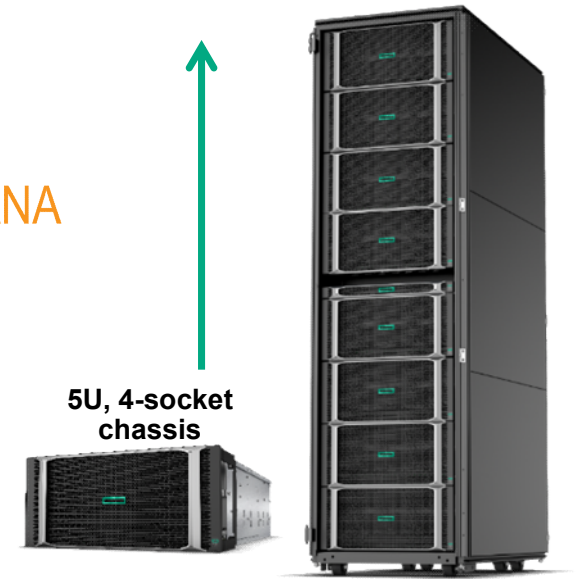

5U, 4-socket chassis

SAP HANA

SUSE ORACLE LINUX KVM vmware

Unparalleled Scale
- Modular scale-up architecture
- Scales seamlessly from 4 to 32 sockets as a single system with both Gold and Platinum processors
- Designed to provide 768GB-48TB of shared memory
- High bandwidth (13.3GB/sec- bi-directional per link)/low latency (<400ns) HPE Flex Grid
- Intel ® Xeon® Scalable (Skylake) processors with up to 28 cores

Unbounded I/0
- Up to 128 PCIe standup cards, LP/FH PCIe

Optimum Flexibility
- 4-socket chassis building blocks, low entry cost; HPE nPARs
- Nvidia GPUs, **Intel SDVis**
- 1/10/25 Gbe, 16GbFC, IB EDR/Ethernet 100gb, Omni-Path
- SAS, Multi-Rail LNet for Lustre; NVMe SSD
- MPI, OpenMP

Extreme Availability
- Advanced memory resilience, Firmware First, diagnostic engine, self-healing
- HPE Serviceguard for Linux

Simplified User Experience
- HPE OneView, IRS, Openstack
- HPE Proactive Care

intel XEON PLATINUM inside

Kinmax 擎昊科技

What’ s new with HPE Superdome Flex?

Support for hard partitions (HPE nPars)	– HPE nPars support, 48 different configurations – Scale from 4 to 32 sockets and 768GB to 48TB memory – 128GB DIMMs pre-enabled
Enhanced ecosystem	– Certified with Windows 2016 – HPE OneView 4.1 Monitoring – HPE Serviceguard for Linux and HPE WASL enhancements
Additional processors	– Seven new Intel Xeon Scalable processor options including three Platinum and four Gold options
Increased I/O options	– Nvidia Quadro P6000 Accelerator – HPE 2.4TB SAS 12G 10K SFF RW DS HDD

Kinmax
擎昊科技

C Telecom

First HPE Superdome Flex win in Asia – A Taiwanese Telecom

Deal Overview
First HPE Superdome Flex win in Asia – leading Taiwanese Telecom

4 socket testing system for Unix to Linux migration project

Country
Taiwan

Industry
Telecommunications

Application/Use Case
Telco billing
Unix to Linux migration

HPE Solution
HPE Superdome Flex
4 socket, 768GB memory
HPE Pointnext services

Customer Challenges

– C Telecom is undergoing a Unix to Linux migration project for their critical in-house telco billing system, currently running with an Informix database

– They were looking for a reliable standard 4-socket server for the migration testing system

Winning Strategy

– Position Superdome Flex immediately after launch to meet customer requirements for reliability and scalability

Why HPE and HPE Superdome Flex?

– HPE Superdome Flex offered the reliability capabilities and room for growth required by the customer at an entry point that was competitive and within budget

– HPE had delivered a reliable solution for years and the customer trusted HPEs capability to continue delivering a stable solution for their critical billing application

Account Team Tips

Why did we win?
HPE Superdome Flex has strong product features at the right cost structure. In addition, HPE´s long term trusted relationship with the customer, allowed us to win the deal.

What did we learn?
Use HPE Superdome Flex entry point to satisfy the needs of customers requiring high RAS at a competitive initial acquisition cost

Who did we beat?
All x86 providers

Kinmax
擎昊科技

如何發揮in-memory Computing 的所有價值

Kinmax
擎昊科技

Harness the full value of in-memory computing

HPE Superdome Flex Use Cases

SAP HANA

- SAP non HANA to SAP HANA
- SAP BW/4HANA
- SAP S/4HANA
- SAP HANA on Cloud

- Market leadership, over 12K systems shipped
- Unrivaled HANA portfolio
- SAP co-innovation and expertise

ORACLE

- Reduce licensing costs
- Unix migrations
- Exadata replacement
- Mix OLTP and analytics

- Over 140K Oracle enterprise customers run on HPE
- #1 in scale-up capacity for Oracle DB in-memory
- Consolidation, migration svcs

Microsoft SQL Server

- Mission critical SQL Server workloads
- Scale-up SQL Server
- SQL Server on Linux

- Co-engineering with Microsoft
- Broadest SQL Server solutions
- Over 29K joint partners

In-memory HPC

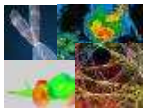

- Genomics
- Computer-aided engineering
- Computational chemistry
- Cyber security
- Financial risk management
- Fraud prevention
- Large data visualization

- Innovation leadership with SGI and HPE IP
- 20+ yrs in-memory experience
- Deep workload expertise

Kinmax 擎昊科技

Accelerate your journey to SAP HANA with HPE Superdome Flex

Innovation for the most demanding mission-critical SAP HANA use cases

Modular, 4 socket building blocks let you add capacity without forklift upgrades and easily convert from scale out to up without re-platforming

Scalable Flexibility that enables:

- Infrastructure that is right sized for you
- Frictionless SAP HANA adoption journey
- Agility to react to evolving business needs

Investment Protection for:

- Optimized infrastructure spend
- Seamless capacity upgrades
- Ability to repurpose and redeploy infrastructure

Improved cost of ownership through:

- Easier manageability
- Reduced Datacenter Footprint
- Flexible architecture

Kinmax 擎昊科技

Scale up Oracle with Superdome Flex

Reduce costs and complexity for enterprises of any size

Right-size your Oracle environment

- Mission-critical solution starting at 4/8 sockets, simply scale up as needed
  - One server to manage
  - Supports Oracle Linux and OVM
- UNIX migration – cut Oracle licensing costs **up to 50%**
- Scale-up compute and get 99.999% HA – **avoid cluster latency & costly Oracle RAC fees**
  - Exadata replacement
  - **Improve performance without tuning** or proprietary software

Become a real-time enterprise

- **Take action using current data** to transform your business
  - Eliminate ETL delays for real-time analytics
- **Unmatched scale-up capacity** for *Oracle in-memory*
  - In-memory workloads run faster on fewer processors, reducing Oracle database license costs
- Scale massively, to handle ever-large data sets
  - **Unlock the value in the Intelligent Edge**
- +140K Oracle customers run on HPE

Quantify bottlenecks with a no-charge Oracle Performance and Cost Assessment
to reduce license and support costs

Slash Oracle licensing, free up budget needed for innovation

Solve your most demanding SQL Server challenges

Gain real-time insights on your operational data

Microsoft SQL Server

Modernize business processing

Accelerate analytics

Optimize infrastructure

Scale-up compute for your largest workloads

- Avoid cluster latency
- Support high transaction rates on critical OLTP applications

Breakthrough in-memory performance

- Accelerate transactions
- Leverage current transaction data for real-time analytics

Right size for every mission-critical workload

- Start at 4-sockets, and grow with your needs
- Gain tremendous savings from Oracle replacement

Avoid inefficiency, cost and complexity

- Simplify managing your environment by consolidating hundreds of SQL Server instances onto a single serve

Kinmax 擎昊科技

Solution: HPE In-memory High Performance Computing

HPE Superdome Flex provides globally shared memory with seamless scale up capacity

HPE Superdome Flex Server

- Solve complex, data-intensive problems **holistically** at **unparalleled** scale with **single-system** simplicity
- Complete **more jobs in less time**
- Free HPC teams from managing clusters to **accelerate time to discovery**

Use cases include:

- Genomics
- Bioinformatics
- Computer-aided engineering
- Computational chemistry
- Cyber security
- Financial risk management
- **Fraud detection and prevention**
- Real-time, streaming graphs
- Large data visualization
- In-memory databases

Kinmax 擎昊科技

United States Postal Service – real time fraud detection

UNITED STATES POSTAL SERVICE

Kinmax 擎昊科技

Sample of fraud detection customer systems

Large memory UV systems

Site	System	Cores	Memory
U.S. Postal Service	UV2000	3,328	64TB
PayPal	UV300H	756	48TB
U.S. Postal Service	UV2000	2,304	48TB
U.S. Postal Service	UV2000	2,304	48TB
PayPal	UV300H	756	48TB
PayPal	UV300H	120	24TB
Ebay	UV2000	96	6TB
Ebay	UV300H	120	6TB
U.S. Postal Service	UV2000	64	2TB

滿足從critical data 轉變成real-time insights 所需要的效能及可擴充能力

Kinmax
擎昊科技

Superdome Flex Chassis Front View

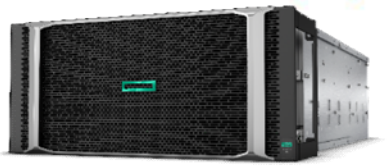

4 Hot-swap Cooling Fan Assemblies (8 x 80mm) N+1 redundancy

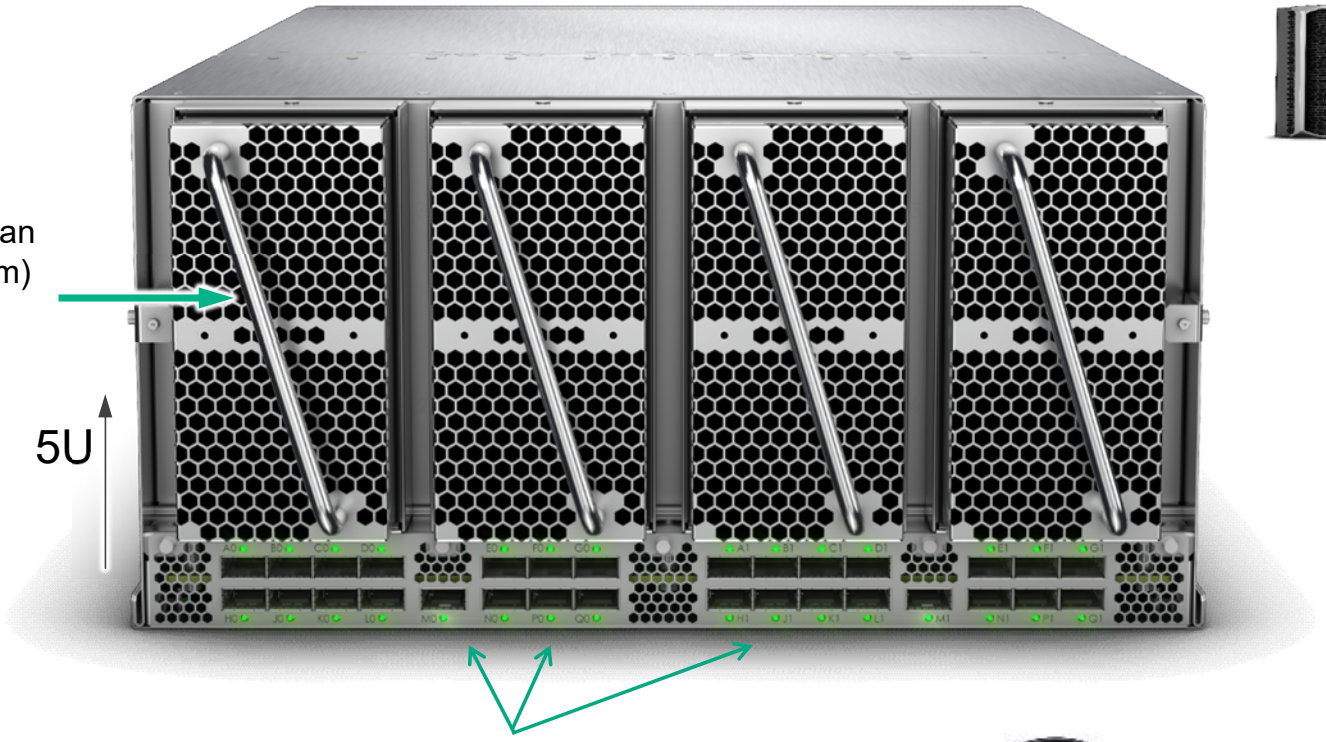

30 HPE Fabric Cable Ports

- Infiniband (EDR) electrical interface and cabling (not an Infiniband network)
- Fault tolerant if cable removed

