



**Hewlett Packard  
Enterprise**

# **HPE Superdome Flex**

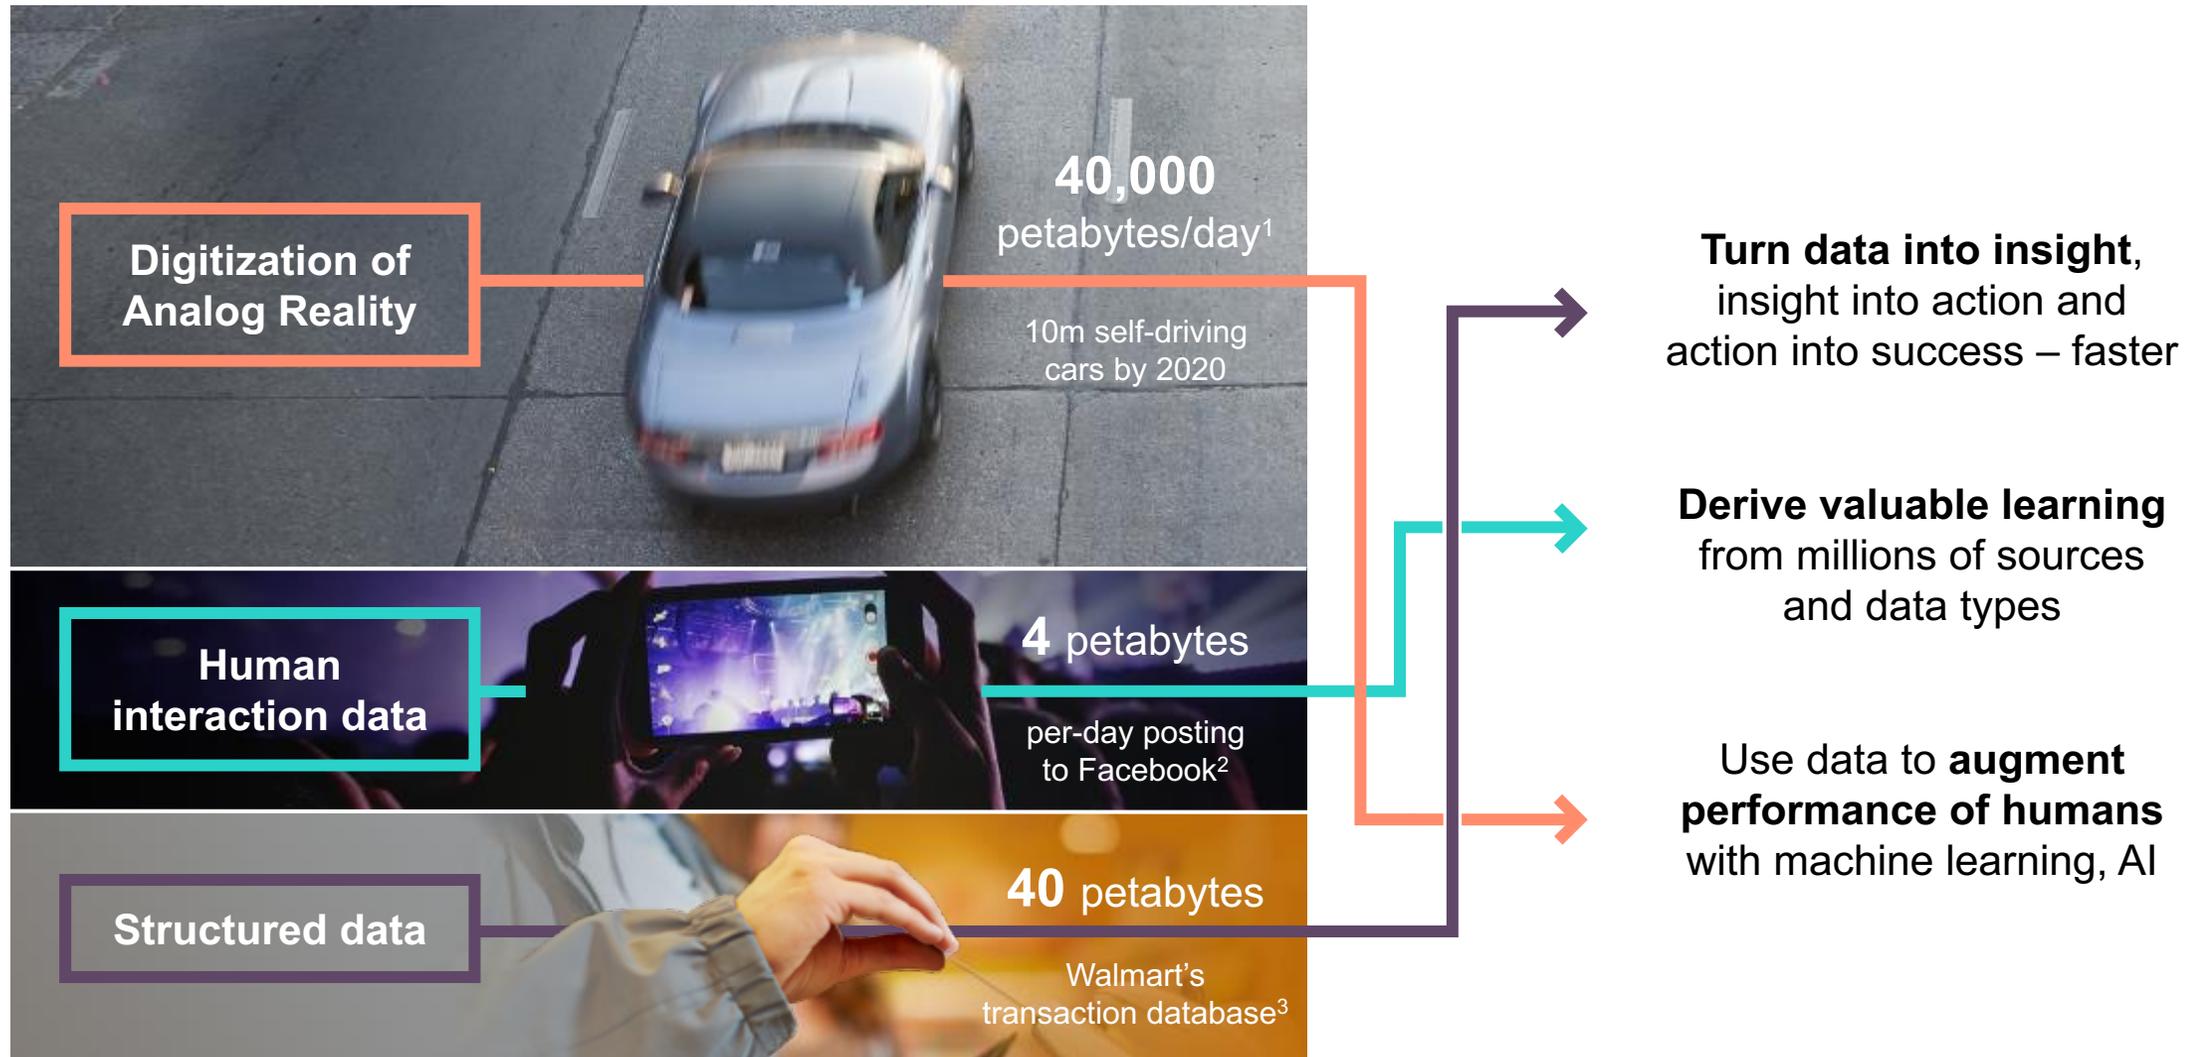
---

# Forward-looking statements

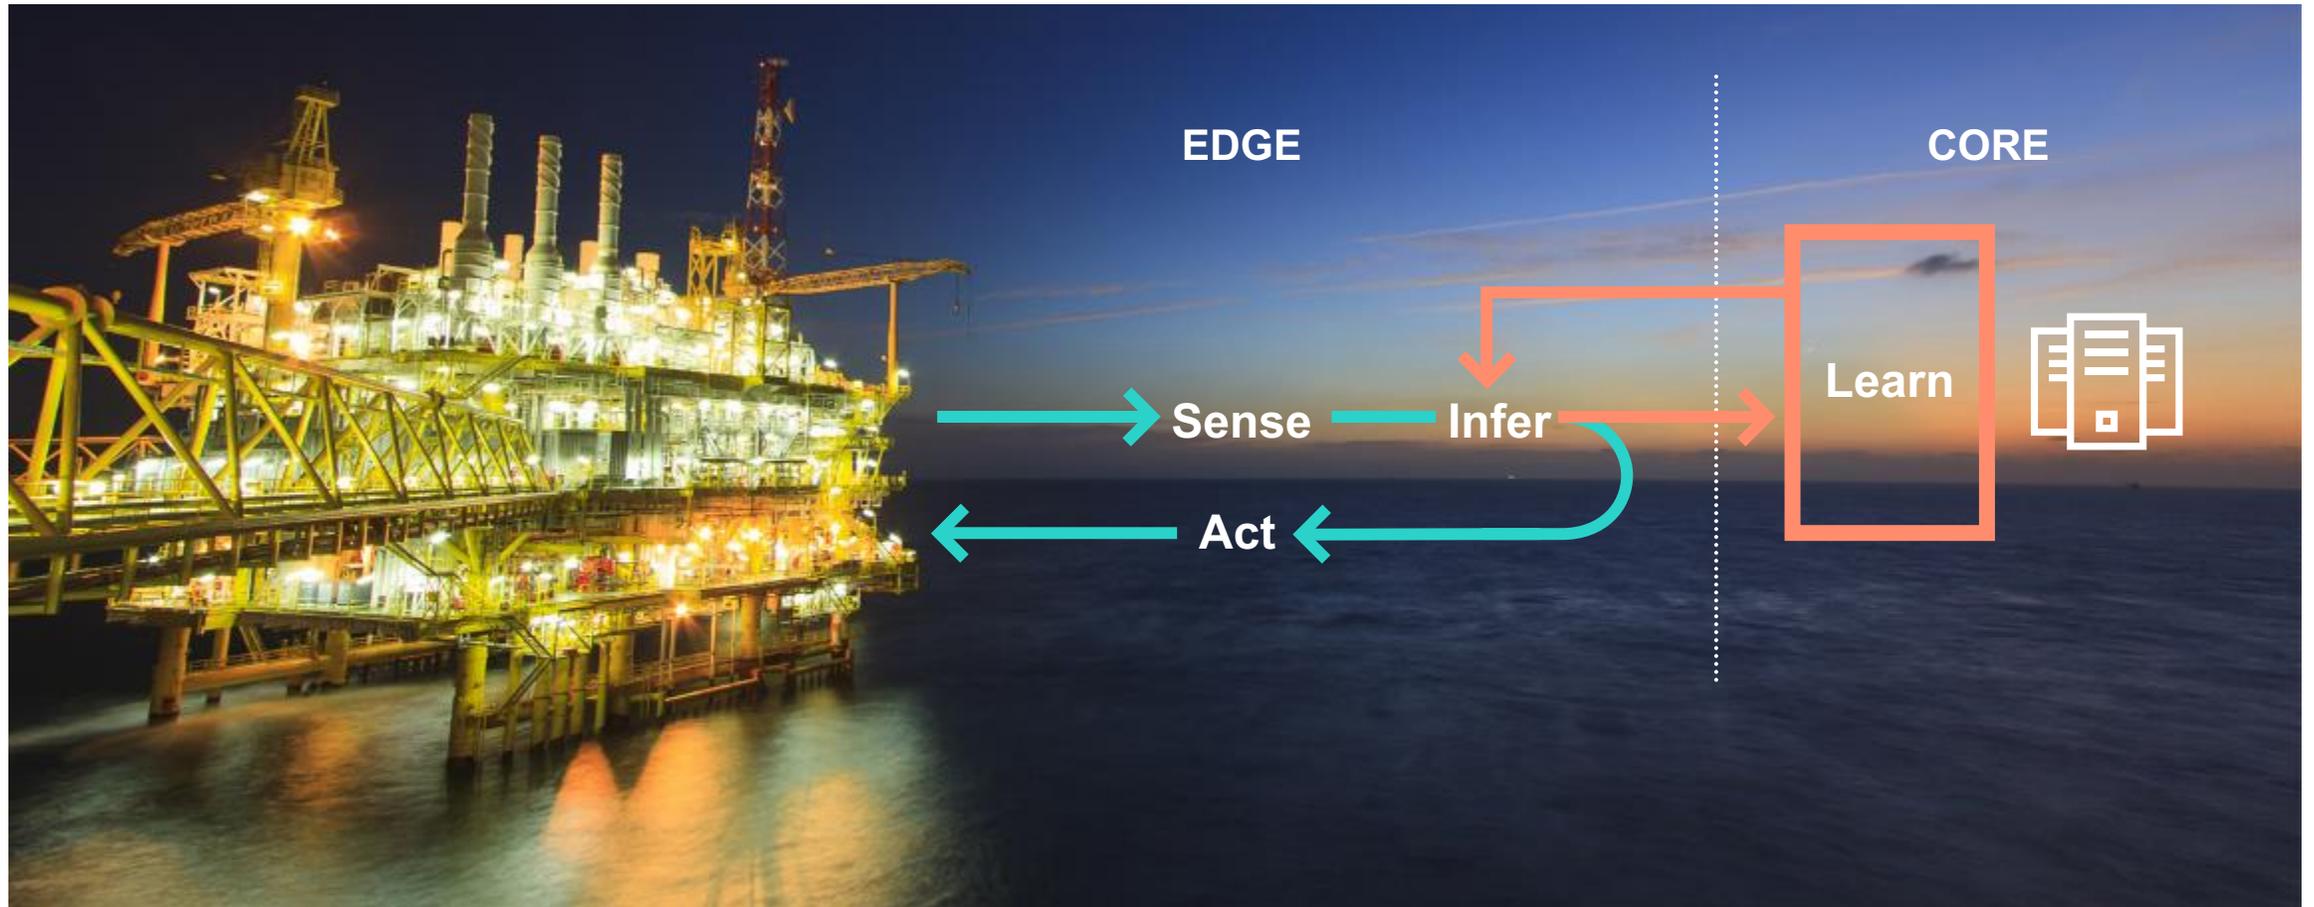
This is a rolling (up to three year) roadmap and is subject to change without notice.