Kinmax 擎昊科技

Superdome Flex Chassis Rear View

Base chassis pictured

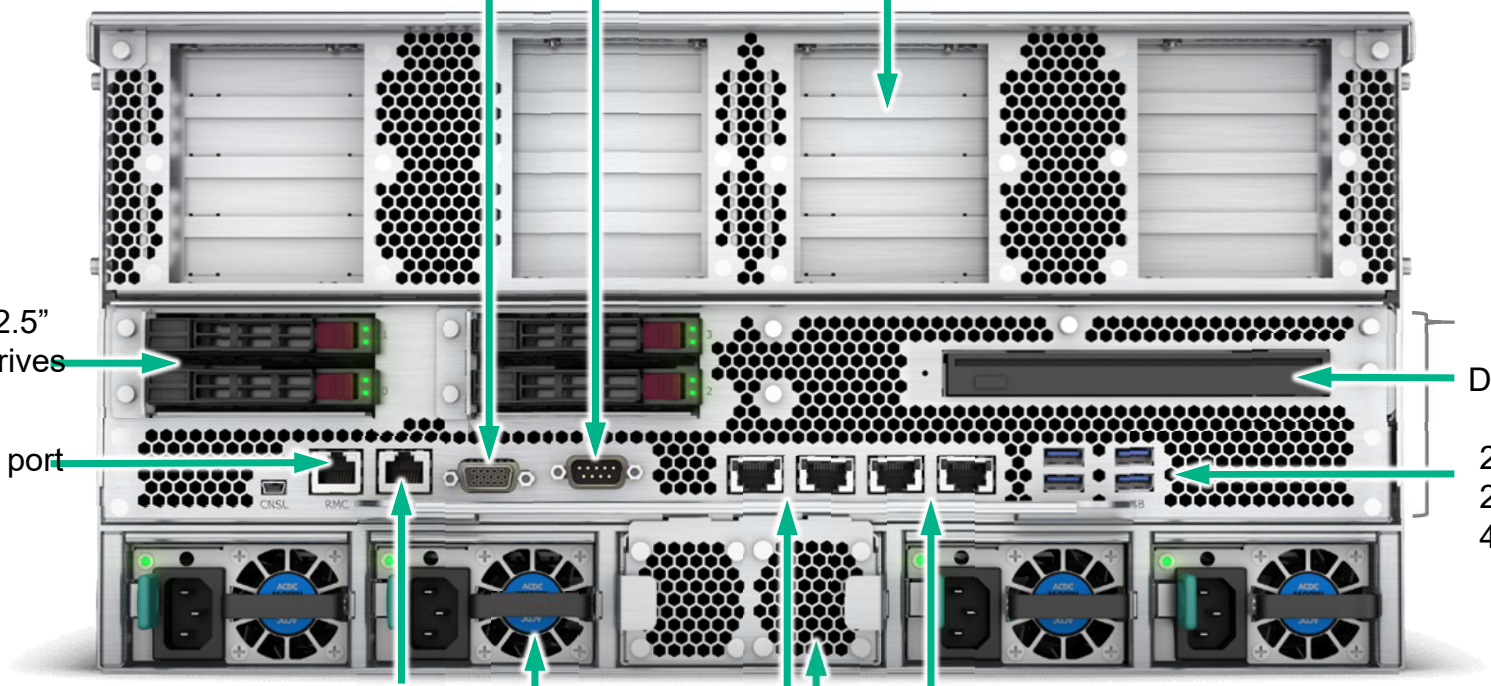

* For Expansion chassis, the Base I/O area is replaced with single RMC, Mgmt, and BMC ports

Kinmax 擎昊科技

HPE Superdome Flex: Chassis Internal Top View

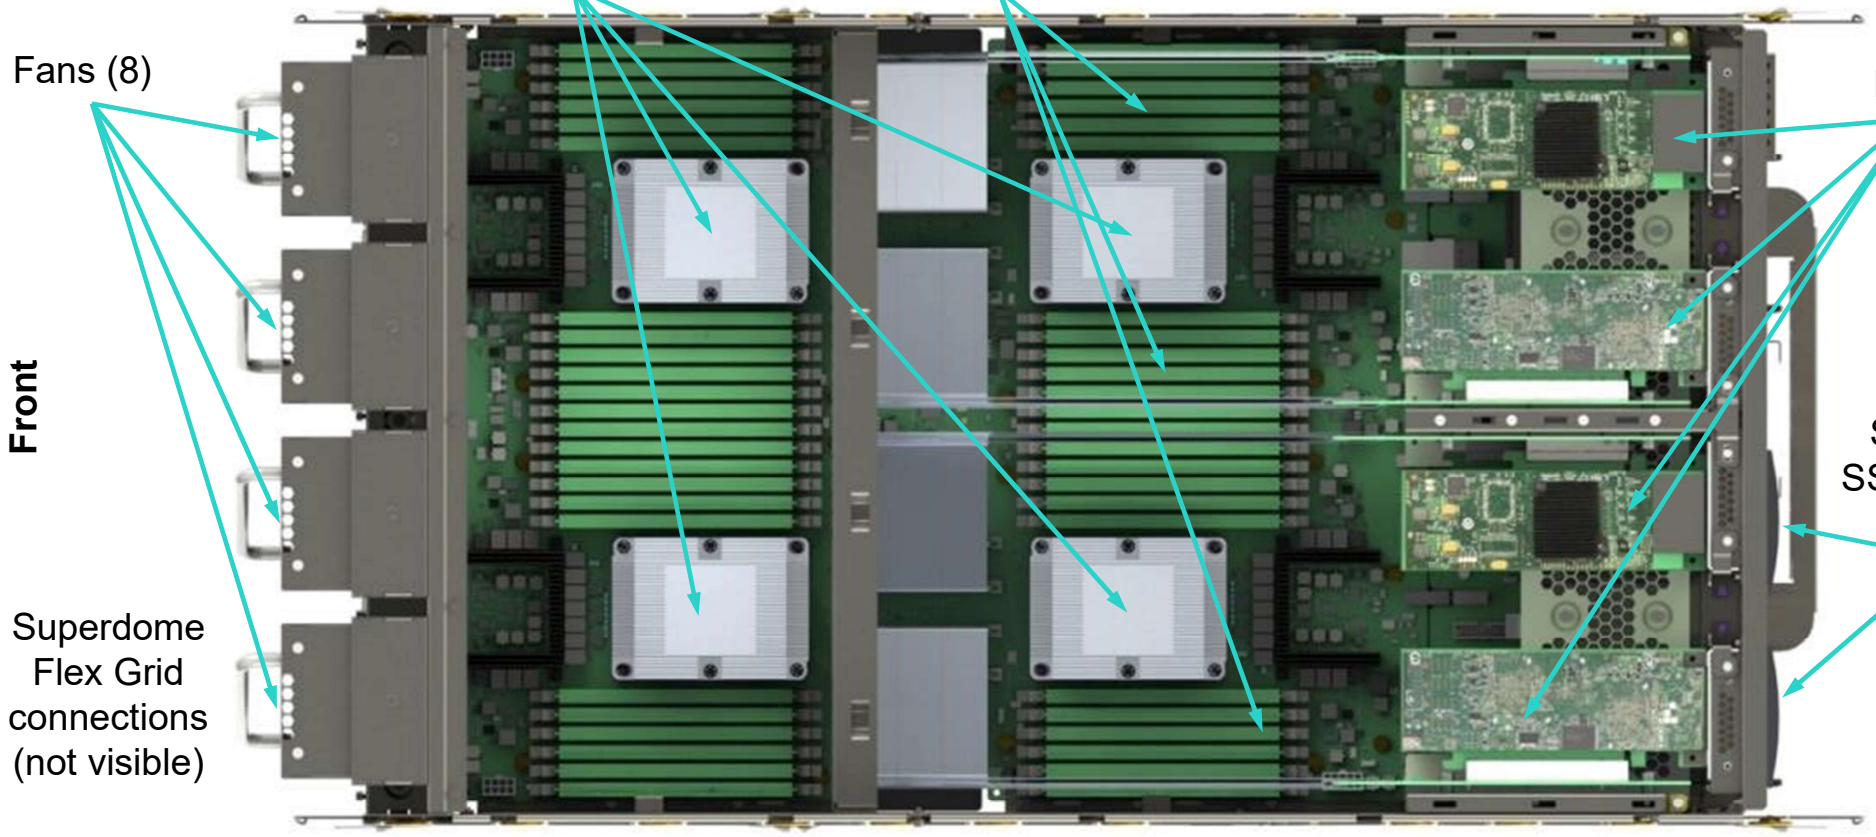

System rails allow chassis to slide forward and backwards in the rack for easy service access

Turn massive amounts of data into business-fueling insights

– Scale easy and economically, regardless of your business size

– Start small and grow seamlessly at your own pace

– Avoid over-provisioning and disruptive upgrades

– Add compute power without sacrificing performance

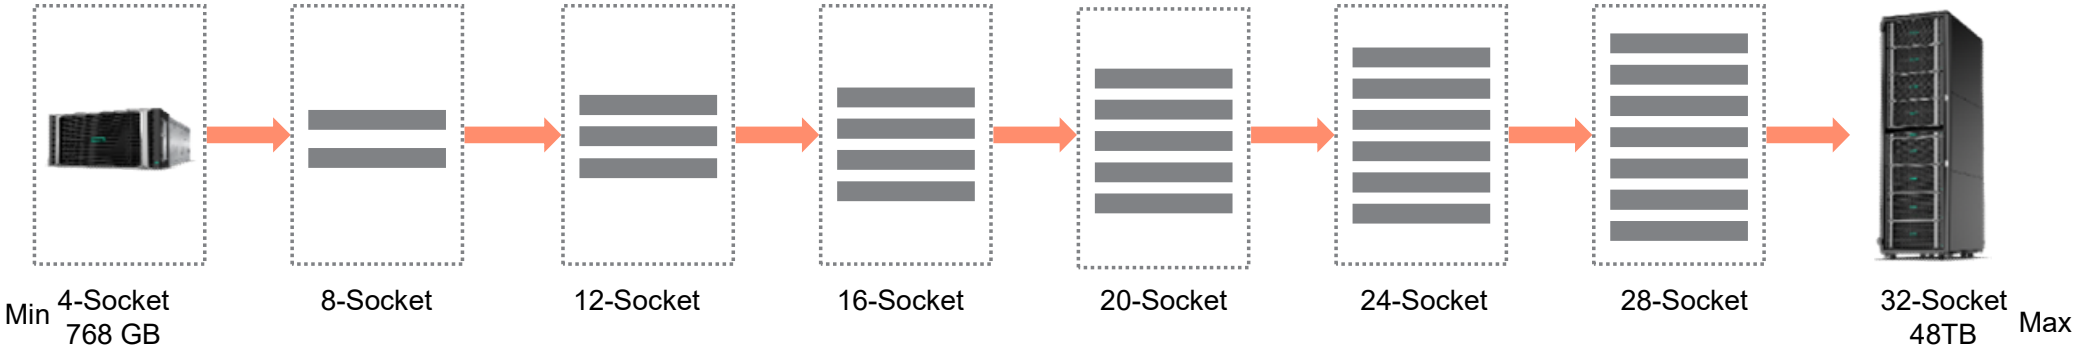

One modular building block, one system, one architecture

HPE Superdome Flex Technology

HPE Superdome Flex ASICs

- Scale beyond capabilities of Intel UPI links

Global memory

- All memory is shared by all processors

Hardware cache coherency

- Copy consistency at ASIC speed and directory-based

High bandwidth – low latency

- 13.3 GB/second; max read latency <400ns!

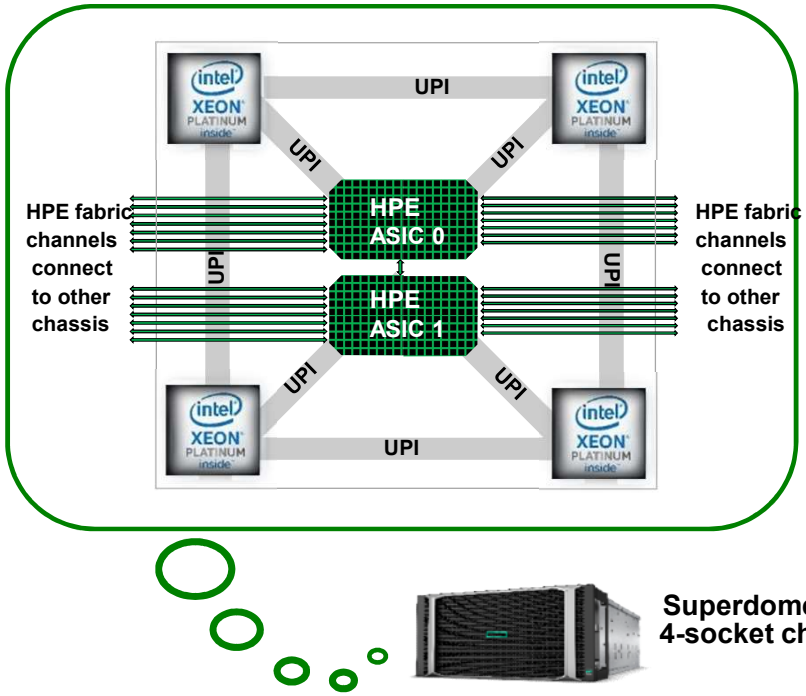

Kinmax 擎昊科技

Unique point-to-point design maximizes performance

Extreme processing speed at scale

Scales from 4 to 32 sockets

- Supports up to 8 chassis, with 4 sockets per chassis
- Includes 16 Flex ASICs
- Point-to-point, ‘all-to-all’ Grid link between system ASICs – unique in the industry
- Lower latency and increased Bandwidth over previous solutions, and competitive systems - delivering extreme performance

Compute resources provided

- 32 sockets
- 384 DIMM slots: up to 48 TB with 128 GB DIMMs
- 128 PCIe Gen3 card slots (56 x16, 72 x8) maximum

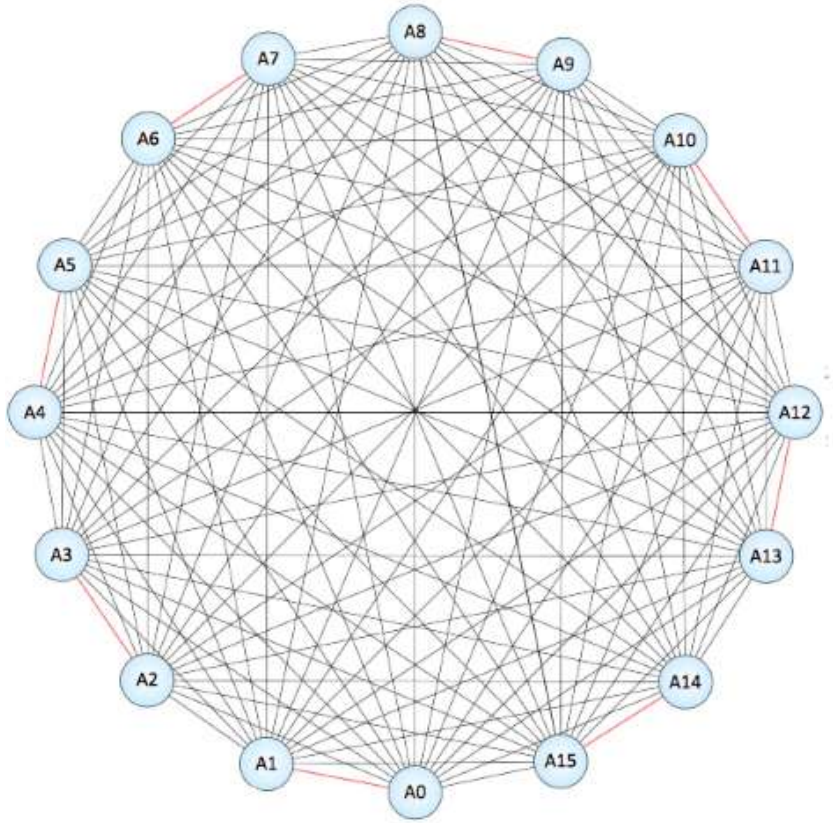

Designed for the future: Changing the rules of what’s possible, today and tomorrow