This document contains forward looking statements regarding future operations, product development, product capabilities and availability dates. This information is subject to substantial uncertainties and is subject to change at any time without prior notification. Statements contained in this document concerning these matters only reflect Hewlett Packard Enterprise's predictions and / or expectations as of the date of this document and actual results and future plans of Hewlett Packard Enterprise may differ significantly as a result of, among other things, changes in product strategy resulting from technological, internal corporate, market and other changes. This is not a commitment to deliver any material, code or functionality and should not be relied upon in making purchasing decisions.

# Turn data into insights with advanced data and analytics platforms



# 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

# Taking action in real-time demands a powerful digital core



# 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



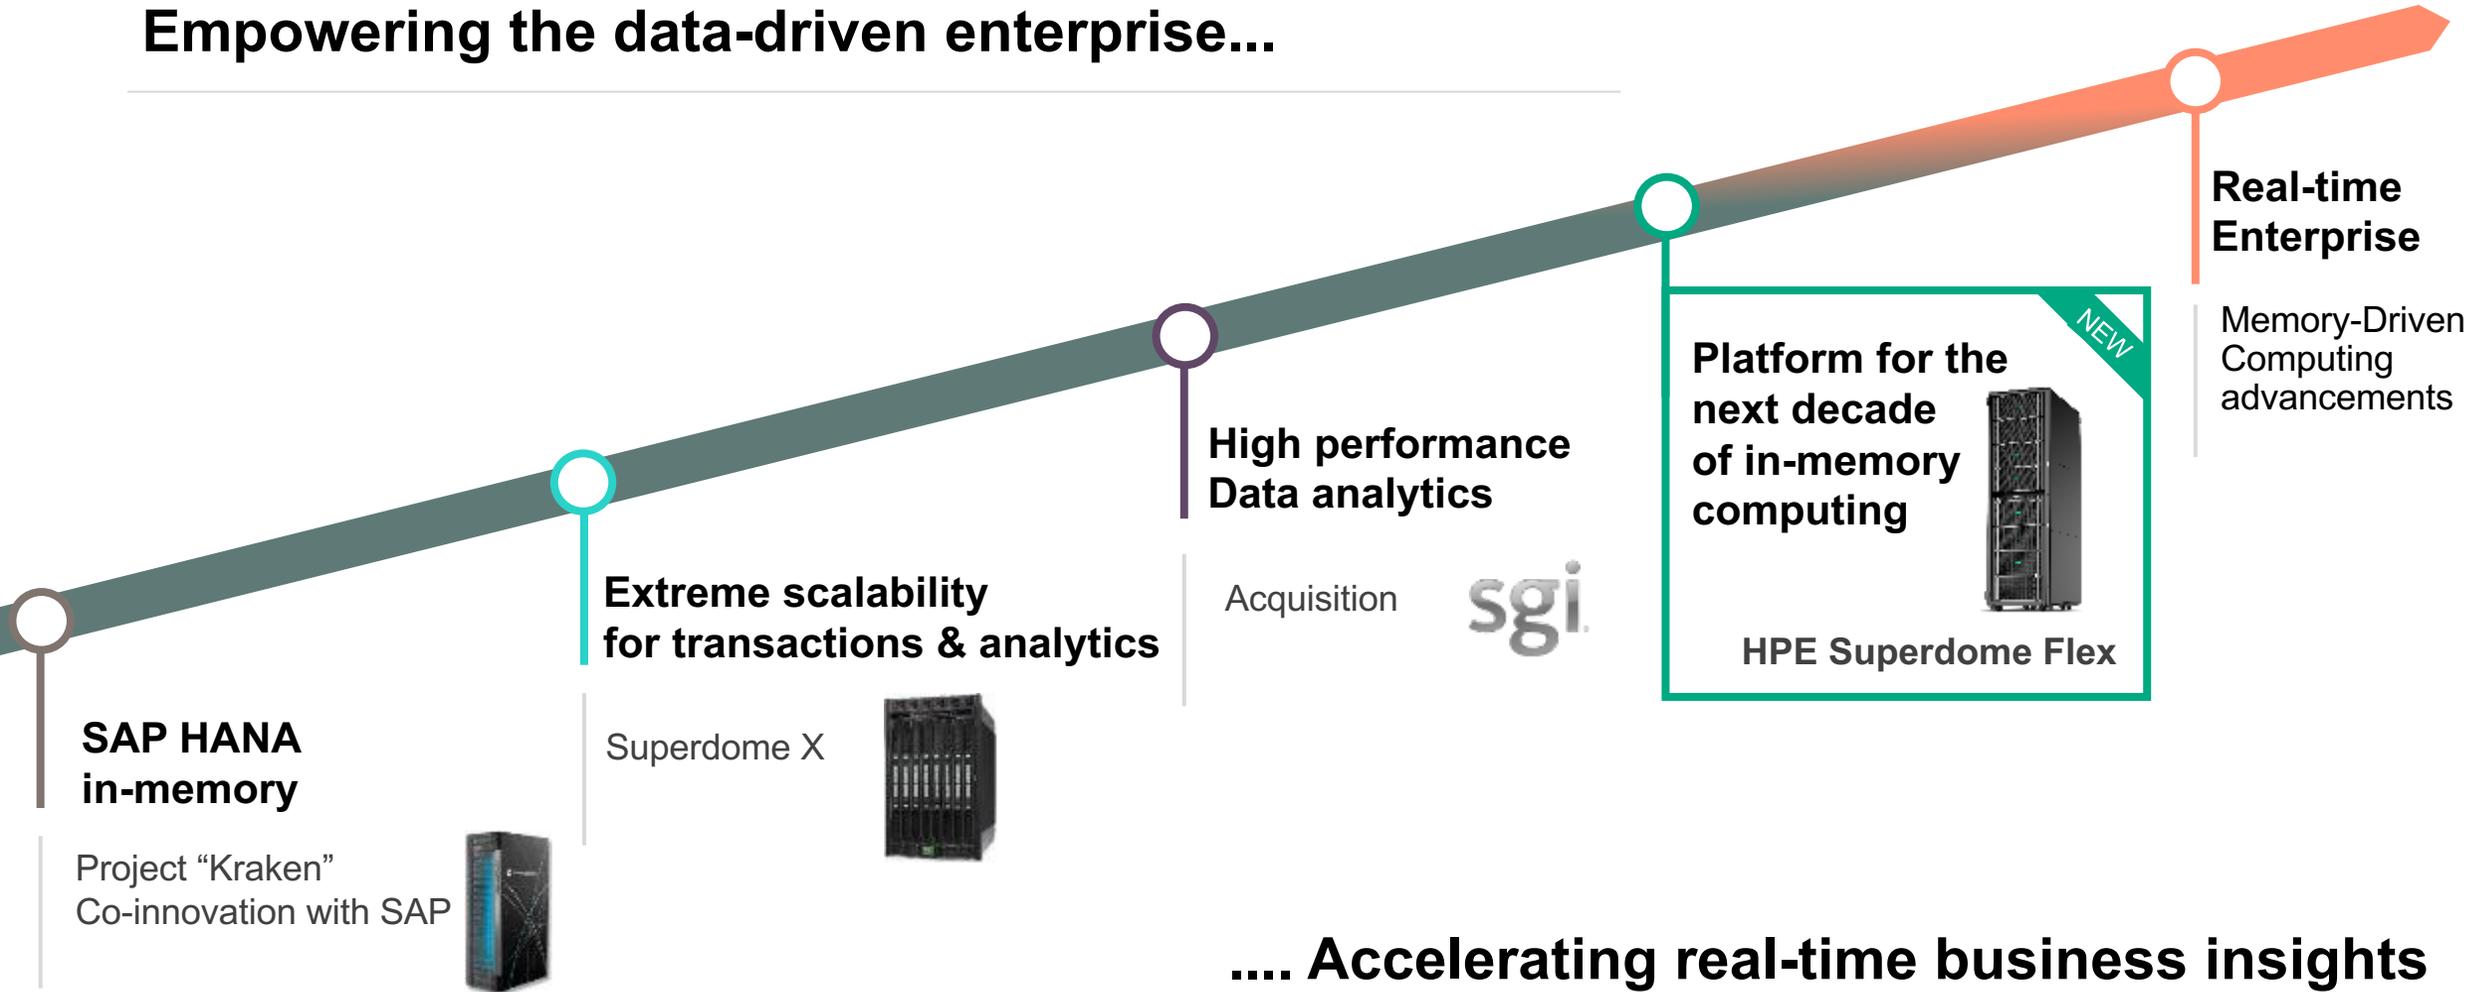
Real-time demand/supply matching



Immediate response to security threats

# Advancing the real-time enterprise journey

Empowering the data-driven enterprise...