Kinmax 擎昊科技

Differentiation1: SD Flex Uniquely Scalable Architecture

Technology highlights:

- SD Flex ASIC
- 4S building block
- Grid architecture

4S building block (6TB)

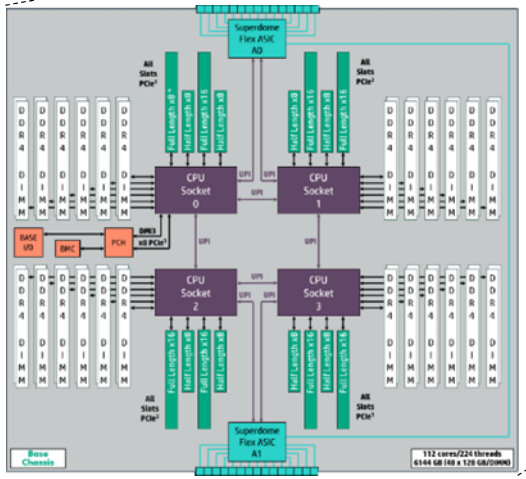

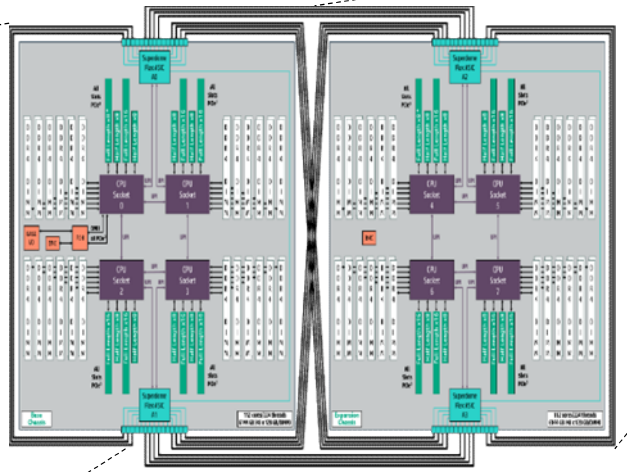

8S architecture (12TB

32S architecture (48TB)

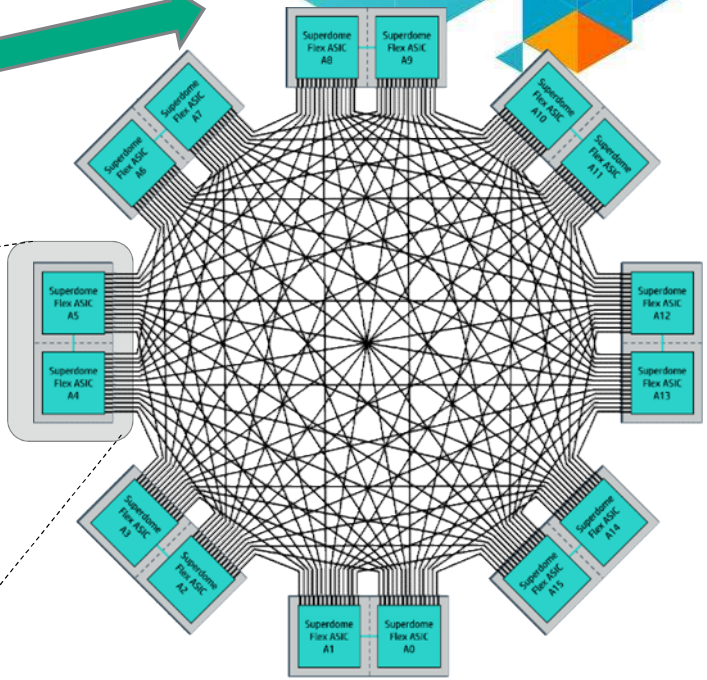

Customer's benefits:

- Meeting most demanding Ent. workload
- Fastest for real-time analytics
- Economical benefit of Gold/Platinum CPU choice
- Low 4S entry point with full investment protection

Kinmax 擎昊科技

Differentiation2: Uncompromised IO Architecture

Technology highlights:

- Fully utilizing 3 independent stacks of x16 PCIe 3.0: (2) x x16 + (2) x x8 (except base chassis with base IO)
- Innovative PCIe 3.0 riser cards and Z-axis approach
- Choice of 16 lots and 12slots IO subsyste

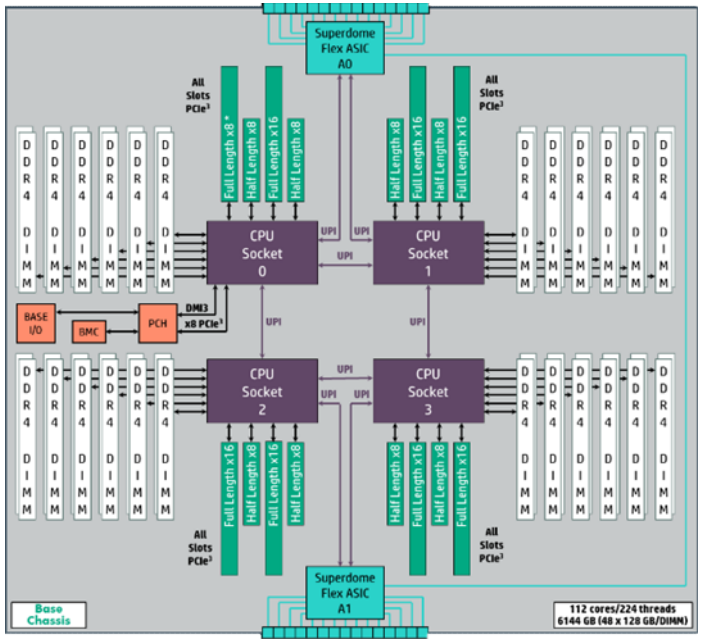

Physical view of riser cards

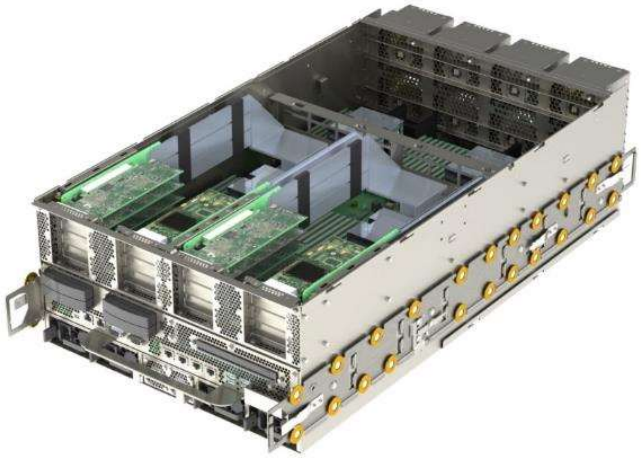

16slots chassis layout

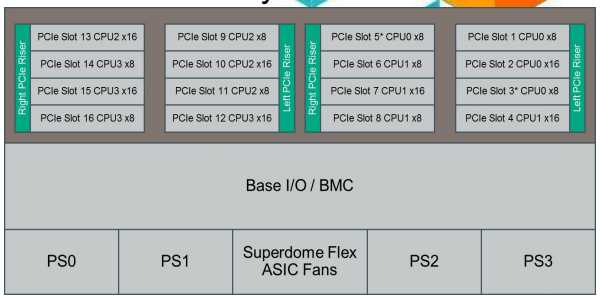

12slots chassis layout

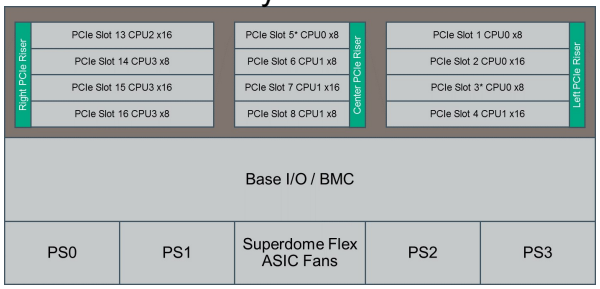

Customer's benefits:

- Unblocking performance
- Low latencies
- Flexible choices of IO cards
- Accommodating enterprise IO workloads / consolidation

Kinmax 擎昊科技

Superdome Flex and GPU opportunities

Deep Learning

16GPU

8GPU

PCIe

SXM2

- ✓ Up to 16 Tesla GPUs installed with a single OS
- ✓ Docker Deep Learning Framework
- ✓ Analysis can be done in the system without data movement

Analytics and inference where real time cross correlation matters

Multiple data streams plus historical data in memory
Compute + GPU + powerful I/O parallel data ingest

Examples such as Fraud, intelligence, compliance, tracking subjects across multiple video feeds..

4x Nvidia Tesla V100

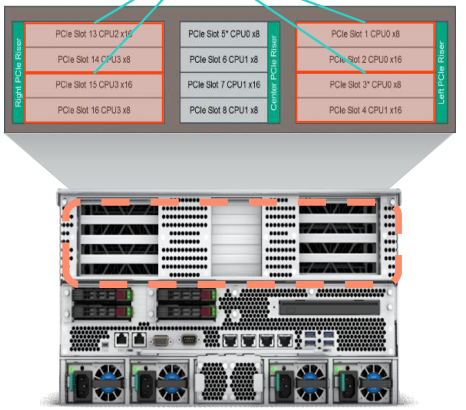

Kinmax 驚昊科技

HPE Superdome Flex Enabled Skylake Processors

System Config #1
Requires Platinum 81xx SKUs

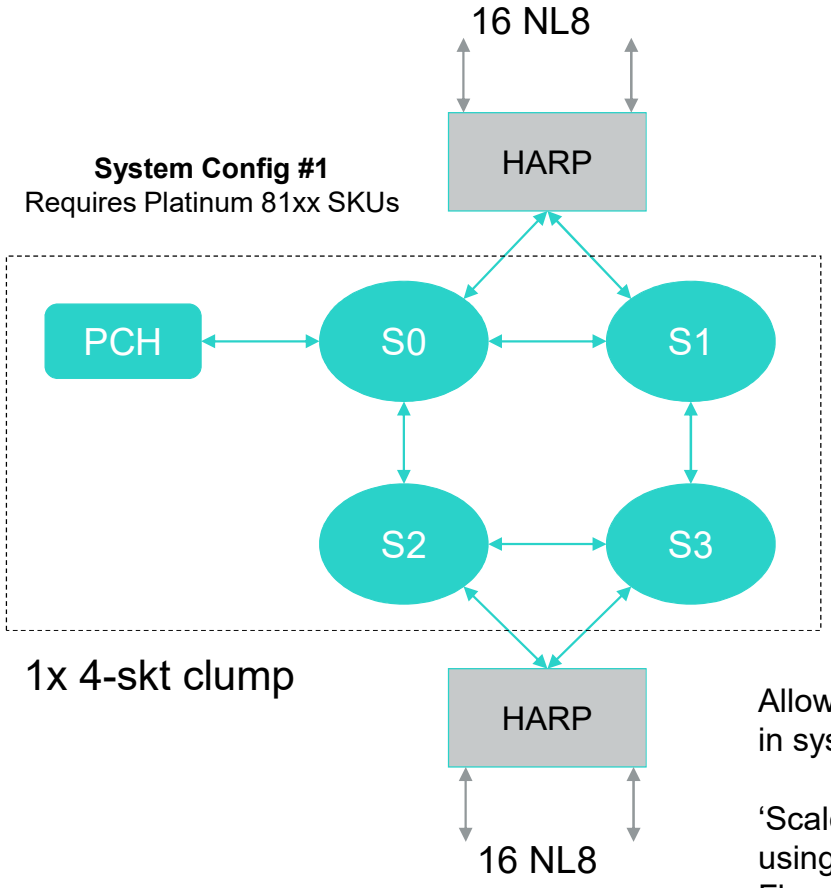

1x 4-skt clump

HPE Value-add!!

System config #2
'Auto-enables' both
Platinum 81xx and Gold 61xx

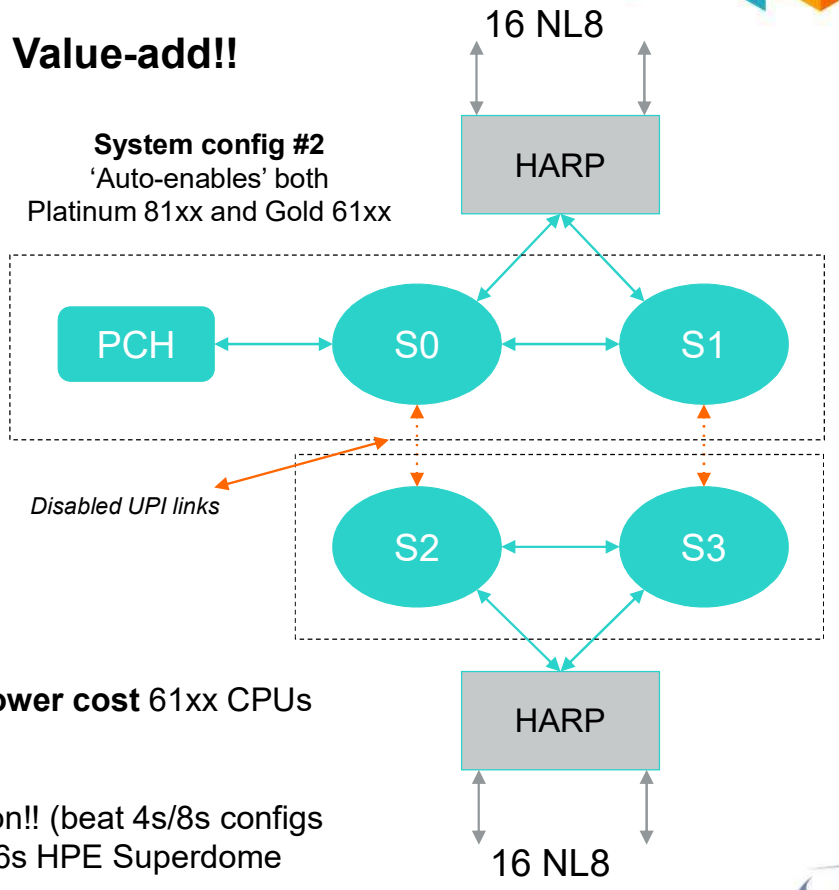

Allows the use of **much lower cost** 61xx CPUs in systems up to 32s.

'Scale past' the competition!! (beat 4s/8s configs using lower cost 8s/12s/16s HPE Superdome Flex configs)

Kinmax 擎昊科技

What is NUMA?

- NUMA: Non-Uniform Memory Access

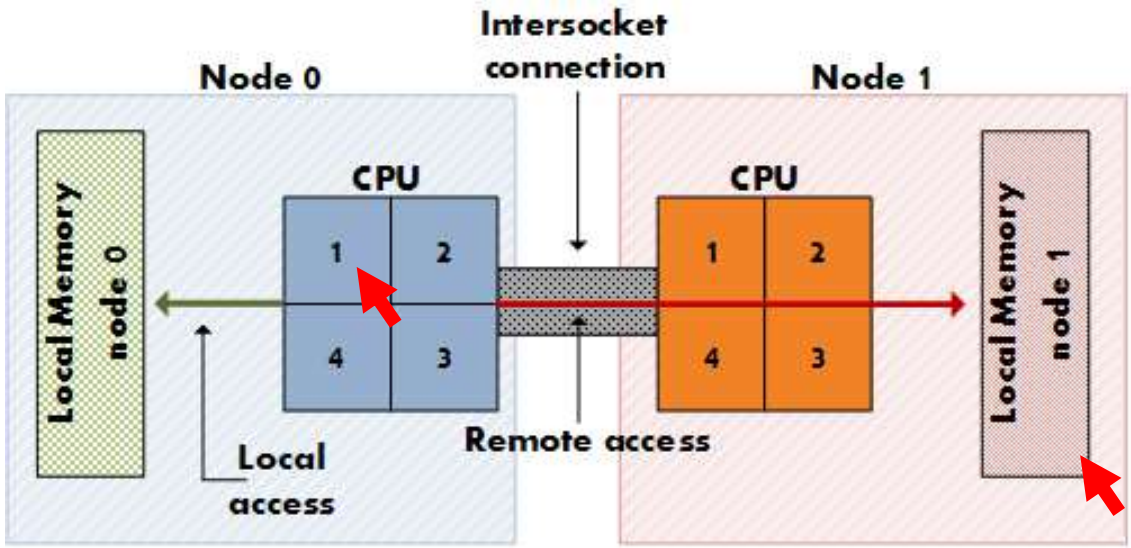

- 2-Socket x86 NUMA topology
- Superdome X
  conceptual NUMA topology
- **50% of the memory is node local memory**
- IVB-EX and Haswell-EX have a **1.6x memory latency** when accessing remote memory

- CPU 1 from Node 0 (blue) can access Local Memory (green) directly
- To access Node 1 memory the request must travel through an inter-socket connection
  - This extra connection or “hop” causes longer memory access latencies for remote memory
- Scheduling on a core from the node containing the process memory results in the best performance

4-socket Conceptual NUMA Topology

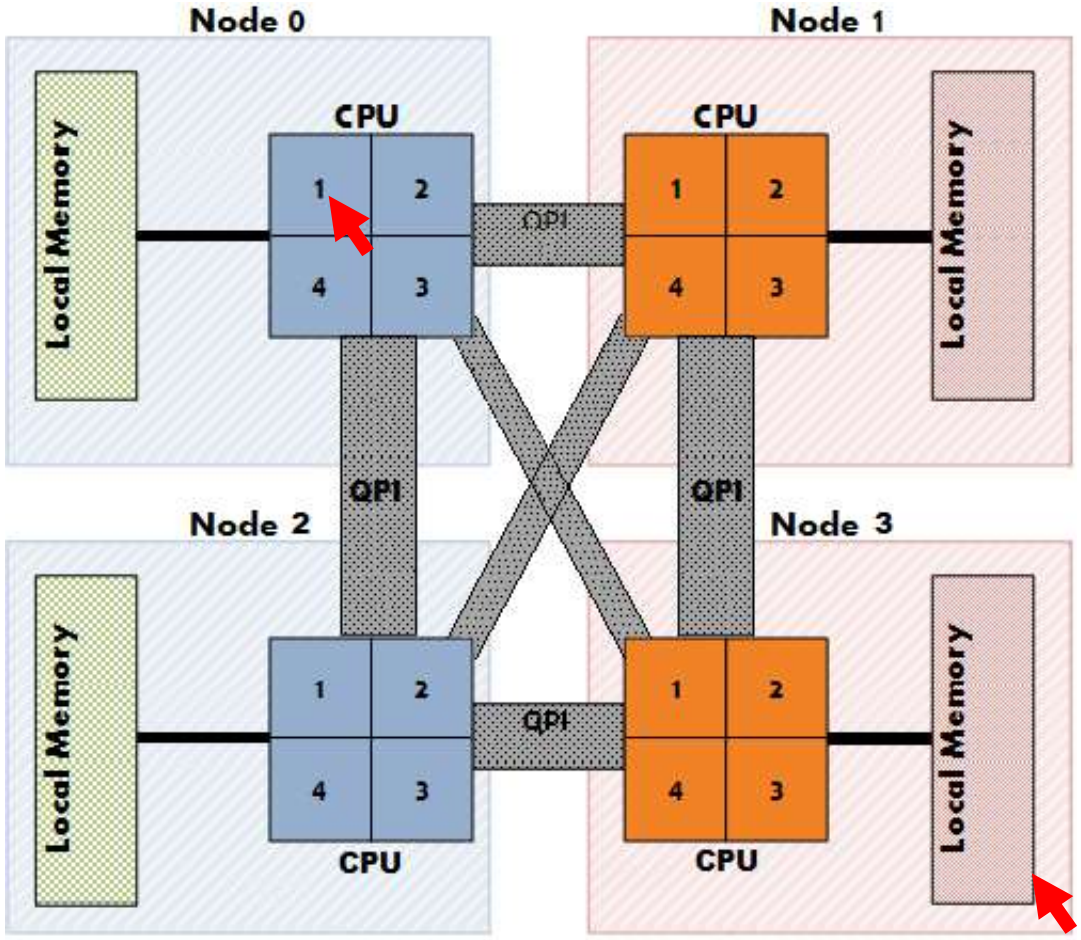

- 4-Socket Sever NUMA topology
- **25% of the memory is node local memory**
- IVB-EX and Haswell-EX have a **1.6x memory latency** when accessing remote memory

8-socket Superdome X Conceptual NUMA Topology

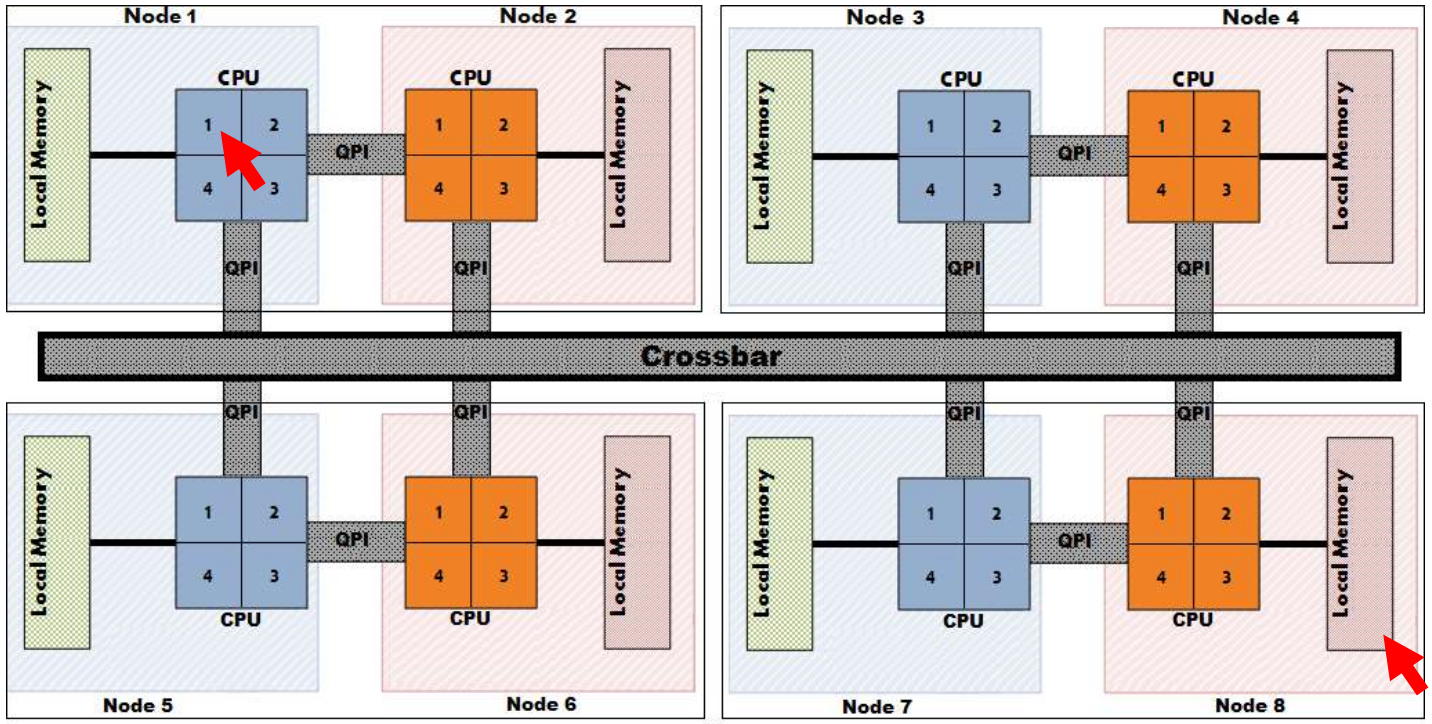

- 8-socket Superdome X conceptual NUMA topology
- **12.5% of the memory is node local memory**
- IVB-EX and Haswell-EX have a **1.6x memory latency** when accessing remote memory on-blade (buddy socket memory)
- Remote off-blade memory access has a **3.0x memory latency**

Kinmax
擎昊科技

What is ATX (HPE Application Tuner Express)?

- ATX is a utility that will make NUMA unaware applications more NUMA aware
  - No application changes are needed!
- ATX controls the distribution of an application's processes and threads in a NUMA environment
  - Several different NUMA node and CPU launch policies are provided to obtain an optimal distribution.
- ATX is similar to the Linux `numactl` command:
  - `numactl` will constrain an application to a set of nodes
  - ATX will distribute an application around a set of nodes
- Benefit of ATX varies by platform and application
  - Higher socket count platforms benefit more than lower socket count platforms
  - NUMA-aware applications benefit less than applications without NUMA awareness

Increase the performance of your Linux workloads

HPE Application Tuner Express

Aligns data in memory close to the computing core resulting in increased performance

Real-world results

- –20% improvement with Oracle DB for a Manufacturing Data Warehouse workload
- –A Major Airline saw a 33% response time improvement for online seat reservations
- –Cerner: 10% response time increase for a custom application on Oracle

Internal testing

– SAP ERP workloads with greater then 50% performance improvement

– Up to 60% scale-up database improvement

- Large scale up servers running Oracle database gained between 30 and 59% (OLTP)
- 20% improvement with Data Warehouse (OLAP)

– 30% performance gain on Apache SPARK on 2-socket server cluster

HPE Superdome Flex Server sets #1 and #2 records on SPEC CPU2006 benchmark for 32 and 16 sockets

Key takeaways

- **First published SPEC CPU2006 results with 32 and 16 Intel Xeon Scalable processors**
- **SPECint_rate_base2006:**
  - **#1 and #2 32S results**
  - **#1 and #2 16S results**
- **SPECfp_rate_base2006:**
  - **#1 and #2 32S results**
  - **#1 and #2 16S results**

Top five 32S and 16S SPECint_rate_base2006 results

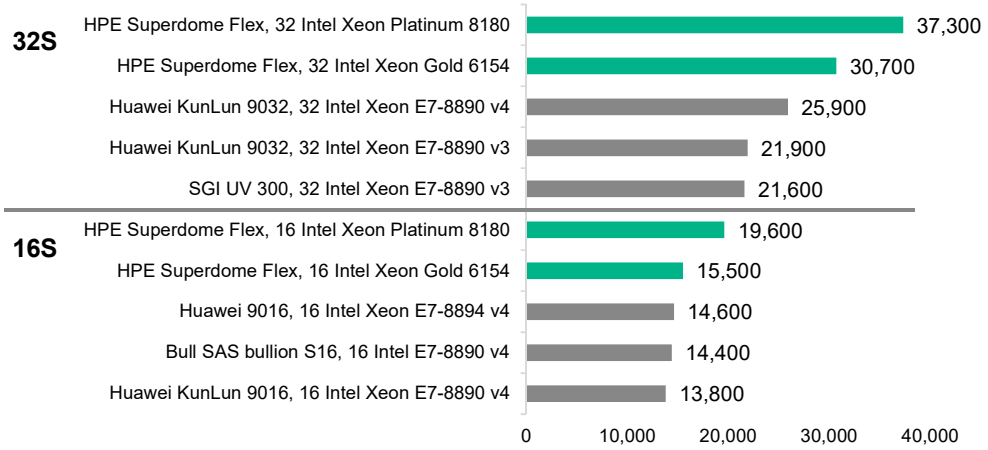

Top five 32S and 16S SPECfp_rate_base2006 results

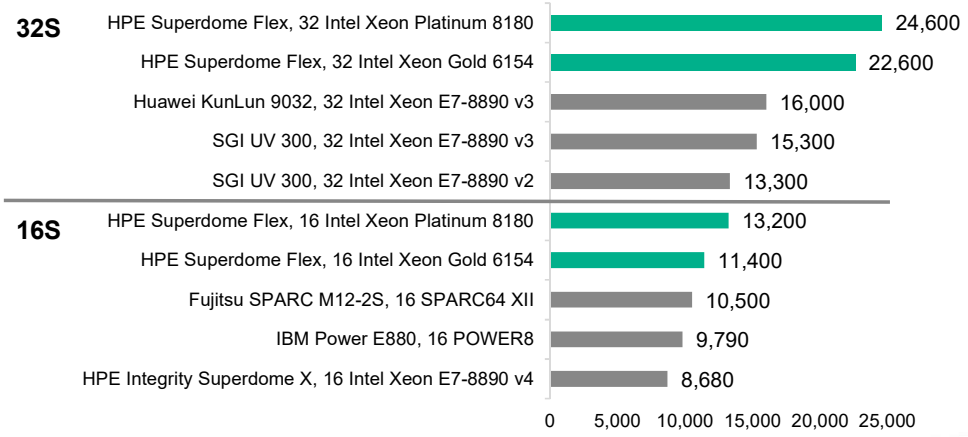

針對Apps及公司的重要系統
提供彈性的compute power

Kinmax
擎昊科技

Flexibility to scale, configure, convert, deploy

The possibilities are endless

1 Meet your individual workload requirements

4-Socket

Configuration choices:
- Memory size and capacity
- Processor and core count
- Amount and type of I/O