.... Accelerating real-time business insights

# Visionary innovation to address the data dilemma

Embrace the possibilities of a world transformed by exponential data growth

Turn data into actionable insights in real time



Speed to insight

Keep pace with evolving business demands



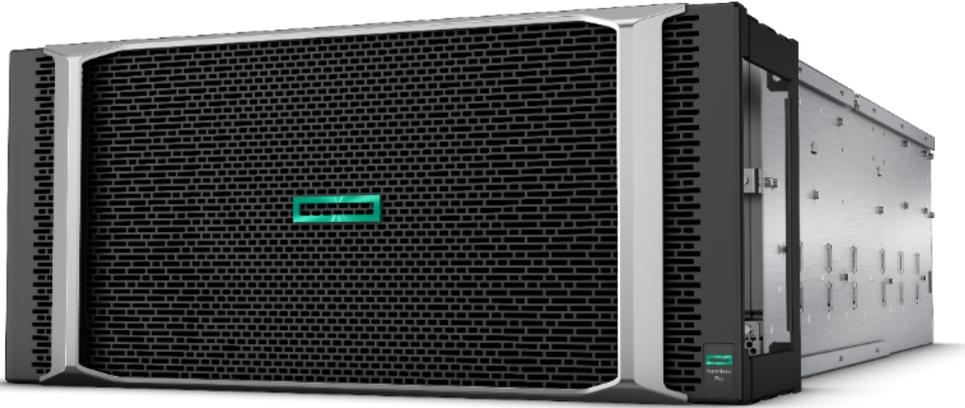
Flexibility and agility

Safeguard your mission-critical workloads



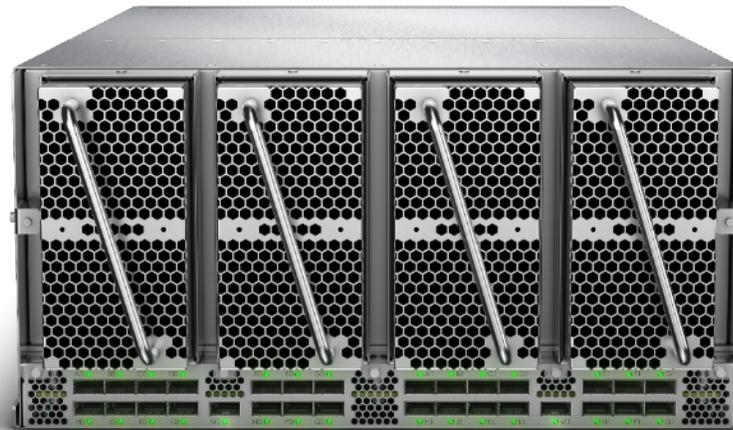
Continuous business

# HPE Superdome Flex Server: Chassis



Superdome Flex  
Chassis

Superdome Flex  
Full rack  
(32-socket configuration)



Superdome Flex Chassis  
(front – no bezel)

# 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 and hard partitioning 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



# Harnessing the full value of in-memory computing

# 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**

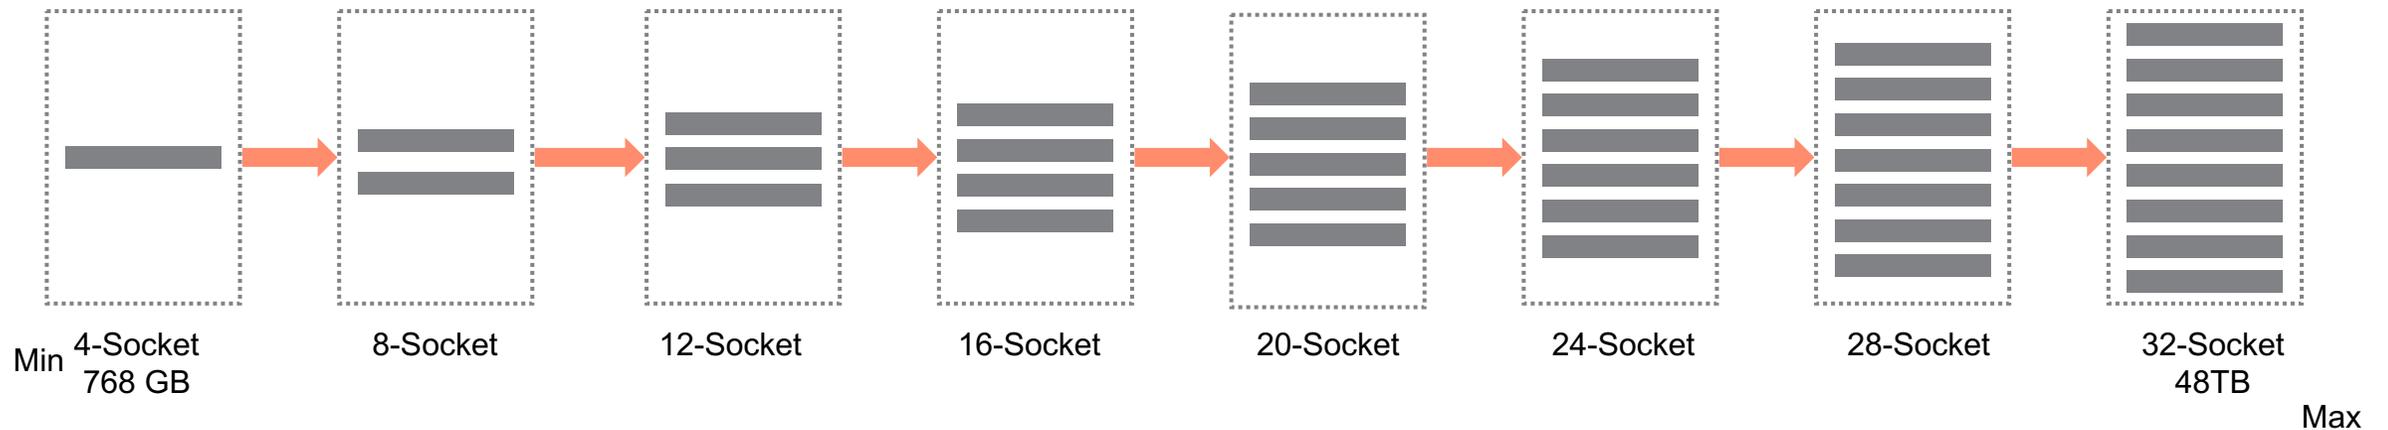


### Use cases include:

- 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

# 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**

# Unique point-to-point design maximizes performance

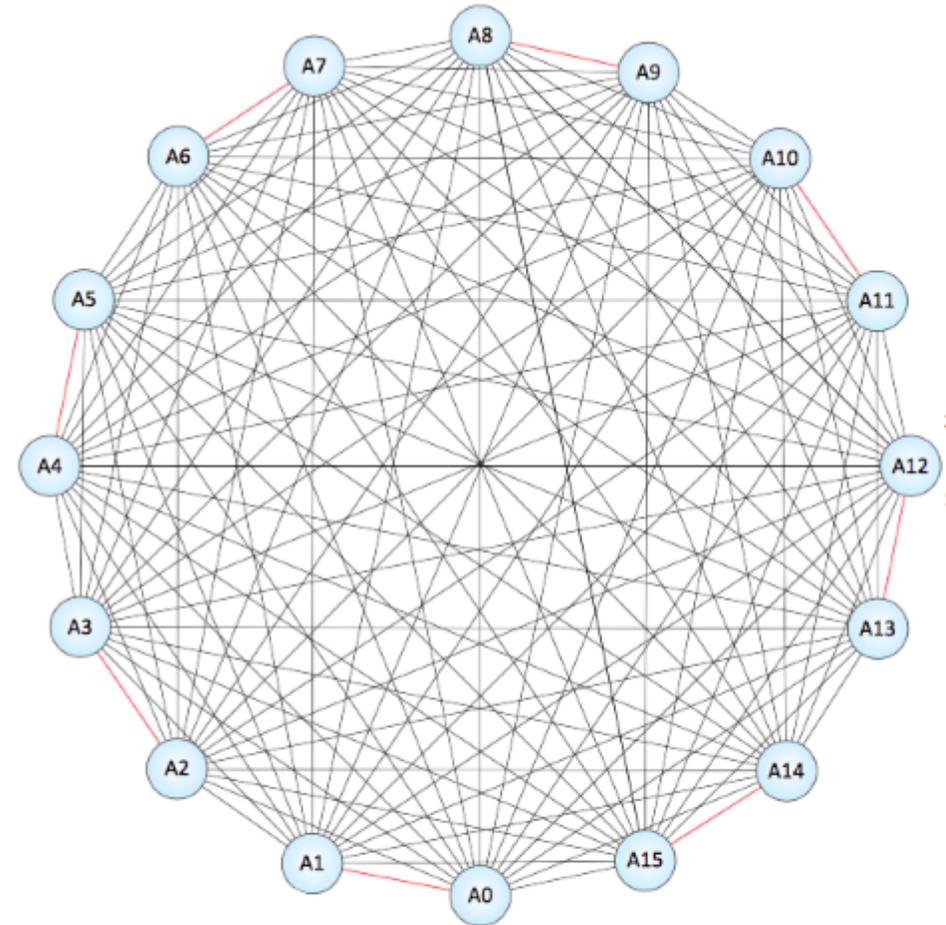
Extreme processing speed at scale

## 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



# HPE Superdome Flex system architecture

## Modular Design

- System designed around a 5U 4-socket building block which can be easily scaled from 4-32 sockets in 4-socket increments
- External Superdome Flex Grid interconnect cables provide a point-to-point interconnect fabric between chassis to guarantee low latency and high bandwidth
  - ~213.33 GB/s bisection crossbar bandwidth @ 8S
  - ~426.67 GB/s bisection crossbar bandwidth @ 16S
  - ~853.33 GB/s bisection crossbar bandwidth @ 32S
- New Flex Grid HA features coming to achieve fault tolerance and make online repairs possible
- Flexible stand-up PCIe Gen3 card format supported, in two different options that allow all low profile cards or 4 full-height/double-wide (up to 300W) cards

### Compute

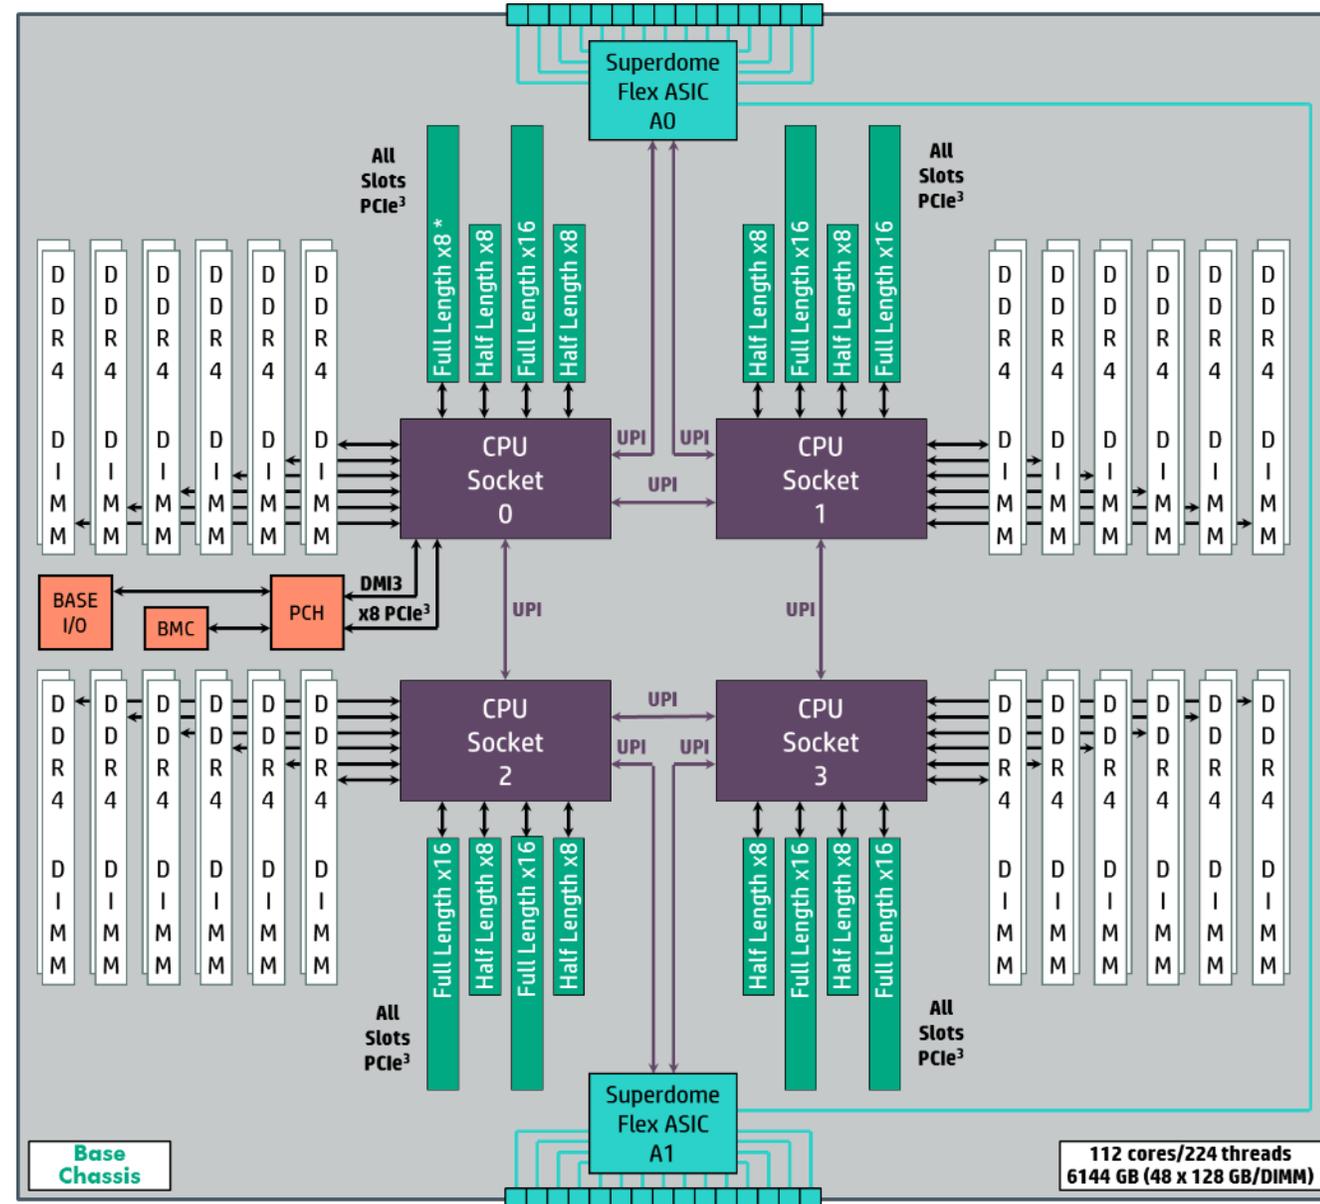
32 CPUs  
896 cores  
1792 threads

### Memory

384 DIMMs  
48 TB capacity

### I/O

72 (x8) slots  
56 (x16) slots  
boot storage



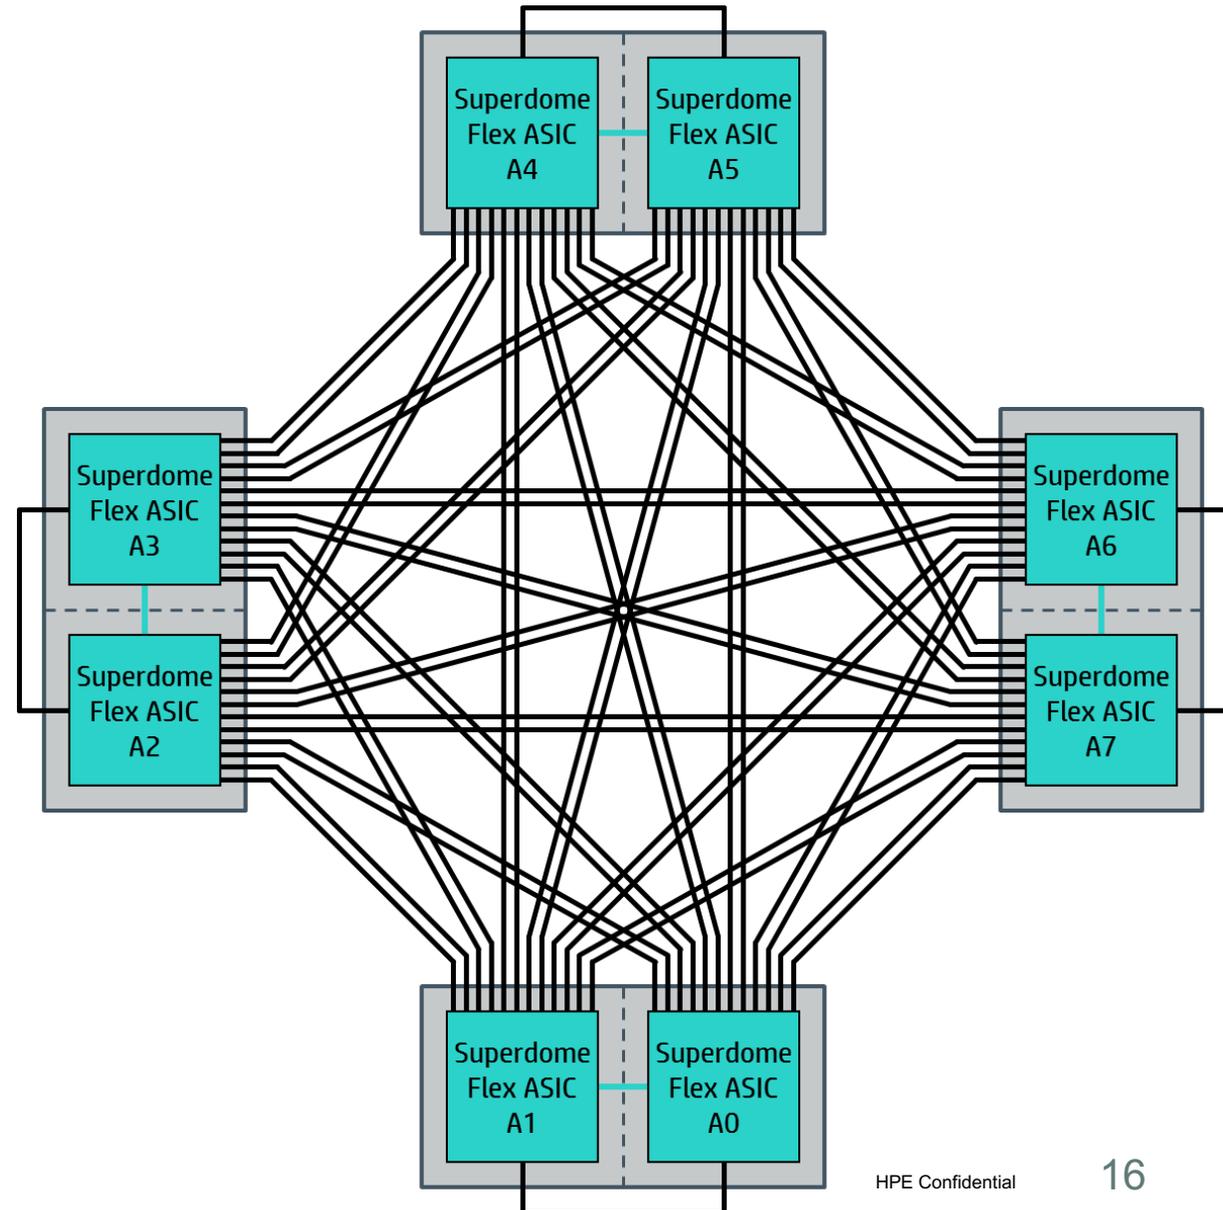
# HPE Superdome Flex Server: 16-socket architecture

## Scaling up to 16-sockets:

- Takes 4 chassis (RMC required)
- Includes 8 Superdome Flex ASICs
- Uses 52 Superdome Flex Grid cables
- Two point-to-point Superdome Flex Grid links between ASICs
- Provides ~426.67 GB/s of bisection BW (32 links are cut)

## Compute resources provided:

- 16 CPUs (up to 448 cores / 896 threads)
- 192 DIMM slots (up to 24TB w/128GB DIMMs)
- 64 PCIe Gen3 card slots (28 x16, 36 x8) max.



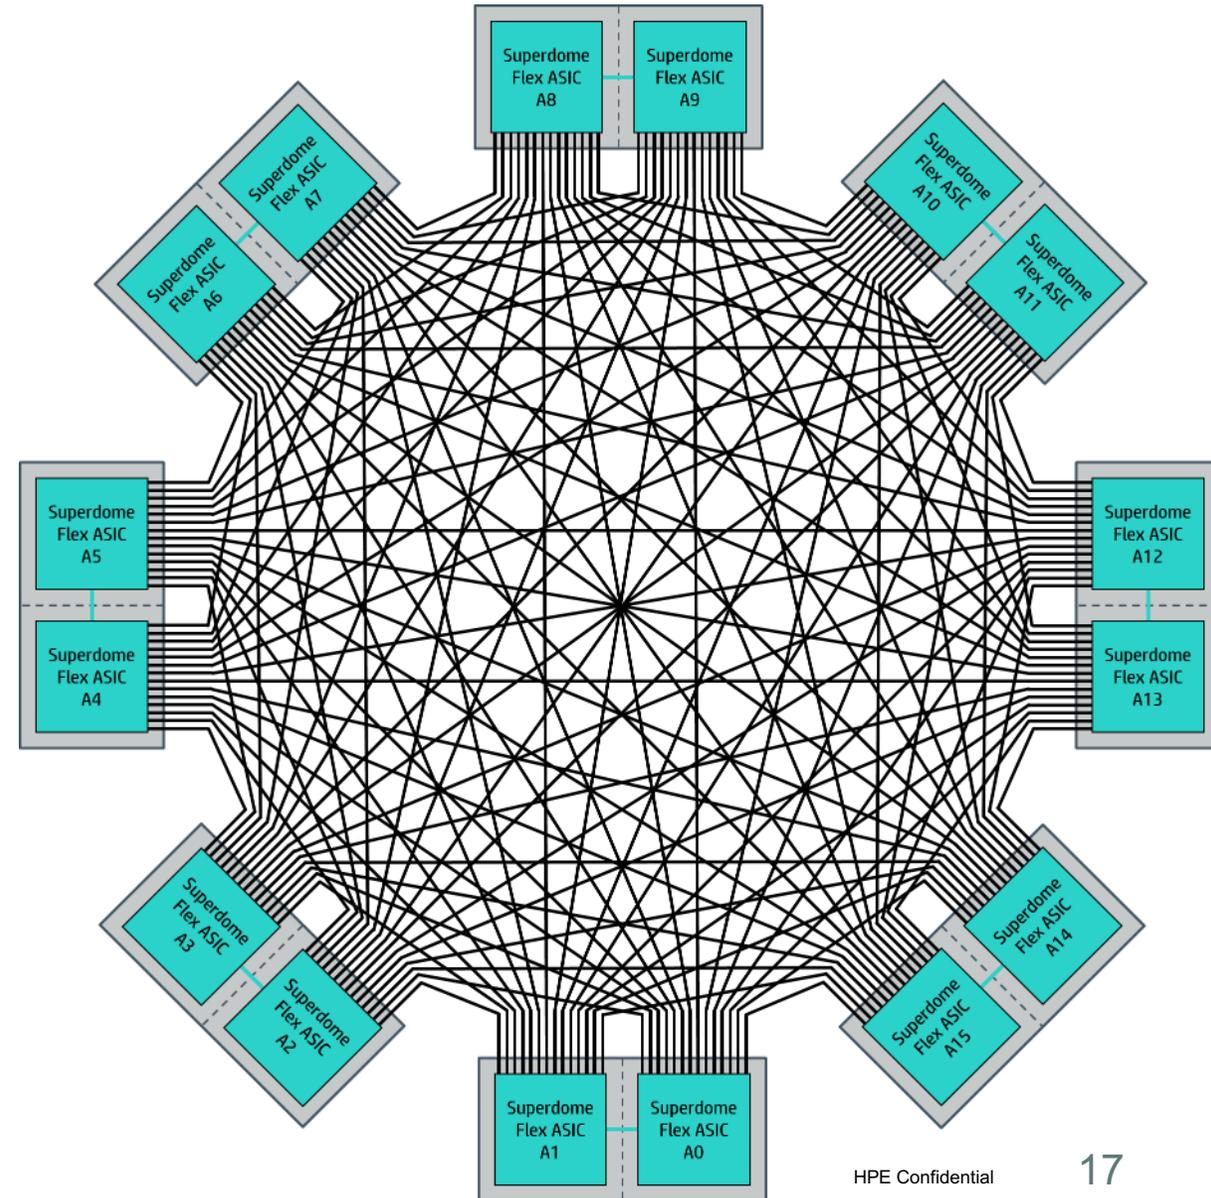
# HPE Superdome Flex Server: 32-socket architecture

## Scaling up to 32-sockets:

- Takes 8 chassis
- Includes 16 Superdome Flex ASICs
- Uses 112 Superdome Flex Grid cables
- One point-to-point Superdome Flex Grid link between ASICs
- Provides ~853.33 GB/s of bisection BW (64 links are cut)

## Compute resources provided:

- 32 CPUs (up to 896 cores / 1792 threads)
- 384 DIMM slots (up to 48TB w/128GB DIMMs)
- 128 PCIe Gen3 card slots (56 x16, 72 x8) max.