”Building blocks” configuration choice

2 Grow seamlessly with your business demands

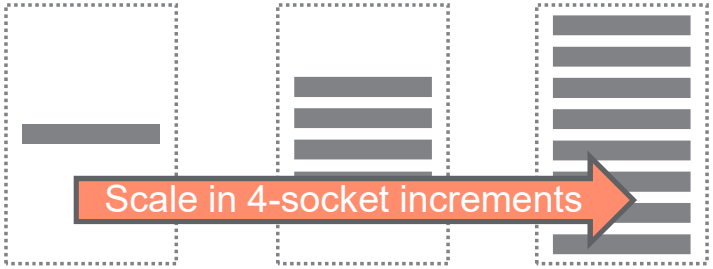

Never outgrow, avoid over-provisioning: 4-32 socket scale

3 Easily transform your environment

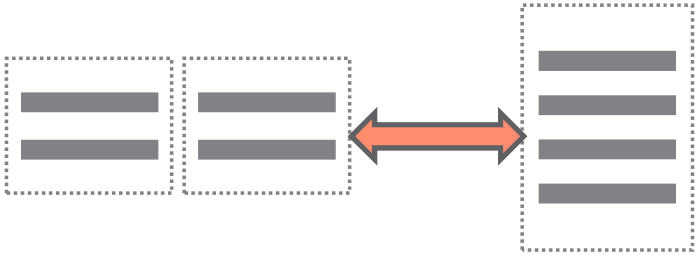

Convert between scale-out and scale-up

4 Deploy separate environments in the same frame

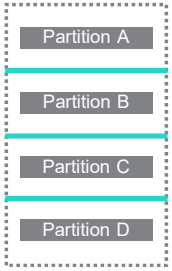

Use case examples:
- SAP HANA, SAP non HANA
- OLTP app layer, database layer
- LOB A (Marketing), LOB B (Finance)
- Customer 1, Customer 2 (Service provider)

Support for electrically isolated hard partitions*

Kinmax 擎昊科技

Open management framework simplifies hybrid IT environments

Open Source Management

DMTF Redfish

openstack

HPE Management Ecosystem

- HPE OneView
- HPE Insight Remote Support
- HPE SUM

HPE POINTNEXT

Choice and flexibility to manage your HPE Superdome Flex deployment

Kinmax 擎昊科技

Deployment choice with scalable or partitionable configurations

Scalable

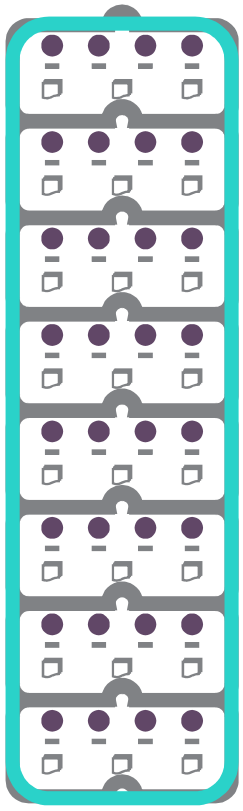

Partitionable

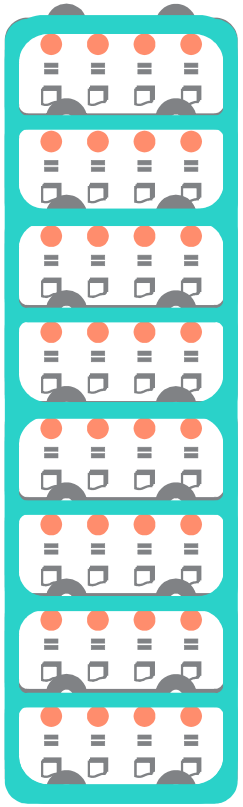

Scalable for single instance workloads

– Support a single, scalable system in 4 socket increments from 4 to 32 sockets with a **single OS instance**

Partitionable* to deploy separate environments in the same system

– Support multiple, independent hard partitions (**HPE nPars**) within a single system

– Hard partitions are configured in varying 4s (per chassis) increments from 4s to 32s

– Each hard partition runs its own OS instance, independently from other hard partitions

The unique value of HPE nPars

Hard partitions* add flexibility and cost efficiencies

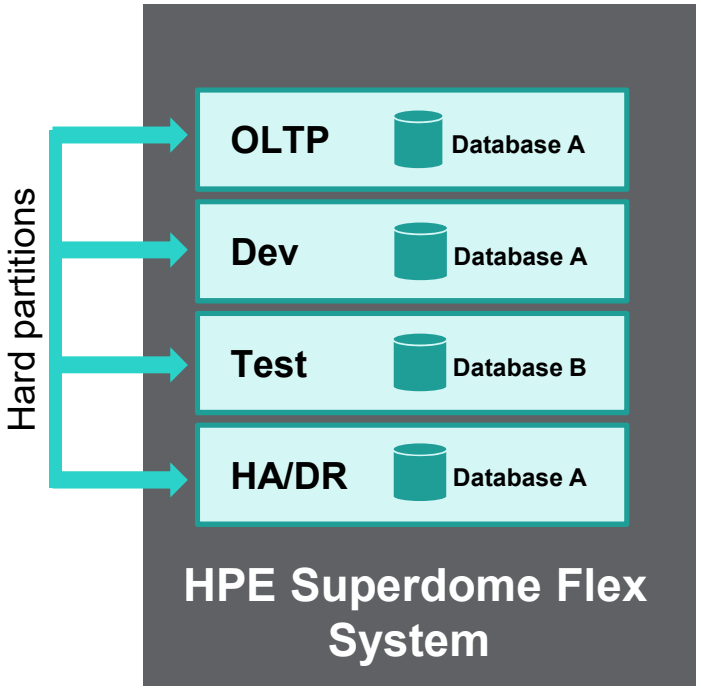

Lower your TCO	Optimize software costs by using HPE nPars
Maximize resource utilization	Create different development, test, and production environments within a single enclosure
Minimize downtime	Take one partition offline, perform maintenance and/or reconfiguration, while other partitions continue to run undisturbed 20x greater reliability
Protect your data	Electronic isolation provides a high degree of security between partitions

擁有專業及提供最高的可靠度
來保護重要的應用

Kinmax
擎昊科技

Safeguarding your critical workloads with HPE Superdome Flex

Proven Superdome Reliability framework delivers mission-critical availability

- Partnerships and expertise
- HPE Pointnext
- HPE Serviceguard
- HPE Workload Aware Security for Linux
- Operating System
- Error Analysis Engine
- Online optimization and repair
- Hard partitioning (HPE nPars)
- 'Firmware First' architecture
- Advanced memory resiliency
- Fault-tolerant fabric
- Redundant components

Up to 100% application availability

- Deep HPE mission-critical expertise, co-engineering with software partners and comprehensive HPE Pointnext services portfolio provide full solution availability

Error identification, reporting, recovery

- Best-in-class predictive fault handling initiates self-repair without operator assistance. Expanded protection with Serviceguard for Linux HA/DR clustering software

Five nines (99.999%) single-system availability

- HPE IP augments Intel base code to protect from and contain many errors, including memory errors, before interruption occurs at the OS layer.

Kinmax 擎昊科技

Proven and unique HPE Superdome RAS

Only HPE delivers advanced resiliency across all subsystems

Error reporting and handling: Error Analysis Engine

Best-in-class predictive fault handling initiates self-repair without operator assistance

Error containment: ”Firmware first” approach

Contains errors at the firmware level before any interruption can occur at the OS layer

Error detection and system self healing

Unique resiliency across every subsystem protects your critical workloads

Memory	Advanced memory resiliency (ADDDC)
I/O	Advanced PCIe error recovery (LER)
Fabric	Enhanced fabric resiliency (hot cable swap)
Processor	Advanced processor error handling (eMCA Gen2)

Linking directly into operating system software R&D

Unmatched partnership model maximizes mission critical capabilities

Driving innovations with partners and suppliers

- Processor advancements
- Hardware innovation
- Data center ecosystem

Decades of mission-critical experience

Hewlett Packard Enterprise

RESEARCH & DEVELOPMENT

Contributing what matters for mission-critical

- Upstream communities
- vmware
- SUSE

Kinmax 擎昊科技

Simplify security compliance

HPE Workload Aware Security for Linux (WASL)

Today, delivering a SAP HANA environment requires manual security compliance

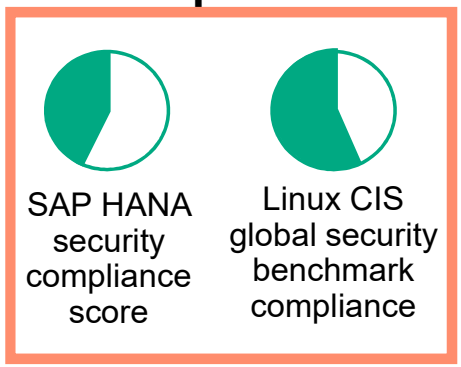

The 276 page SAP HANA security guide can take days or weeks
to follow and implement

HPE WASL

Push Button Security

- Assesses Security
- Hardens Linux O/S
- Hardens SAP HANA

Total solution addressing security compliance of both the Linux Operating System (SLES/RHEL) as well as the workload (SAP HANA)

HPE WASL can automatically bring compliance up to 90%

SAP HANA security compliance
~70 automated rules
Available in one click!

HPE WASL reduces security compliance deployment time from days to minutes

Kinmax 擎昊科技

Superdome Flex 規格

Kinmax
擎昊科技

Superdome Flex Specifications

	Description	
System	1 (4s) to 8 (32s) chassis per system; each chassis supports four (4) Intel® Xeon® Scalable processors	
Processors	Intel Xeon Platinum 8180 processor	28-cores/2.5GHz/205W/38.5M
	Intel Xeon Platinum 8180M processor	28-cores/2.5GHz/205W/38.5M (1.5TB)
	Intel Xeon Platinum 8176 processor	28-cores/2.1GHz/165W/38.5M
	Intel Xeon Platinum 8176M processor	28-cores/2.1GHz/165W/38.5M (1.5TB)
	Intel Xeon Platinum 8160 processor	24-cores/2.1GHz/150W/33M
	Intel Xeon Platinum 8160M processor	24-cores/2.1GHz/150W/33M (1.5TB)
	Intel Xeon Platinum 8156 processor	4-cores/3.6GHz/105W/16.5M
	Intel Xeon Platinum 8158 processor	12-cores/3.0GHz/150W/24.75M
	Intel Xeon Gold 6146 processor	12-cores/3.2GHz/165W/24.75M
	Intel Xeon Gold 6144 processor	8-cores/3.5GHz/150W/24.75M
	Intel Xeon Gold 6154 processor	18-cores/3.0GHz/200W/24.75M
	Intel Xeon Gold 6152 processor	22-cores/2.1GHz/140W/30.25M
	Intel Xeon Gold 6132 processor	14-cores/2.6GHz/140W/19.25M
	Intel Xeon Gold 6130 processor	16-cores/2.1GHz/125W/22M
Memory	48 DDR4 DIMM slots per chassis Maximum memory: 6 TB (48x 128 GB DIMMs) per chassis; 48TB per system 32 GB, 64 GB, and 128GB DDR4 DIMMs loaded in groups of 12 DIMMs	

Superdome Flex Specifications

	Description
Base IO (base chassis)	2x 10GbE ports, 2x 1GbE ports, 4x USB 3.0 ports, serial and MGMT ports
Internal drive slots	Up to four (4) 2.5” SATA/SAS HDD or SSD with option for hardware RAID
IO expansion options	16 PCIe 3.0 low-profile slots; 7 x16 slots and 9 x8 slots 12 PCIe 3.0 slots; 8 full-height slots (4 x16 & 4 x8) + 4 low-profile slots (1 x16 and 3 x8) Zero (0) slot, compute only
Management	Optional 1U Rack Management Controller (RMC) for CLI; Redfish® API
Operating systems	Red Hat Enterprise Linux (RHEL) 7 SUSE Linux Enterprise Server (SLES) 12 Oracle Linux 7 VMware Microsoft Windows Server 2016 Standard and Datacenter
Form Factor	5U server chassis; width: 17.5” (44.5cm); depth: 32.5” (82.6cm)

Superdome Flex supported storage

- **Internal storage:** Four (4) 2.5” drive bays to support SATA SSDs or SAS HDDs/SSDs.
  - 6G SATA SSDs use embedded chip (Intel RSTe) with SW RAID (w/boot support)
  - 12G SAS HDDs/SSDs use PCIe RAID card (internal) with HW RAID (w/boot support)
- **SAS**: HPE SAS JBOD (e.g. D3700) are supported with PCIe RAID card (external) (w/boot support)
- **Fibre Channel**: HPE FC arrays (e.g. 3PAR, XP, MSA, Nimble) support w/PCIe FC HBAs (w/boot support)
- **Third party storage**: Storage vendor takes the lead in documenting interoperability
- Reference SPOCK for additional details

Superdome Flex Base Chassis (rear)
with 16-slot PCIe riser

HPE Superdome Flex 的效益

Kinmax
擎昊科技

HPE Superdome Flex 摘要

- 為專為核心系統或高穩定度需求環境設計的產品x86伺服器。
- 可提供較一般x86伺服器更高的穩定度，可大幅減少系統當機產生的損失。
- 可提供垂直擴充能力，未來效能不足可原機擴充最大到8倍的運算能力，有效節省投資。
- 產品生命週期為一般x86伺服器2倍，可有效節省伺服器因擴充需求被迫汰換的風險。

Kinmax
擎昊科技

HPE Superdome Flex擴充能力比較

Kinmax
擎昊科技

擴充能力比較

HPE Superdome Flex vs. 4-Socket x86 Server

Superdome Flex 可由四個CPU及768GB記憶體起始，未來可以模組化方式擴充到32個處理器及48TB記憶體

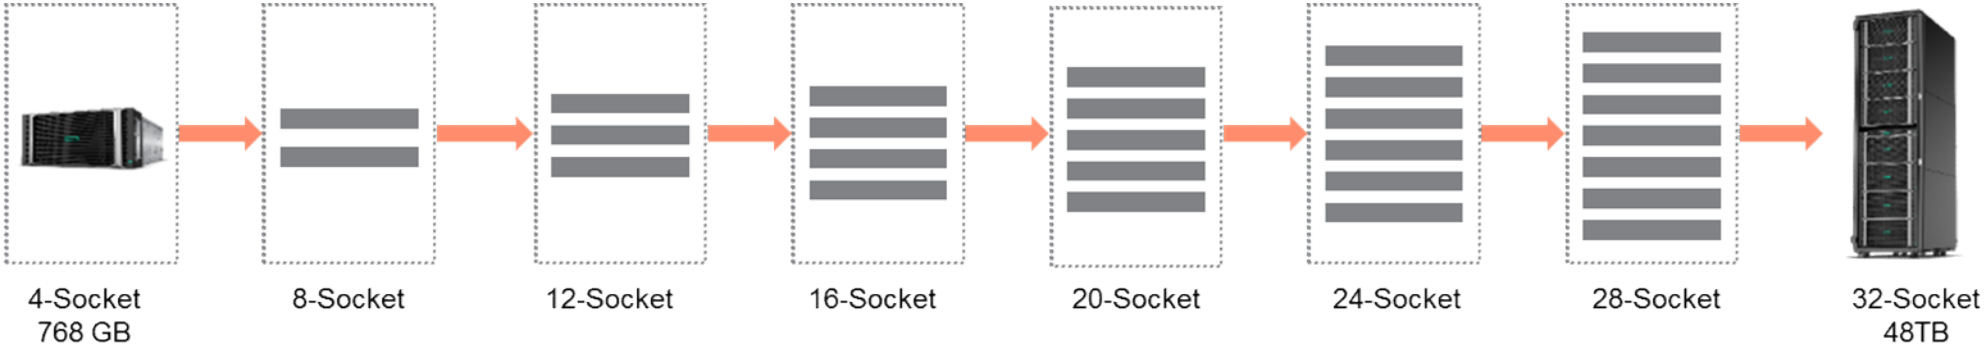

Kinmax
擎昊科技

HPE Superdome Flex
與
標準x86主機
RAS比較
(Reliability, Availability, Serviceability)

Kinmax
擎昊科技

End-to-end RAS protects high value applications and data

HPE Superdome Flex RAS features at a glance

Chassis

- Firmwa
- Redun
- Flex G negotia
- Adapti
- CRC p
- System
- Socket
- Chassi
- Hard P

Process

- Enhan
- Integer
- ECC c
- Regist
- Improv
- UPI lin
- UPI rol
- Core le
- Poison Data Containment

HPE Superdome Flex: key areas of RAS superiority

RAS	SD Flex	Standard x86
Firmware-first	✓	x
Automatic error logging	✓	x
Auto self-healing (Analysis Engine)	✓	x
Disabling / deconfiguration of failed FRUs	✓	x
Onboard fault analyzer	✓	x
Automatic restart	✓	x
Advanced processor error handling (eMCA Gen2)	✓	x
Advanced memory resiliency (ADDDC)	✓	x
Enhanced fabric resiliency (Adaptive routing)	✓	x
Advanced PCIe error recovery (LER)	✓	x
Hard Partitions (nPars)	✓	x

C needs

e

ainment and card

containment

- PCIe link retraining and recovery

Kinmax 擎昊科技

主機比較

RAS on Chassis-Level features

RAS on Chassis-Level features	Superdome Flex RAS features	x86 RAS features
Firmware-first error handling	V	X
Redundant, hot-swappable power supplies (N+N or N+1) and fans	V	N+1
Flex Grid link failover, link-level retry, dynamic link tuning and bandwidth negotiation	V	X
Adaptive routing finds bad fabric links and routes traffic around failures	V	X
CRC protection per micro-packet and fast retry for transient errors	V	X
Systemic transient errors, triggering retry or recovery attempt	V	X
Socket indictment at boot	V*	X
Chassis de-configure at boot	V*	X
Hard Partitions (nPar)	V	X

* Coming soon

主機比較

RAS on Memory features

RAS on Memory features	Superdome Flex RAS features	x86 RAS features
Proactive memory (patrol and demand) scrubbing	V	V
Adaptive DDDC, SDDC+1 for Mission Critical and HPC needs Superdome Flex thresholds)	V	X
Address/Cmd parity error resiliency	V	X
Doubled refresh rates at high temp	V	X
DIMM de-configuration at boot	V	X
Memory error logging/history in management firmware	V	X
OS-level page deallocation	V	X
Memory error storm response [coming soon]	V	X
DRAM Post Package repair [coming soon]	V	X
Rank and Bank sparing	V	X

主機比較

RAS on I/O capabilities

RAS on I/O capabilities	Superdome Flex RAS features	x86 RAS features
PCIe Live Error Recovery (LER); PCIe root port containment and card error recovery	V	X
PCIe "Stop and Scream"; PCIe root port corrupt data containment	V	X
PCIe end-to-end CRC checking [coming soon]	V	X
PCIe corrupt data containment (data poisoning)	V	X
PCIe link CRC error check and retry	V	X
PCIe link retraining and recovery	V	X

主機比較

RAS on Processor coverage

RAS on Processor coverage	Superdome Flex RAS features	x86 RAS features
Enhanced Machine Check Gen2 architecture and recovery	V	X
Integer pipeline/instruction pipeline retry capability	V	X
ECC coverage on all internal caches and cache tags	V	X
Register/TLB parity protection	V	X
Improved error (Viral) containment, aiding system survivability	V	X
UPI link-level retry/restart/recalibrate	V	X
UPI rolling CRC check for transient errors	V	X
Core level corrupt data containment and core disable at boot time	V	X
Poison Data Containment	V	X

HPE Superdome Flex
長期產品生命週期

Kinmax
擎昊科技

Intel Xeon CPU Roadmap

Kinmax
擎昊科技

Mission-critical x86 solutions roadmap

Superdome Flex可提供5年以上的產品生命週期保證

	Now	Future
CS500 for SAP HANA 	**Entry point for SAP HANA deployments** 4-socket, 4TB scale-up 68TB scale-out	**Higher** performance and **Greater** density Up to 4-socket scale-up
MC990 X **MC990 X TDI for SAP HANA** 	**Modularity for extreme scale** 4s modular scalability to 32-socket, 48TB HANA: 20-socket/20TB scale-up	Even more scalable!
Superdome X **CS900 for SAP HANA** 	**Highest availability with large memory capacity** 16-socket, 48TB HANA: 16-socket/16TB scale-up, 192TB scale-out	Now 2x memory!
Superdome Flex **Superdome Flex TDI for SAP HANA** 	**Next Generation Mission-critical x86**	**HPE Innovation platform for the next decade including 'Machine' technology for Memory-Driven Computing** **Scale-up to 32-socket** **4-socket & 8-socket nodes for scale-out**

Continued Investment
- Offering customer choice during technology transition
- Ensuring solution ecosystem maturity for critical workloads

NOTE: This is a rolling (up to three year) roadmap and is subject to change without notice

HPE Superdome Flex
與
4-Socket x86 Server
效益分析

Kinmax
擎昊科技

整體效益分析(一)

功能特色	HPE Superdome Flex	4-Socket x86 Server	效益說明
效能擴充能力	每箱Intel Xeon 處理器4顆, 每次一箱加4顆,最多可加到8箱32顆處理器	最多4顆處理器	DL580已無法再擴充CPU。 SD Flex垂直擴充功能，可在未來性能不足時繼續擴充，可節省一倍以上投資。
處理器規格	Intel Xeon-G 6144 (3.5GHz/8-core/150W)	Intel Xeon-Gold 6134 (3.3GHz/8-core/130W)	SD Flex提供較高效能之CPU型號。
記憶體容錯	Adaptive DDDC	SDDC+1	SD Flex 提供更高級的記憶體容錯機制，可大幅降低系統當機造成系統維運終止所產生的損失。
系統穩定度	99.999 (downtime 5m 15.6s per year)	99.95 (downtime 4h 22m 58.5s per year)	SD Flex整體RAS容錯機制較4-Socket x86 Server多，可大幅提升系統穩定度

Kinmax 擎昊科技

整體效益分析(二)

功能特色	HPE Superdome Flex	4-Socket x86 Server	效益說明
產品生命週期	提供5年以上產品可有零件擴充保證	一般CPU產品週期，一年半至兩年即停產	SD Flex長期產品週期保證，可確保系統未來效能不足可有零件擴充，較一般x86主機亦可節省一倍以上之主機重購或系統升級的風險
整體評估	Superdome Flex 為提供核心系統專用機主，初期建置成本高於4-Socket x86 Server，但後續穩定、擴充性及長期生命週期，可大幅降低整體維運成本2倍以上。	4-Socket x86 Server初期投資成本較低，但對於核心系統的維運會造成較大的成本。	

Kinmax
擎昊科技

Q&A

Kinmax
擎昊科技

Backup

Kinmax
擎昊科技

HPE SuperDome Flex – RAS features

Kinmax
擎昊科技

Safeguarding critical workloads with HPE Superdome Flex

Proven Superdome Reliability framework delivers mission-critical availability

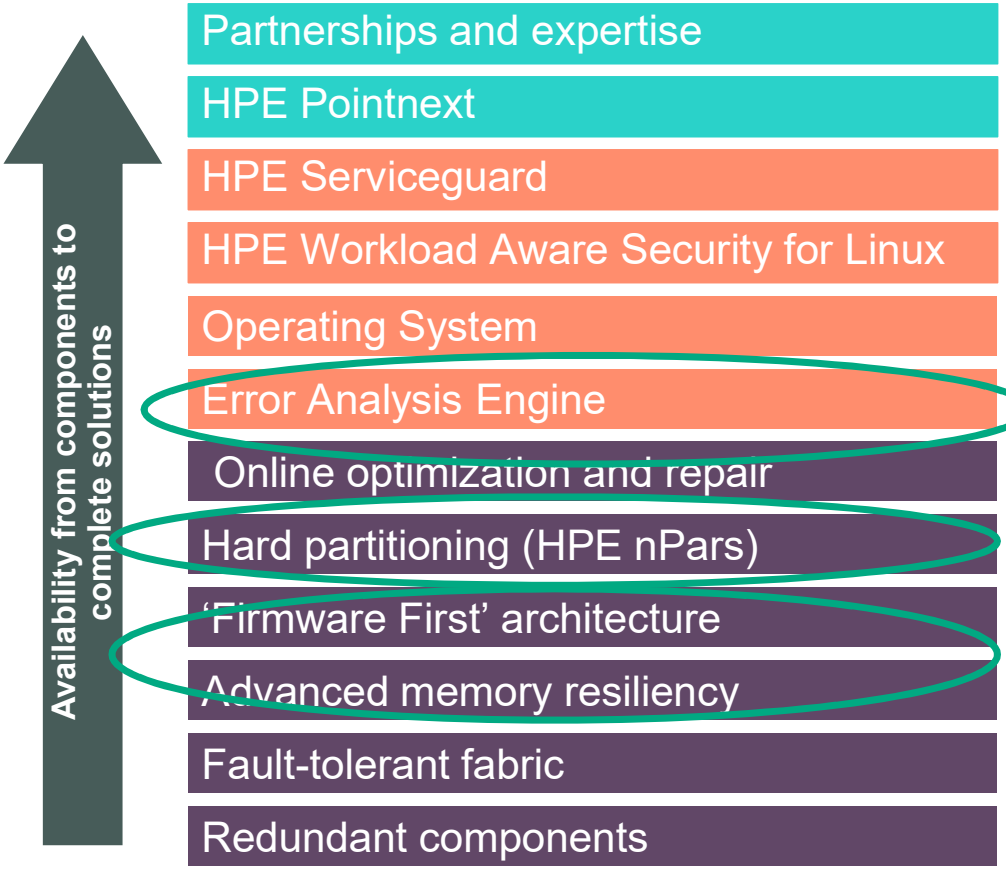

Up to 100% application availability

- Deep HPE mission-critical expertise, co-engineering with software partners and comprehensive HPE Pointnext services portfolio provide full solution availability

Error identification, reporting, recovery

- Best-in-class predictive fault handling initiates self-repair without operator assistance. Expanded protection with Serviceguard for Linux HA/DR clustering software

Five nines (99.999%) single-system availability

- HPE IP augments Intel base code to protect from and contain many errors, including memory errors, before interruption occurs at the OS layer.