**Flexible compute power for the apps that  
are the lifeblood of your enterprise**

# Flexibility to scale, configure, convert, deploy

The possibilities are endless

## 1 Meet your individual workload requirements



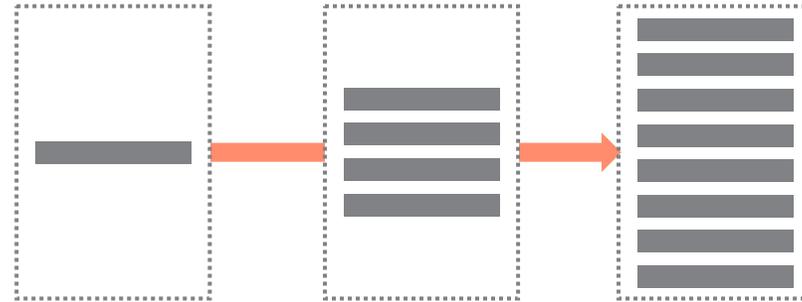
4-Socket

### Configuration choices:

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

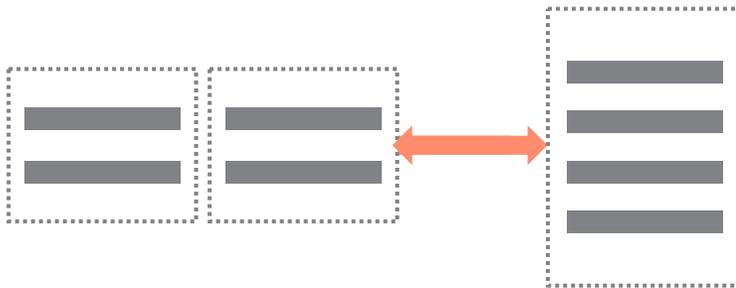
Individual "building block" configuration choice

## 2 Grow seamlessly with your business demands



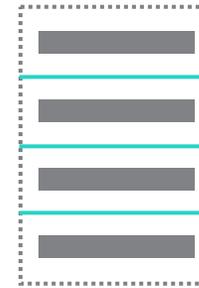
Never outgrow, avoid over-provisioning: 4-32 socket scale

## 3 Easily transform your environment: Scale-out or up



Convert between scale-out and scale-up

## 4 Deploy separate environments in the same frame



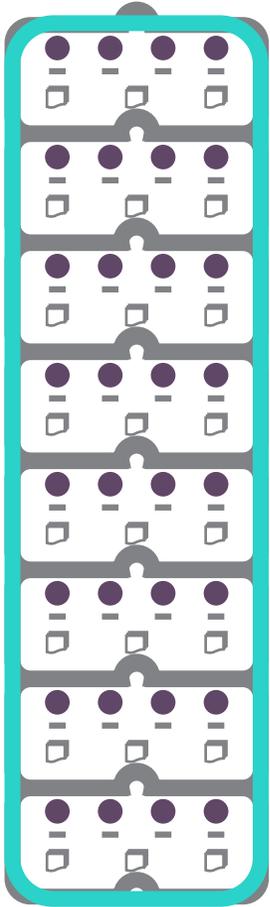
### Use case examples:

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

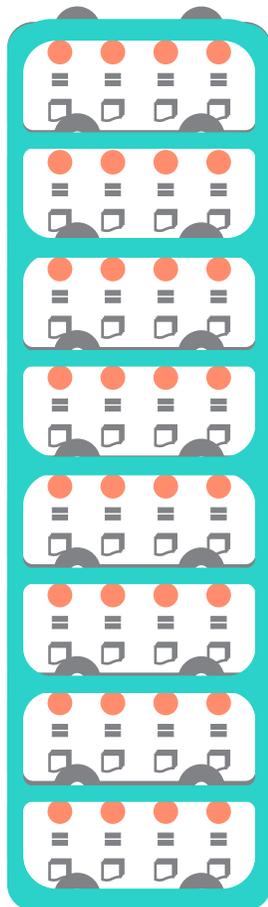
Support for electrically isolated hard partitions

# 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**

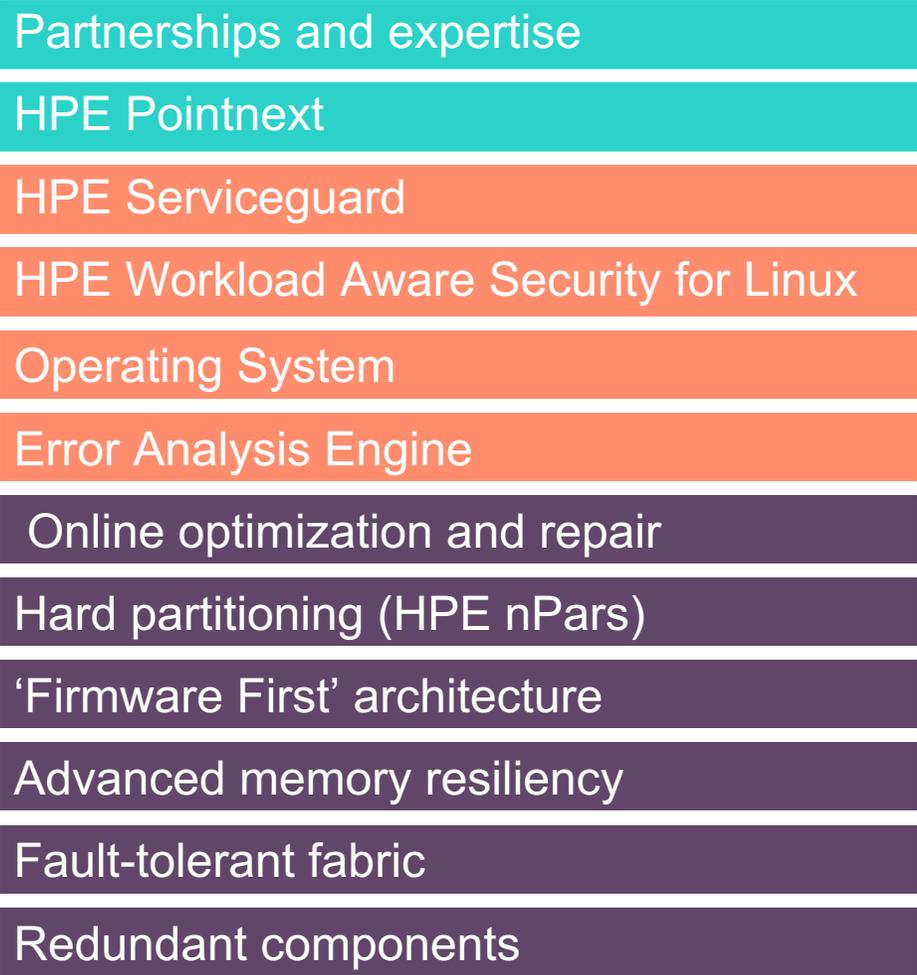
### Partitionable to deploy separate environments in the same system

- 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

**Reliability and expertise to safeguard your  
critical workloads**

# Safeguarding your 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.

# 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

<b>Memory</b>	Advanced memory resiliency (ADDDC +1)
<b>I/O</b>	Advanced PCIe error recovery (LER)
<b>Fabric</b>	Enhanced fabric resiliency (hot cable swap)
<b>Processor</b>	Advanced processor error handling (eMCA Gen2)



# Fortify Linux workloads with HPE Serviceguard

Industry leading high availability and disaster recovery solution for critical x86 workloads

## Automated

- Zero scripting integration with major databases and applications
- Site Aware disaster recovery solutions for complex, distributed workloads

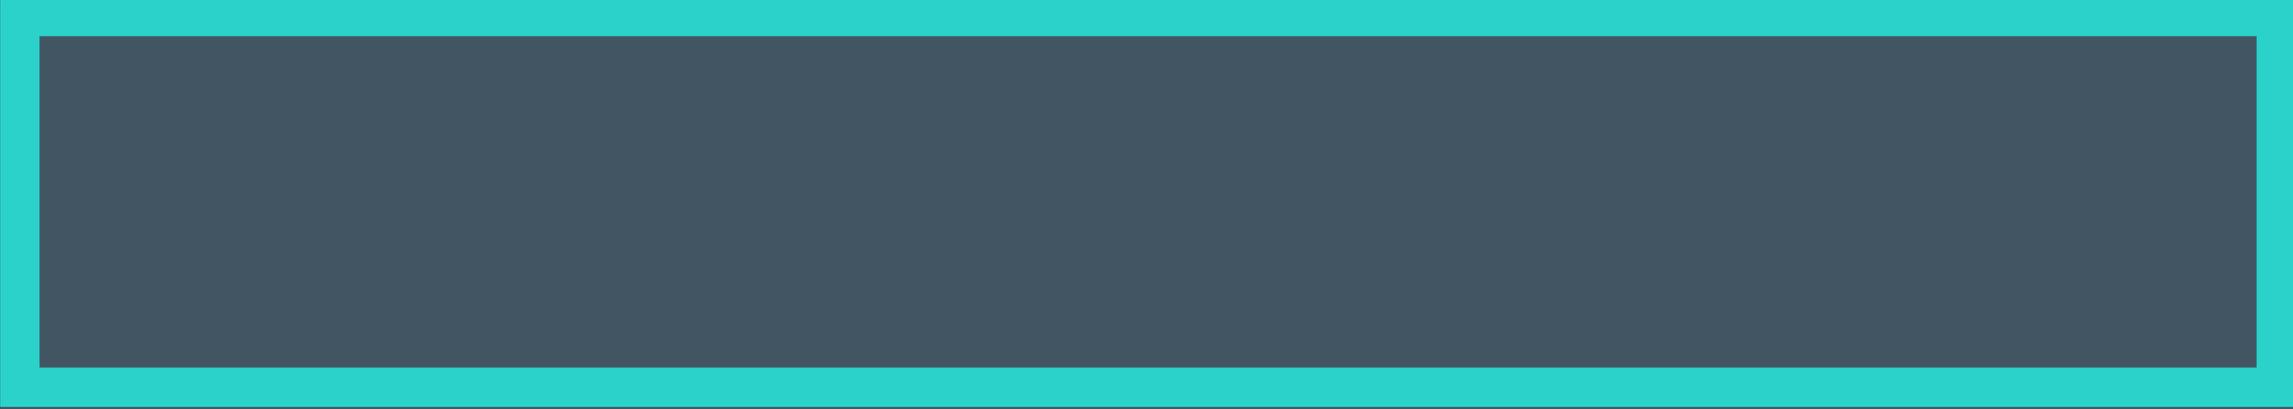
## Efficient

- Intelligent clustering and reduced capex for disaster recovery with smart quorum
- Enhanced and powerful GUI for analytics, simulation and verification

## Versatile

- Support for traditional (Oracle, SQL & DB2) and in-memory databases (SAP HANA)
- Supports RHEL and SLES distros with any x86\_64 server
- VMware and KVM hypervisor support