Kinmax 擎昊科技

RAS features discussion and showcase scenarios

1. nPar capability:
   - Highlight SD Flex capability of consolidation and flexibility through nPar → live demo
   - Note: multiple nPar feature available at release 1.2
2. Memory RAS capability:
   - Highlight SD Flex differentiations in keeping in-memory data safe and stable
   - SW protection through MCA recovery → live demo
   - Memory de-configuration → live demo
3. CPU de-configuration (WIP):
   - Socket de-configuration availability at release 1.1
   - Core de-configuration availability at release 1.2

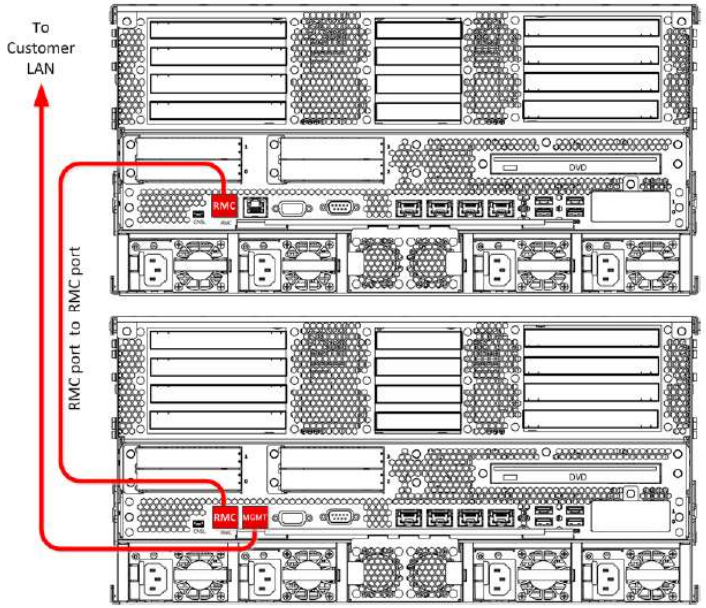

Kinmax
擎昊科技

SD Flex Hard partitioning: electrical isolation and flexibility

8S SD Flex with two chassis:
- Two npars, per 4S / 3TB RAM configuration

Modifying nPar configuration:
- Delete nPar1
- Extend nPar0 to be with 8S / 6TB RAM

Rebooting new nPar configuration:
- OS and application getting double resources
- Observing reboot time

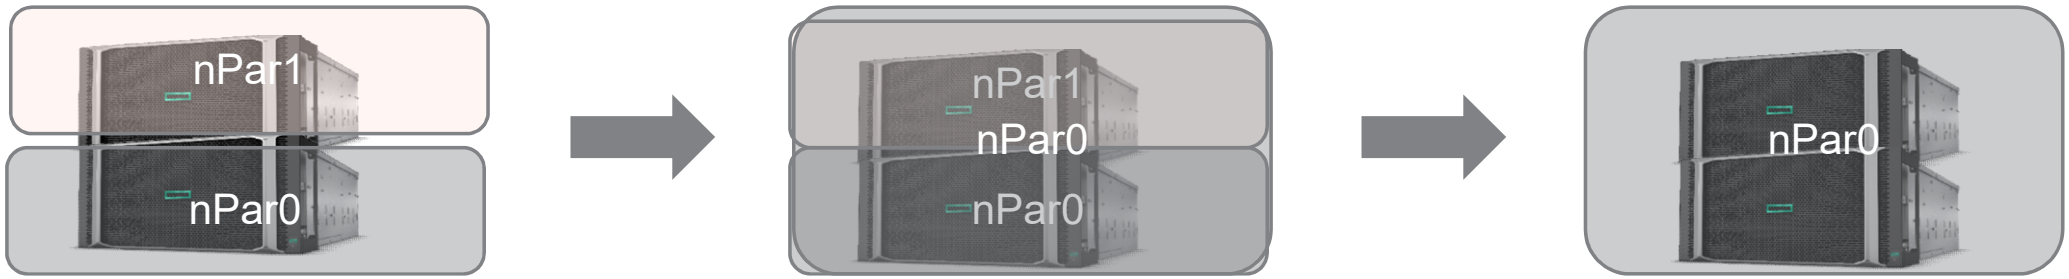

Kinmax 擎昊科技

Initial 2 nPar configuration→ Removing one nPar

eRMC-cli> show npar

```
SDFLexRMC eRMC:r001i06c cli> show npar

Partitions: 2

Par   Run        Status     # of  HT  RAS CPUs      Memory (GB)    IO Cards  Boot     Boot
Num   State                 BMCs          In/OK     In/OK          In/OK     Chassis  Slots
===== ========== ========== ===== === === ========= ============== ========= ======== ============
p0    Off        OK         1     on  on  4/4       3072/3072      4/4       r001i06b 3,5
p1    Off        OK         1     on  on  4/4       3072/3072      4/4       r001i11b 3,5

 * In/OK = Installed/OK
```

Initial 2nPars

eRMC-cli> remove npar pnum=1

```
SDFLexRMC eRMC:r001i06c cli> remove npar pnum=1

SUCCESS

Partitions: 1

Par   Run        Status     # of  HT  RAS CPUs      Memory (GB)    IO Cards  Boot     Boot
Num   State                 BMCs          In/OK     In/OK          In/OK     Chassis  Slots
===== ========== ========== ===== === === ========= ============== ========= ======== ============
p0    Off        OK         1     on  on  4/4       3072/3072      4/4       r001i06b 3,5

 * In/OK = Installed/OK


BMCs: 2

BMC        Rack UPos Par   Power      Status     CPUs      Memory (GB)    IO Cards
           Num       Num   State                 In/OK     In/OK          In/OK
========== ==== ==== ===== ========== ========== ========= ============== =========
r001i06b   001  06   p0    off        OK         4/4       3072/3072      4/4
r001i11b   001  11   -     off        OK         4/4       3072/3072      4/4

 * In/OK = Installed/OK
```

One nPar remains

One chassis not assigned

Kinmax 擎昊科技

Adding a chassis into nPar0 to make 8S/6TB partition

eRMC-cli> modify npar pnum=0 add-chassis=r001i11b

```
SDFLexRMC eRMC:r001i06c cli> modify npar pnum=0 add_chassis=r001i11b

SUCCESS


Partitions: 1

Par   Run         Status     # of  HT   RAS CPUs      Memory (GB)     IO Cards  Boot      Boot
Num   State                  BMCs          In/OK     In/OK           In/OK     Chassis   Slots
===== ========== ========== ===== === === ========= =============== ========= ========= ============
p0    Off         OK         2     on   on  8/8       6144/6144       8/8       r001i06b 3,5

 * In/OK = Installed/OK


BMCs: 2

BMC        Rack UPos Par   Power      Status     CPUs      Memory (GB)     IO Cards
           Num       Num   State                 In/OK     In/OK           In/OK
========== ==== ==== ===== ========== ========== ========= =============== =========
r001i06b   001  06   p0    off        OK         4/4       3072/3072       4/4
r001i11b   001  11   p0    off        OK         4/4       3072/3072       4/4

 * In/OK = Installed/OK
```

nPar0 size is doubled with 2 chassis

Kinmax 擎昊科技

Booting time (8S/6TB)

First reboot:
- To EFI: 8min (including nPar reconfiguration)
- From EFI to OS: 4min 30sec

Subsequent boot:
- From reboot to EFI: 2min 30sec
- From EFI to OS: 4min 30sec

[rhel-OS] # lscpu

```
[root@sdflex-rhel73 ~]# lscpu
Architecture:          x86_64
CPU op-mode(s):        32-bit, 64-bit
Byte Order:            Little Endian
CPU(s):                448
On-line CPU(s) list:   0-447
Thread(s) per core:    2
Core(s) per socket:    28
Socket(s):             8
NUMA node(s):          8
Vendor ID:             GenuineIntel
CPU family:            6
Model:                 85
Model name:            Intel(R) Xeon(R) Platinum 8180 CPU @ 2.50GHz
Stepping:              4
CPU MHz:               3800.000
BogoMIPS:              5014.07
Virtualization:        VT-x
L1d cache:             32K
L1i cache:             32K
L2 cache:              1024K
L3 cache:              39424K
NUMA node0 CPU(s):     0-27,224-251
NUMA node1 CPU(s):     28-55,252-279
NUMA node2 CPU(s):     56-83,280-307
NUMA node3 CPU(s):     84-111,308-335
NUMA node4 CPU(s):     112-139,336-363
NUMA node5 CPU(s):     140-167,364-391
NUMA node6 CPU(s):     168-195,392-419
NUMA node7 CPU(s):     196-223,420-447
[root@sdflex-rhel73 ~]#
```

OS view of nPar0:
8S / 448CPU
(HT on)

eRMC-cli> show npar

```
SDFLexRMC eRMC:r001i06c cli> show npar

Partitions: 1

Par  Run        Status     # of  HT  RAS CPUs      Memory (GB)     IO Cards     Boot      Bo
Num  State                  BMCs          In/OK     In/OK           In/OK        Chassis   Slots
===  =========  =========  ====  ==  === ========  ==============  ========     ========  ==========
p0   EFI Boot   OK         2     on  on  8/8       6144/6144       8/8          r001i06b  3,5

 * In/OK = Installed/OK
```

[rhel-OS] # free -m

```
[root@sdflex-rhel73 ~]# free -m
              total        used        free      shared  buff/cache   available
Mem:        6002498       39608     5962127          21         762     5960581
Swap:         20479           0       20479
[root@sdflex-rhel73 ~]#
```

OS view of nPar0:
6TB RAM

Returning back to 2 nPars with 4S/3TB per nPar

eRMC-cli> modify npar pnum=0 remove-chassis=r001i11b

```
SDFLexRMC eRMC:r001i06c cli> modify npar pnum=0 remove_chassis=r001i11b

SUCCESS

Partitions: 1

Par   Run        Status     # of  HT  RAS CPUs      Memory (GB)     IO Cards  Boot     Boot
Num   State                 BMCs          In/OK     In/OK           In/OK     Chassis  Slots
===== ========== ========== ===== === === ========= ============== ========= ======== ============
p0    Off        OK         1     on  on  4/4       3072/3072       4/4       r001i06b 3,5
 * In/OK = Installed/OK

BMCs: 2

BMC        Rack UPos Par   Power      Status     CPUs      Memory (GB)     IO Cards
           Num       Num   State                 In/OK     In/OK           In/OK
========== ==== ==== ===== ========== ========== ========= ============== =========
r001i06b   001  06   p0    off        OK         4/4       3072/3072       4/4
r001i11b   001  11   -     off        OK         4/4       3072/3072       4/4
 * In/OK = Installed/OK
```

nPar0 reduced to one chassis; 4S/3TB

Creating additional nPar1

eRMC-cli> create npar pnum=1 ras=on chassis=r001i11b

```
SDFLexRMC eRMC:r001i06c cli> create npar pnum=1 ras=on chassis=r001i11b

SUCCESS

Partitions: 2

Par   Run        Status     # of  HT  RAS CPUs      Memory (GB)     IO Cards  Boot     Boot
Num   State                 BMCs          In/OK     In/OK           In/OK     Chassis  Slots
===== ========== ========== ===== === === ========= ============== ========= ======== ============
p0    Off        OK         1     on  on  4/4       3072/3072       4/4       r001i06b 3,5
p1    Off        OK         1     on  on  4/4       3072/3072       4/4       r001i11b 3,5
 * In/OK = Installed/OK

BMCs: 2

BMC        Rack UPos Par   Power      Status     CPUs      Memory (GB)     IO Cards
           Num       Num   State                 In/OK     In/OK           In/OK
========== ==== ==== ===== ========== ========== ========= ============== =========
r001i06b   001  06   p0    off        OK         4/4       3072/3072       4/4
r001i11b   001  11   p1    off        OK         4/4       3072/3072       4/4
 * In/OK = Installed/OK
```

Booting time (4S/3TB)

First reboot:
- To EFI: 5min (including nPar reconfiguration)
- From EFI to OS: 2min 20sec

Subsequent boot:
- From reboot to EFI: 2min 30sec
- From EFI to OS: 2min 20sec

eRMC-cli> show npar

```
SDFLexRMC eRMC:r001i06c cli> show npar

Partitions: 2

Par  Run       Status    # of  HT  RAS CPUs     Memory (GB)    IO Cards  Boot     Boot
Num  State               BMCs          In/OK    In/OK          In/OK     Chassis  Slots
===  ========= ========= ===== === === ======== ============== ========= ======== ===========
p0   OS Boot   OK        1     on  on  4/4      3072/3072      4/4       r001i06b 3,5
p1   Off       OK        1     on  on  4/4      3072/3072      4/4       r001i11b 3,5

* In/OK = Installed/OK
```

[rhel-OS] # lscpu

```
[root@sdflex-rhel73 ~]# lscpu
Architecture:          x86_64
CPU op-mode(s):        32-bit, 64-bit
Byte Order:            Little Endian
CPU(s):                224
On-line CPU(s) list:   0-223
Thread(s) per core:    2
Core(s) per socket:    28
Socket(s):             4
NUMA node(s):          4
Vendor ID:             GenuineIntel
CPU family:            6
Model:                 85
Model name:            Intel(R) Xeon(R) Platinum 8180 CPU @ 2.50GHz
Stepping:              4
CPU MHz:               3800.000
BogoMIPS:              5007.52
Virtualization:        VT-x
L1d cache:             32K
L1i cache:             32K
L2 cache:              1024K
L3 cache:              39424K
NUMA node0 CPU(s):     0-27,112-139
NUMA node1 CPU(s):     28-55,140-167
NUMA node2 CPU(s):     56-83,168-195
NUMA node3 CPU(s):     84-111,196-223
```

OS view of nPar0:
4S / 224CPU
(HT on)

[rhel-OS] # free -m

```
[root@sdflex-rhel73 ~]# free -m
              total        used        free      shared  buff/cache   available
Mem:        3000595       20312     2979793          11         488     2978180
Swap:         20479           0       20479
[root@sdflex-rhel73 ~]#
```