Provides out-of-the-box integration with applications, upholds data-integrity, minimizes downtime



# Specifications and roadmap

# Superdome Flex Specifications

	Description												
<b>System</b>	1 to 8 chassis; each supports four (4) Intel® Xeon® Scalable processors												
<b>Processors</b> (Available at initial release)	<table border="0"> <tr> <td>Intel Xeon Platinum 8180 processor</td> <td>28-cores/2.5GHz/205W/38.5M</td> </tr> <tr> <td>Intel Xeon Platinum 8176 processor</td> <td>28-cores/2.1GHz/165W/38.5M</td> </tr> <tr> <td>Intel Xeon Platinum 8156 processor</td> <td>4-cores/3.6GHz/105W/16.5M</td> </tr> <tr> <td>Intel Xeon Platinum 8158 processor</td> <td>12-cores/3.0GHz/150W/16.5M</td> </tr> <tr> <td>Intel Xeon Gold 6154 processor</td> <td>18-cores/3.0GHz/200W/24.75M</td> </tr> <tr> <td>Intel Xeon Gold 6132 processor</td> <td>14-cores/2.6GHz/133W/19.25M</td> </tr> </table>	Intel Xeon Platinum 8180 processor	28-cores/2.5GHz/205W/38.5M	Intel Xeon Platinum 8176 processor	28-cores/2.1GHz/165W/38.5M	Intel Xeon Platinum 8156 processor	4-cores/3.6GHz/105W/16.5M	Intel Xeon Platinum 8158 processor	12-cores/3.0GHz/150W/16.5M	Intel Xeon Gold 6154 processor	18-cores/3.0GHz/200W/24.75M	Intel Xeon Gold 6132 processor	14-cores/2.6GHz/133W/19.25M
Intel Xeon Platinum 8180 processor	28-cores/2.5GHz/205W/38.5M												
Intel Xeon Platinum 8176 processor	28-cores/2.1GHz/165W/38.5M												
Intel Xeon Platinum 8156 processor	4-cores/3.6GHz/105W/16.5M												
Intel Xeon Platinum 8158 processor	12-cores/3.0GHz/150W/16.5M												
Intel Xeon Gold 6154 processor	18-cores/3.0GHz/200W/24.75M												
Intel Xeon Gold 6132 processor	14-cores/2.6GHz/133W/19.25M												
<b>Memory</b>	48 DDR4 DIMM slots per chassis Maximum memory: 3 TB (48x 64 GB DIMMs) per chassis 32 GB and 64 GB DDR4 DIMMs loaded in groups of 12 DIMMs												
<b>Base IO</b> (base chassis)	2x 10GbE ports, 2x 1GbE ports, 4x USB 3.0 ports, serial and MGMT ports												
<b>Internal drive slots</b>	Up to four (4) 2.5" SATA/SAS HDD or SSD with option for hardware RAID												
<b>IO expansion options</b>	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												
<b>Management</b>	Optional 1U Rack Management Controller (RMC) for CLI; Redfish® API												
<b>Operating systems</b>	SUSE® Linux® Enterprise Server 12, Red Hat® Enterprise Linux 7, Oracle Linux 7												
<b>Form Factor</b>	5U server chassis; width: 17.5" (44.5cm); depth: 32.5" (82.6cm)												

# Superdome Flex supported storage

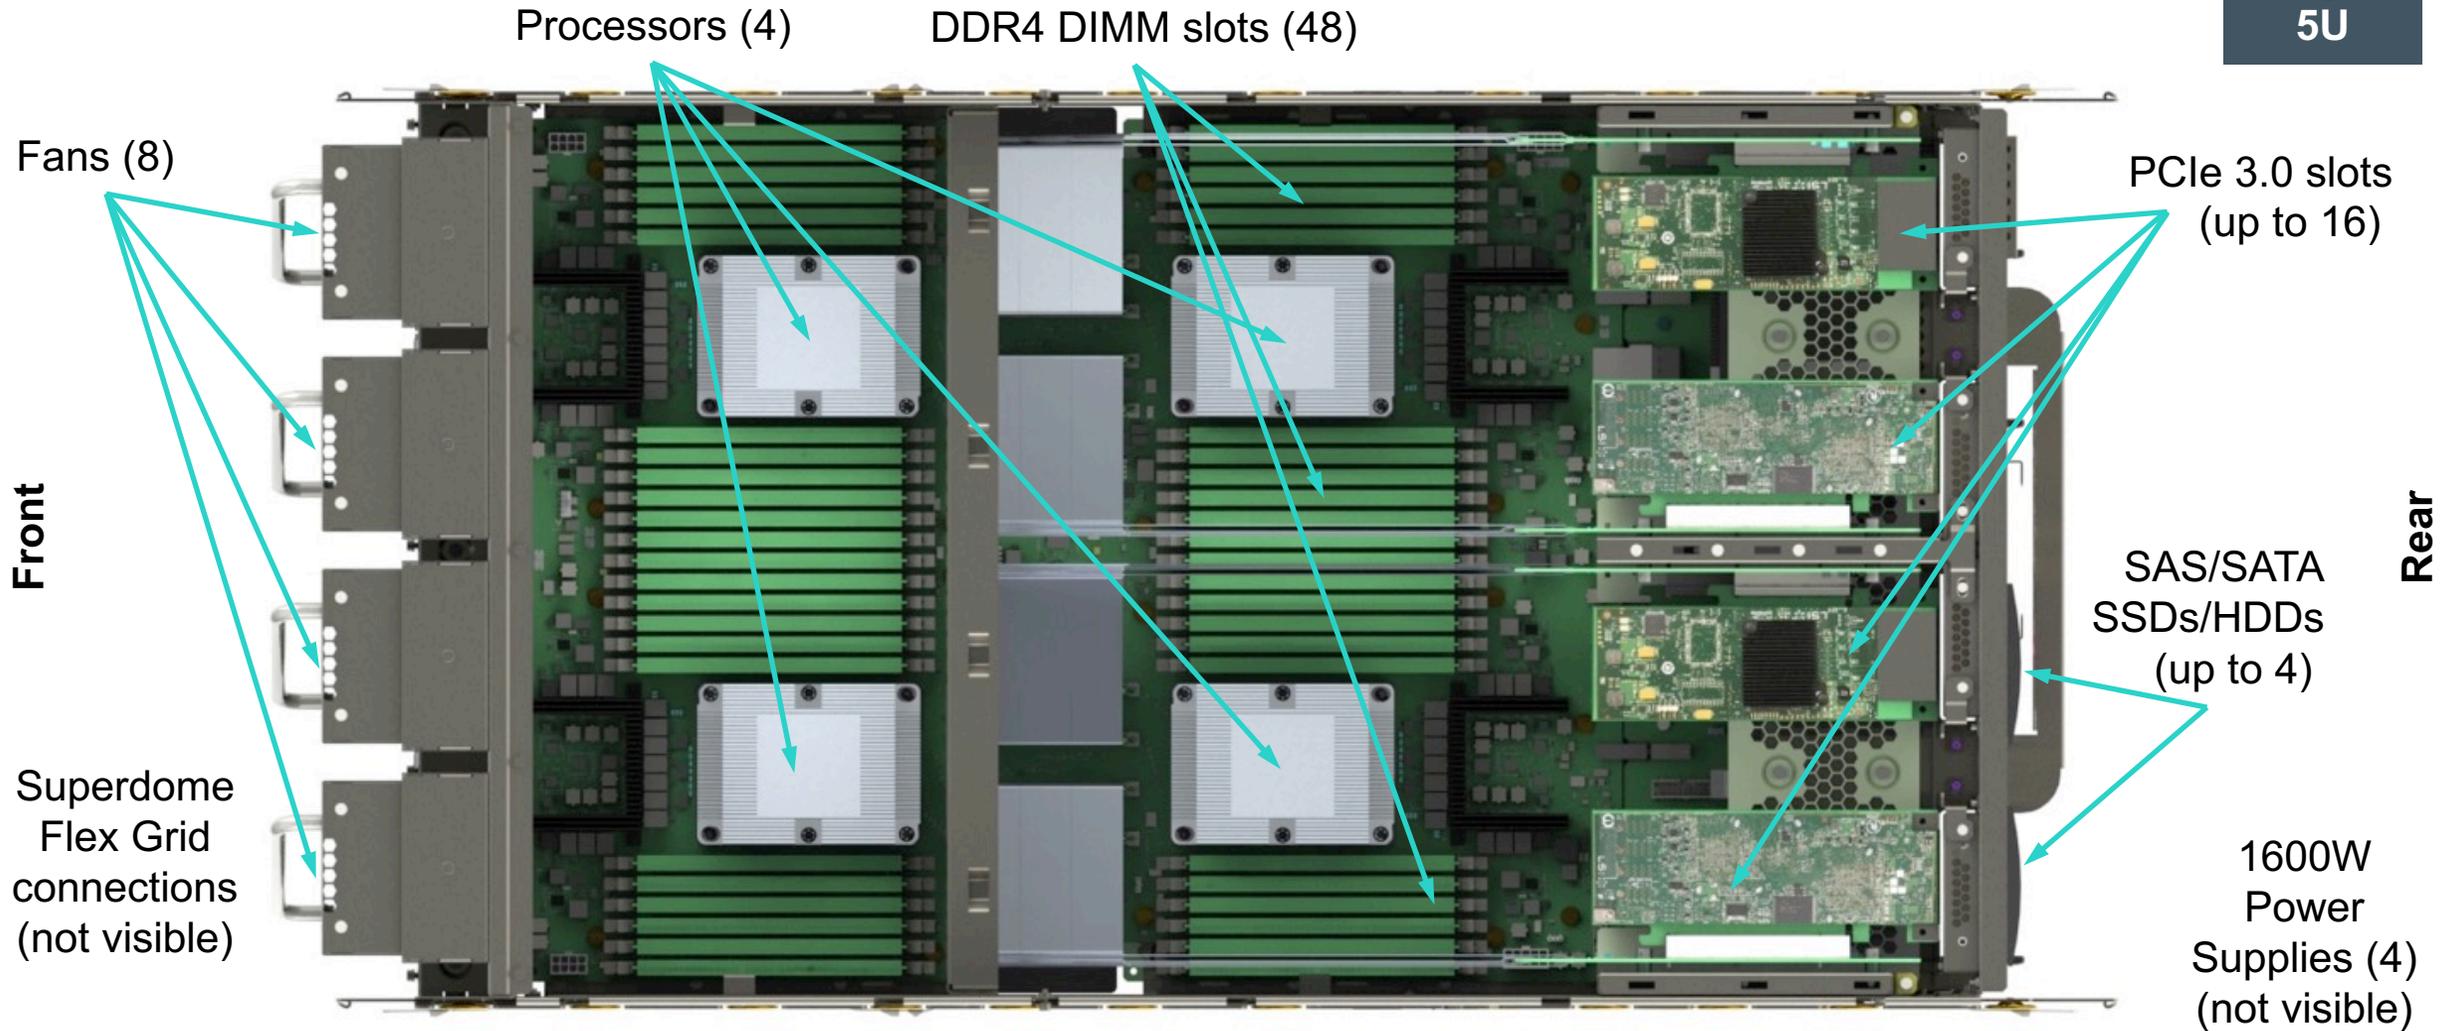
- **Internal storage:** Four (4) 2.5” drive bays to support SATA SSDs or SAS HDDs/SSDs.
  - 6G SATA SSDs use embedded chip (Intel RSTe) with SW RAID (w/boot support)
  - 12G SAS HDDs/SSDs use PCIe RAID card (internal) with HW RAID (w/boot support)
- **SAS:** HPE SAS JBOD (e.g. D3700) are supported with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