OS view of nPar0:
3TB RAM

Superdome Flex memory RAS feature - highlight

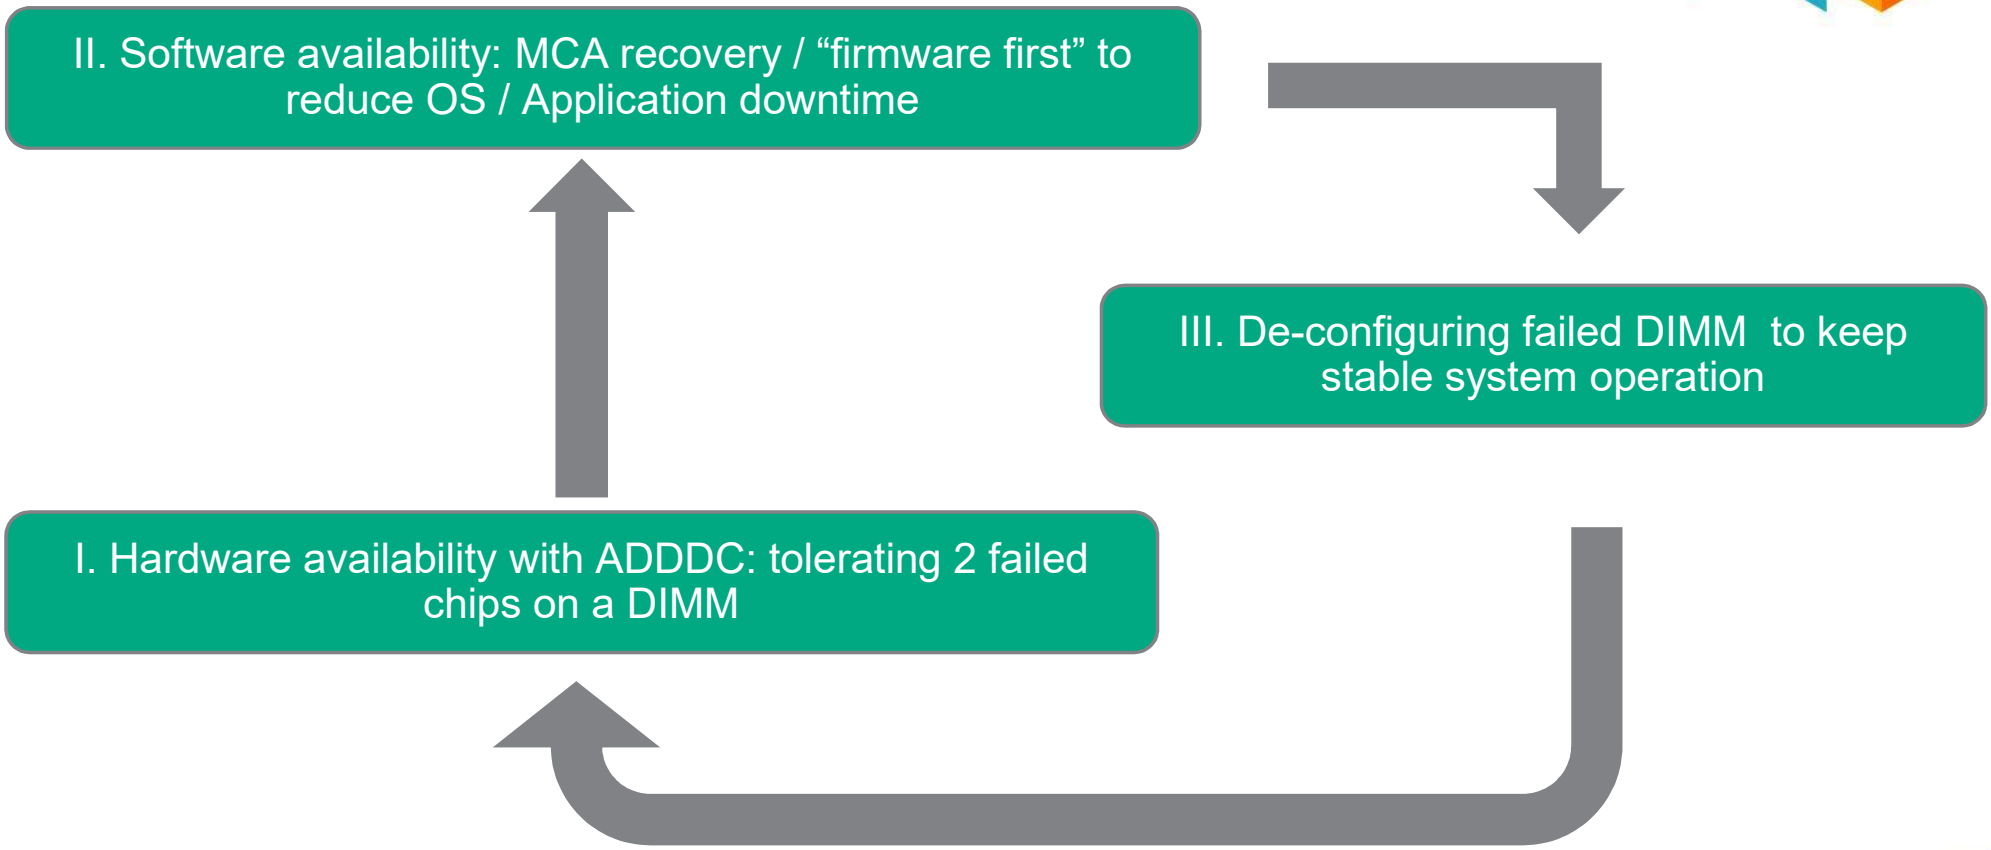

Kinmax 擎昊科技

I. Advanced Double Device Data Correction – explained (ADDDC)

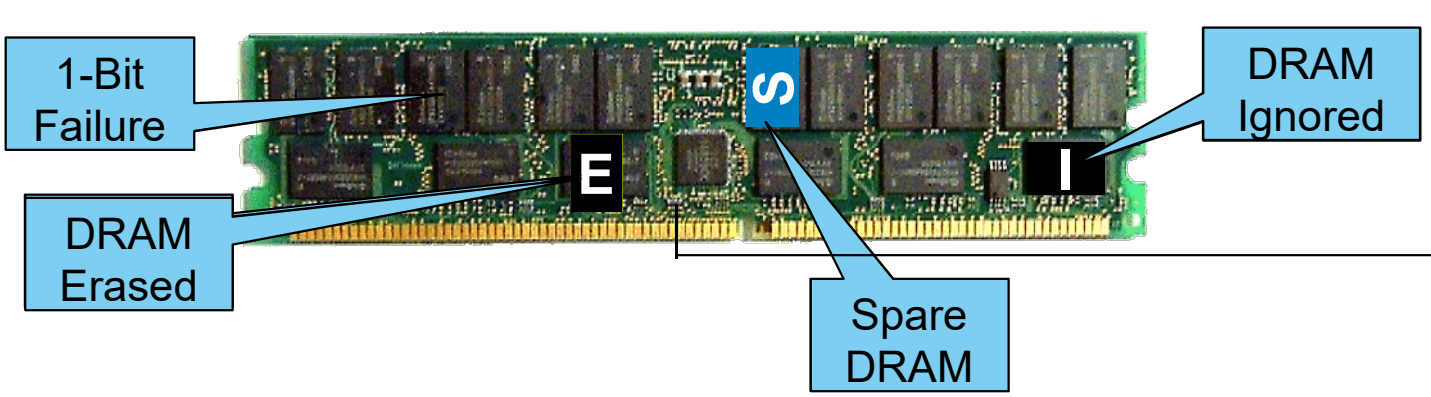

1. Intel SP CPU has 2 independent memory controllers (MC); each MC has 3 fully independent memory channels, with Error Correcting Code (ECC), in performance mode
2. If a DRAM chip fails, then the Memory Controller (MC) can recreate lost data with ECC logic
3. When first chip failure rate hits a threshold, The MC "spares" the first bad DRAM by using the ECC logic to recreate the lost data and store it in a spare DRAM
4. After a DRAM is spared, a second DRAM can fail and the ECC logic can once again recreate lost data
5. The MC can "erase" the second bad DRAM by marking it as known bad to improve the speed of the ECC logic
6. After a DRAM is tagged for erasure, the memory controller can still correct single bit errors

Historically, double-chip sparing provided 1/17th the repair rate, and 1/3rd less memory downtime compared to single-chip sparing, and Intel's implementation has improved since then.

ADDDC advancement compared to DDDC

ADDDC memory subsystem is always configured to operate in performance mode, until there is threshold of correctable errors is reached. That way, the system achieves availability objective, while reducing performance impact.

Adaptive Double DRAM Device Correction (ADDDC)

Intel® Xeon® processor introduces an innovative approach in managing errors that the DDR4 DRAM DIMM may induce through the life of the product. ADDDC is deployed at runtime to dynamically map out the failing DRAM device and continue to provide SDDC ECC coverage on the DIMM, translating to longer DIMM longevity. The operation occurs at the fine granularity of DRAM Bank and/or Rank to have minimal impact on the overall system performance.

With the advent of ADDDC, the memory subsystem is always configured to operate in performance mode. When the number of corrections on a DRAM device reaches the targeted threshold value, with help from the UEFI runtime code, the identified failing DRAM region is adaptively placed in lockstep mode where the identified failing region of the DRAM device is mapped out of ECC. Once in ADDDC, cache line ECC continues to cover single DRAM (x4) error detection and apply a correction algorithm to the nibble.

https://software.intel.com/en-us/articles/new-reliability-availability-and-serviceability-ras-features-in-the-intel-xeon-processor

Kinmax 擎昊科技

II. HPE Superdome Flex MCA recovery: advanced memory error recovery

Basic flow

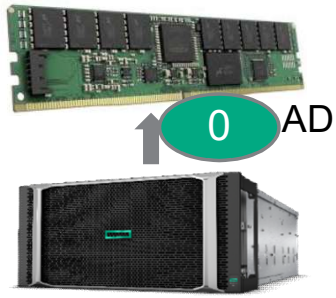

1 Uncorrectable memory error consumed by user application

2 Error is intercepted and contained at firmware level

3 IOs are stopped preventing corrupted data on disk

4 Bad memory DIMMs are flagged and indicted / de-configured from allocated memory pool

5 Kernel sends a “KILL” signal to application. Application with signal handler keep running

0 ADDDC: explained in <I>

Kinmax 擎昊科技

MCA recovery demo: advanced memory error recovery

```
[root@sdflex-rhel73 demo]# ./mca-prepare.sh
[  761.923485] EINJ: Error INJection is initialized.
mount: none is already mounted or /sys/kernel/debug busy
[root@sdflex-rhel73 demo]#
```

```
[root@sdflex-rhel73 demo]# ./mca-recover
vtop(7ffff7ff8000) = 1776290d000

[1]+  Stopped                 ./mca-recover
[root@sdflex-rhel73 demo]# ./mca-inject.sh 1776290d000
[root@sdflex-rhel73 demo]# fg
./mca-recover
[ 2231.742132] mce: [Hardware Error]: Machine check events logged
[ 2231.748651] mce: Uncorrected hardware memory error in user-access at 1776290d000
[ 2231.756917] MCE 0x1776290d: Killing mca-recover:8892 due to hardware memory corruption
[ 2231.765967] MCE 0x1776290d: dirty LRU page recovery: Recovered
recover: sig=7 si=0x7fffffffddb0 v=0x7fffffffdc80
Platform memory error at 0x0x7ffff7ff8000
addr = 0x7ffff7ff8000 lsb=12
Recovery allocated new page at physical 1776d004000
```

Starting sample appl

Injecting error into RAM addr

MCA recovery kicks in to avoid crash

```
Hardware event. This is not a software error.
MCE 0
CPU 141 BANK 1 TSC 57e368681ac
RIP 33:400a69
MISC 86 ADDR 1776290d000
TIME 1522205013 Wed Mar 28 10:43:33 2018
MCG status:RIPV EIPV MCIP
MCi status:
Uncorrected error
Error enabled
MCi_MISC register valid
MCi_ADDR register valid
SRAR
MCA: Data CACHE Level-0 Data-Read Error
STATUS bd80000000100134 MCGSTATUS 7
MCGCAP f000c14 APICID 43 SOCKETID 1
CPUID Vendor Intel Family 6 Model 85
```

Error is logged at OS – appl interrupt triggered

```
SDFLexRMC eRMC:r001i06c cli> show indict


r001i06b

  FF00070106000074 rack1/chassis_u6/cpu1/dimmG0

    Indicted: Yes
        2000-01-23T21:18:01Z  Reason ID: 306
            Summary: Memory uncorrectable data error
            Cause: DIMMs are incorrectly seated, or multiple DRAM chips on DIMMs are failing, or else the DDR channel for
this DIMM failed.
            Action: Verify the DIMMs are properly seated into the DIMM slots.  Inspect the board for any defects or contam
ination in the DIMM connector. If the problem persists, replace the affected DIMMs. In case of a DDR channel failure, repl
ace the chassis.
```

Memory DIMM is indicted at FW layer for further processing

III. HPE Superdome Flex memory de-configuration feature

3 System keeps running and stability with healthy components. Alert is sent for scheduled replacement

2 DIMM is taken out to keep system running, marked as “de-configured” in next boot

1 A memory DIMM fails or malfunctions

Kinmax 擎昊科技

De-configuring failed memory DIMM from the nPar

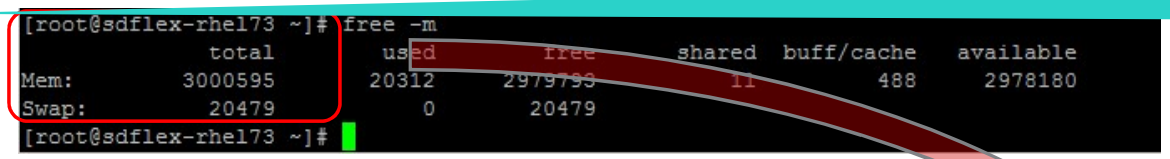

RAM size before deconfig

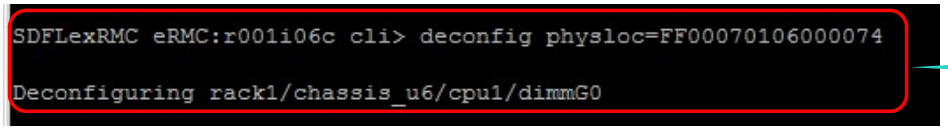

Deconfiguring DIMM with issue

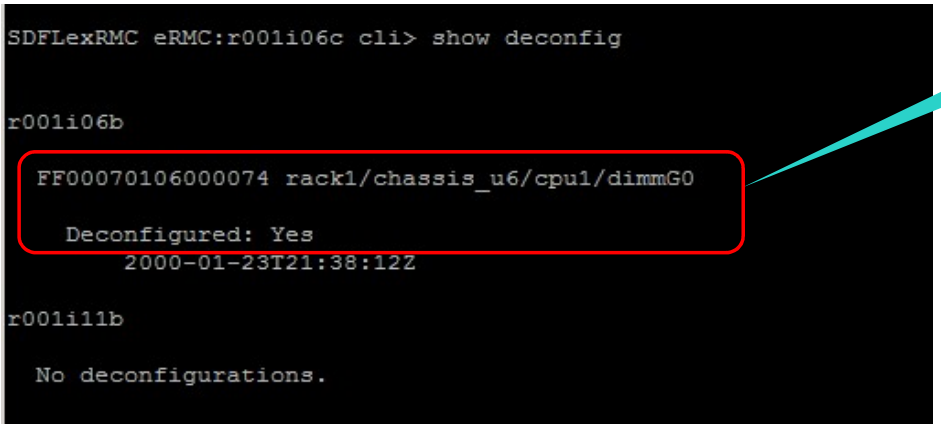

RAM size reduced by 64GB

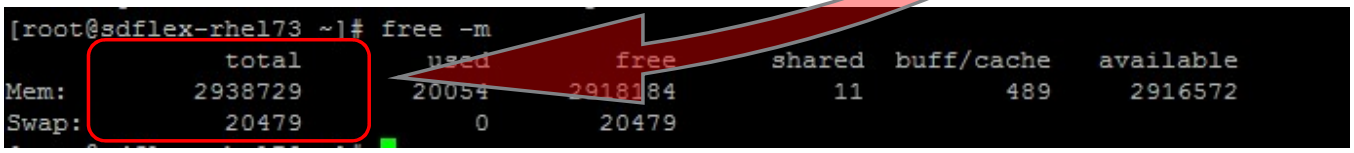

Kinmax 擎昊科技