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

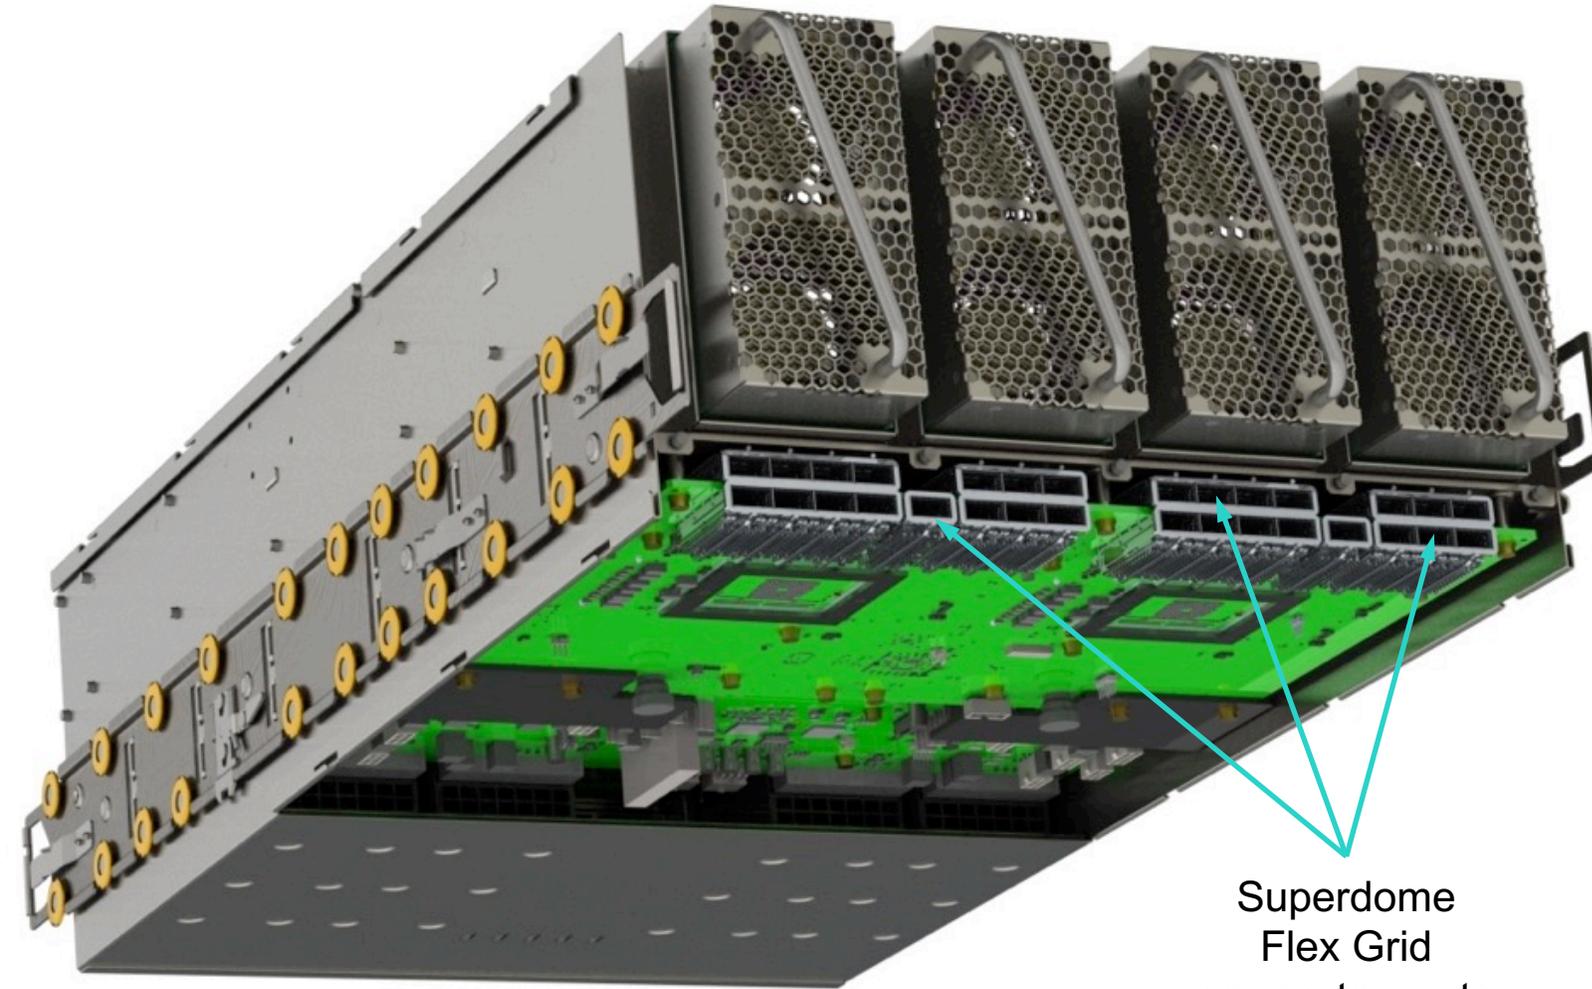
# HPE Superdome Flex: Chassis Internal Top View

Height:  
5U



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

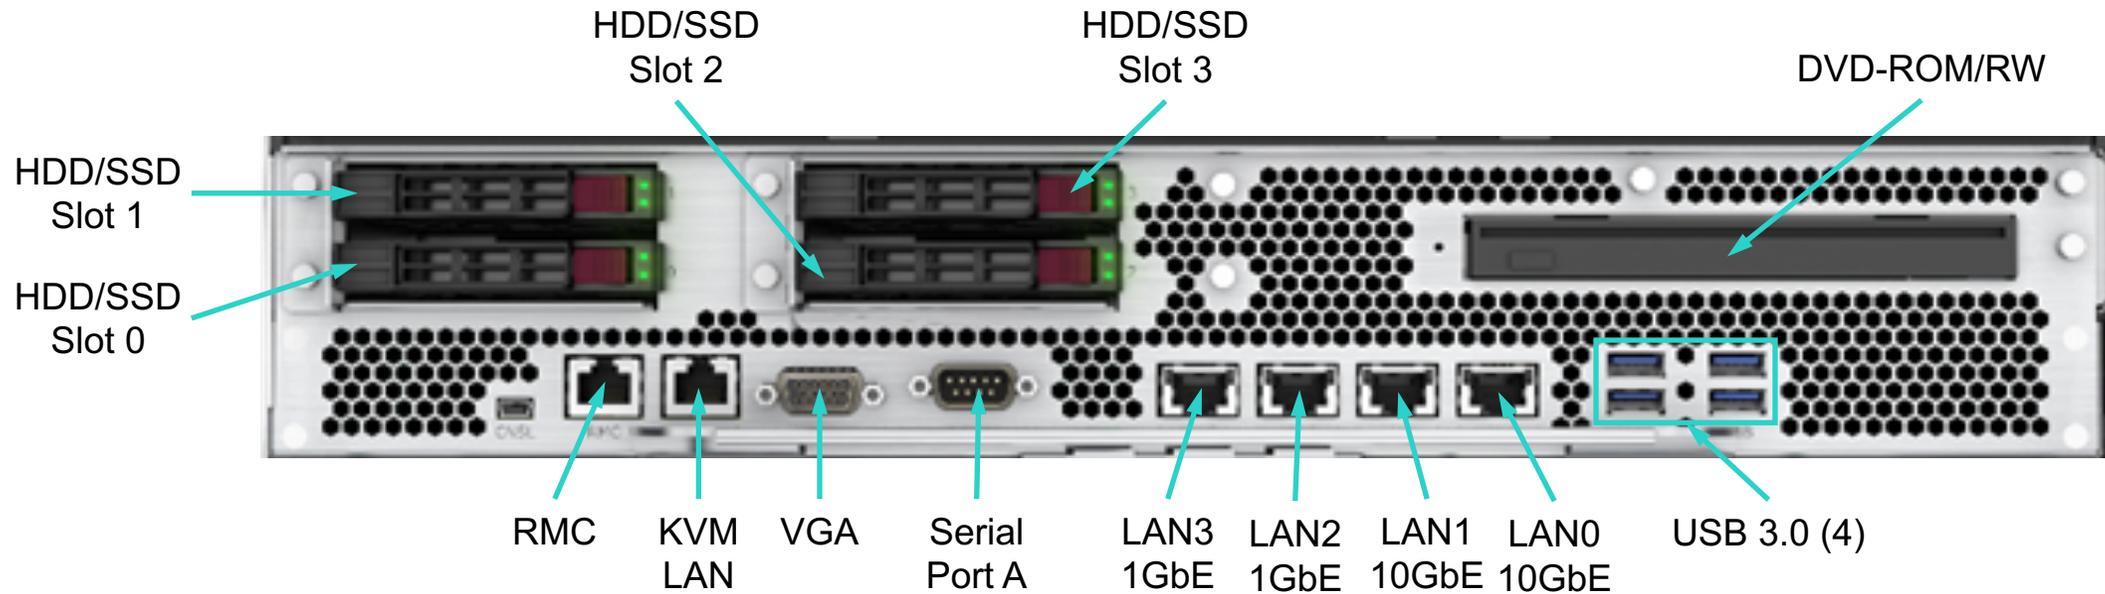
# HPE Superdome Flex: Chassis Internal Bottom View



Superdome  
Flex Grid  
connector ports

- **Superdome Flex Grid connections made at the front of the enclosure**
- **30 Superdome Flex Grid connectors**
  - Cables (IB-EDR) span between chassis to complete fabric interconnection
  - New features provide fault-tolerance, and enable online cable repair
- **8 hot-swappable fans**

# HPE Superdome Flex I/O subsystem: Base I/O



## Base I/O card provides access to:

- One DVD-ROM/RW drive
- The Board Management Controller (BMC)
  - 1 GbE Rack Management Controller (RMC) port
  - 1 GbE Management (KVM) LAN port
- The Platform Controller Hub (PCH)
- Built-in boot storage
  - Four 2.5" hot-swappable HDD/SSDs
- Console Serial, VGA and USB 3.0 ports (4)
- Two 10 GbE LAN ports
- Two 1 GbE LAN ports

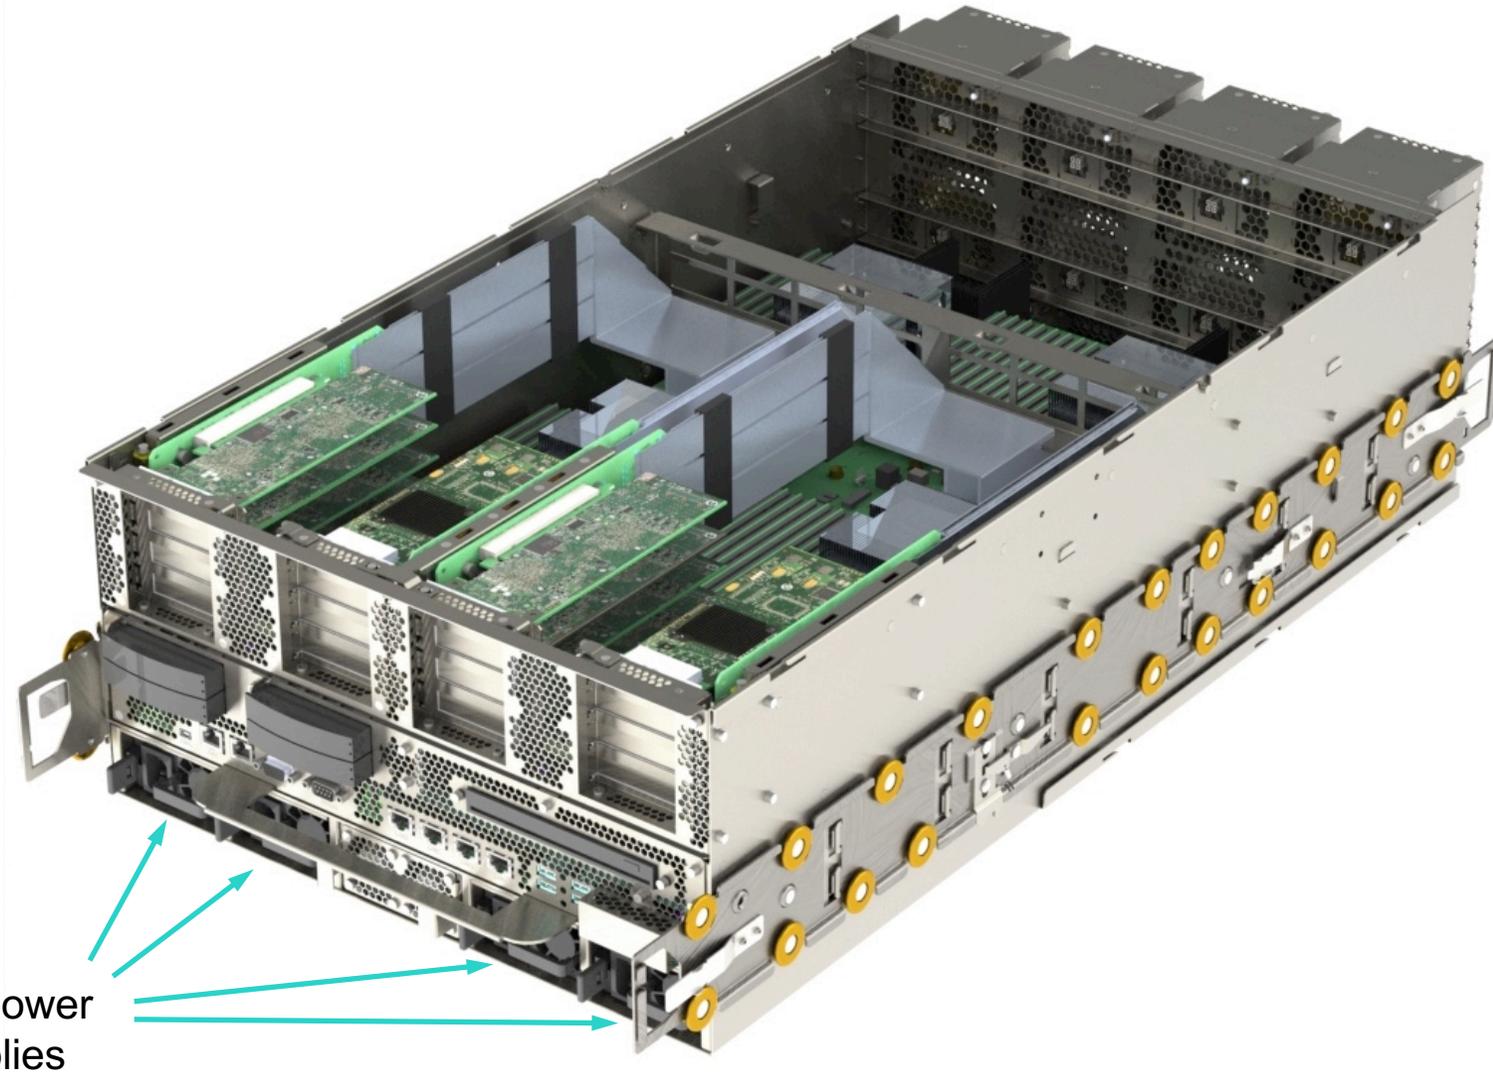
# Superdome Flex I/O subsystem: 16-slot PCIe 3.0 configuration

## 16 low profile card slot option

- 7 full length x16 slots available
- 9 full length x8 slots available

## Uses all available PCIe channels from CPUs !!!

- I/O designs in competitor systems compromise the number of PCIe 3.0 channels they actually connect to
- Superdome Flex design has unique Z-axis approach to escape those channels from the center of the system board
- All done w/o bus repeaters or other performance limiters



---

# The Intel® Xeon® Scalable processor – Skylake

## Some exciting new capabilities and some interesting limitations

### New attributes

- Supports new Intel® Ultra Path Interconnect (UPI)
  - Increased bandwidth and performance over QPI
  - Improved messaging efficiency, multiple requests per packet
- SKUs fit into 1 of 4 available tiers
  - Platinum: top bin (up to 28 cores, highest frequency, largest cache, full power), designed for 8s or smaller, fully featured, and consequently, most expensive
  - Gold: high bin (up to 22 cores, high frequency, standard cache,  $\frac{3}{4}$  power), designed for 4s or smaller, fully featured, but more value oriented in pricing terms
  - Silver/Bronze: designed for 2s or smaller, N/A on Superdome Flex
- Memory Performance improvements
  - Support for DDR4 RDIMMs, LR-DIMMs, and TSV (3DS) DIMMs
  - Direct attach from CPU to DIMM provides lower latency and higher bandwidth via faster data transfer speeds (up to 2667 MT/s)
  - 6 fully independent channels (50% more than E7 v4)
- New Intel® AVX-512 (512 bit) instruction set
  - up to 2x improvement in peak floating point performance over AVX2 (E7 v3) results in arguably the best HPC engine ever
- Extended I/O support:
  - 48 PCIe Gen 3.0 lanes (x16, x8, x4 bifurcatable) – 50% more than E7 v4
  - New MCTP Scaling supports 256 PCIe\* buses, and up to 8 segments

### The “downside” worth mentioning

- New Socket P is 85% larger and no longer tool-less
- UPI incompatible with QPI, so no support for prior chipsets
- Rebalanced Cache Hierarchy (MLC Optimization)
  - Shift cache balance by enlarging mid-level cache
  - Translates into a significantly smaller L3 shared cache
- Per socket memory capacities above 768 GB (i.e. support for 128+ GB DIMMs) requires a special “M” SKU
  - These exclusive SKUs are even more expensive !!!
- No external scalable memory buffer available, so can only support half as many DIMMs per socket (12 not 24)
  - Reduces capacity flexibility since only half (6) or full (12) populations per socket deliver full performance
- No support for Intel® Optane™ memory DIMMs
  - Have to wait for the next processor release, which also requires some HPE Superdome Flex ASIC changes

# HPE Superdome Flex to support multiple processor SKUs

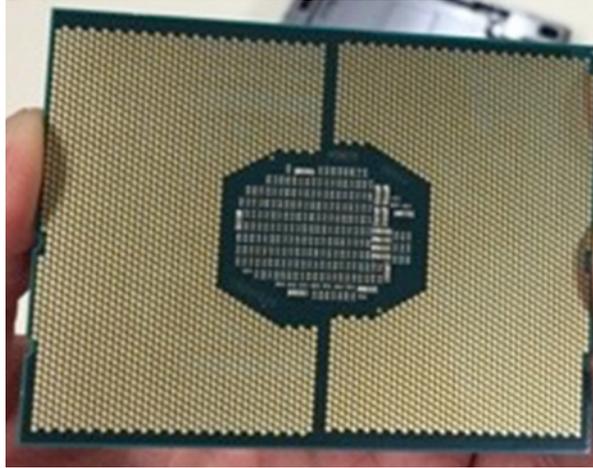
Six Intel Xeon Scalable processors for first release

## 28-core processors

- [Platinum 8180](#) (14 nm)
  - 28 Cores / 56 Threads
  - 2.5 GHz (3.8 GHz) Core Frequency
  - 38.5 MB cache, 10.4 GT/s UPI, 205 W
- [Platinum 8176](#) (14 nm)
  - 28 Cores / 56 Threads
  - 2.1 GHz (3.8 GHz) Core Frequency
  - 38.5 MB cache, 10.4 GT/s UPI, 173 W

## 18-core processor \*

- [Gold 6154](#) (14 nm)
  - 18 Cores / 36 Threads
  - 3.0 GHz (3.7 GHz) Core Frequency
  - 24.75 MB cache, 10.4 GT/s UPI, 200 W



Note: All listed CPU SKUs have 2 embedded Memory Controllers (DDR4), are 46-bit physical address capable and possess 48 embedded PCIe Gen 3 lanes

\* Gold 61xx processors utilized only in 2S+2S system configurations

## 12-core processor

- [Platinum 8158](#) (14 nm)
  - 12 Cores / 24 Threads
  - 3.0 GHz (3.7 GHz) Core Frequency
  - 24.75 MB cache, 10.4 GT/s UPI, 150 W

## 4-core processor

- [Platinum 8156](#) (14 nm)
  - 4 Cores / 8 Threads
  - 3.6 GHz (3.7 GHz) Core Frequency
  - 16.5 MB cache, 10.4 GT/s UPI, 105 W

## 14-core processor \*

- [Gold 6132](#) (14 nm)
  - 14 Cores / 28 Threads
  - 2.6 GHz (3.7 GHz) Core Frequency
  - 19.25 MB cache, 10.4 GT/s UPI, 140 W

---

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

## HPE Superdome Flex RAS features at a glance

### Chassis-Level features

- Firmware-first error handling
- Redundant, hot-swappable power supplies (N+N or N+1) and fans
- Flex Grid link failover, link-level retry, dynamic link tuning and bandwidth negotiation
- Adaptive routing finds bad fabric links and routes traffic around failures
- CRC protection per micro-packet and fast retry for transient errors
- Systemic transient errors, triggering retry or recovery attempt
- Socket indictment at boot [coming soon w/ limitations]
- Chassis deconfigure at boot [coming soon]
- Hard Partitions (nPars) [coming soon]

### Processor coverage

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

### Memory features

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

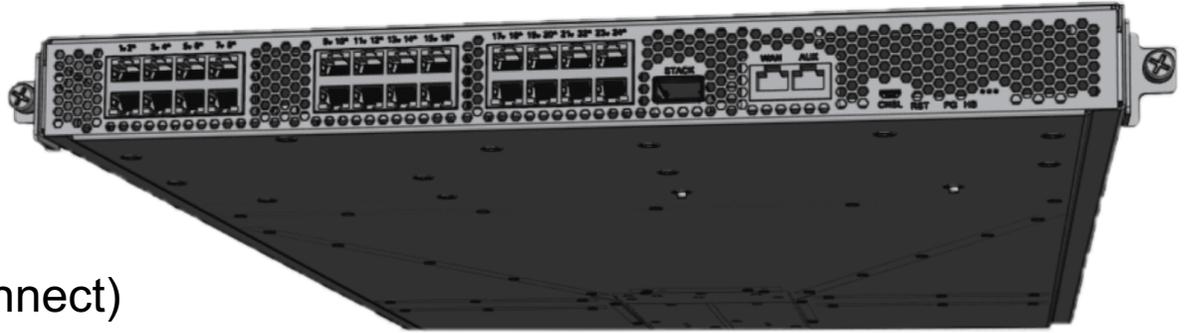
### I/O capabilities

- PCIe Live Error Recovery (LER); PCIe root port containment and card error recovery
- PCIe "Stop and Scream"; PCIe root port corrupt data containment
- PCIe end-to-end CRC checking [coming soon]
- PCIe corrupt data containment (data poisoning)
- PCIe link CRC error check and retry
- PCIe link retraining and recovery

## HPE Superdome Flex: key areas of RAS superiority

<u>RAS</u>	<u>Superdome Flex</u>	<u>Standard x86</u>
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 +1)	✓	x
Enhanced fabric resiliency (Adaptive routing)	✓	x
Advanced PCIe error recovery (LER)	✓	x
Hard Partitions (nPars)	✓	x

# Rack Management Controller



## The RMC in HPE Superdome Flex provides:

- Redfish API and CLI (ability to set, show, power, connect)
- Access to CLI (SET, SHOW, POWER, CONNECT, HELP)
- HPE nPar creation, modification and removal
- System power control and sequencing
- Environmental control and monitoring
- Specific commands to monitor or change system functions
- System status/health monitoring and status reporting
- System boot control and maintenance
- Fault and error message information

**Note:** A new server management capability in HPE Superdome Flex makes the physical RMC optional in 4s or 8s configurations. BMC firmware can now provide an embedded RMC (eRMC) functionality for 1- or 2-chassis complexes



**Hewlett Packard**  
Enterprise

**Thank you**

[www.hpe.com/superdome](http://www.hpe.com/superdome)