Hewlett Packard Enterprise

HPE APOLLO 80 SYSTEMS ARCHITECTURAL INNOVATION ENABLING INNOVATION IN OIL & GAS HPC SIMULATION

Roger Rintala, Apollo 80 & Strategic Alliances Marketing, HPE
Brent Gorda, Strategic HPC Executive, Arm

September 2021

ENERGY INDUSTRY: UPSTREAM AND DOWNSTREAM

Simulation workloads throughout the business

Proprietary applications encode and exploit intellectual property

Dispersed data and dispersed talent

Reservoir & Seismic Simulation
HPC Centers, In region/country

CFD - Process Simulation

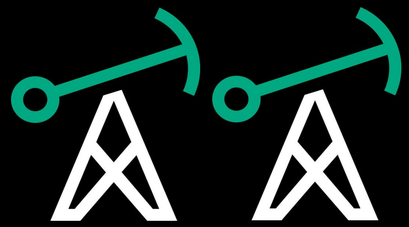

Digital Twins

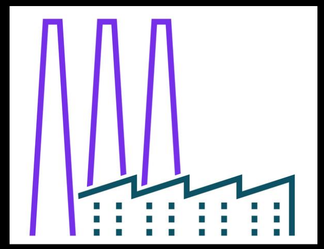

Environmental

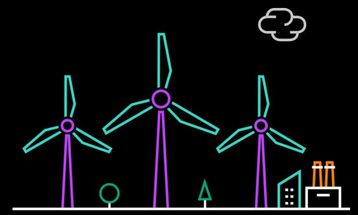

Renewables – siting and performance prediction

Carbon Capture

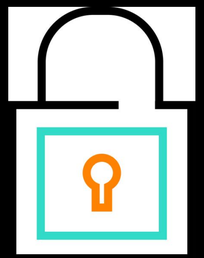

Carbon Storage

CRITICAL HPC CHALLENGES MUST BE ADDRESSED TO SUCCEED

Memory bottlenecks, limited bandwidth, and/or underutilized memory

Low throughput and inability to scale

Complex coding and time-consuming porting of applications

Inefficient power consumption and high price/performance ratio

HPE APOLLO 80 ARM-BASED SYSTEM

Apollo 80 Server

Fujitsu A64FX

- ✓ 4 times the memory bandwidth of typical x86 systems (1TB/s vs. 256MB/s)
- ✓ Delivers over 3.1 double precision teraflops
- ✓ 32 gigabits of direct-attached HBM
- ✓ Capable of >16 GFlops/watt – more power from an air cooled system

- Next-generation Arm solution continues the HPE legacy of providing leading edge, Arm-based, high performance computing.
- Built on Cray and Fujitsu's strong history with vector processing and supercomputing.
- Designed specifically for HPC, powered by Fujitsu A64FX processor for supporting new Arm scalable vector extensions.
- World's first combination of high bandwidth memory (HBM) and SVE 512-bit wide SIMD.
- HPE Apollo 80 is a compact cluster-ready solution with 8 single processor servers in a standard 19" 2U chassis – this enables over 8,000 cores and over 500 teraflops per standard 42U rack.

ACCESSIBLE SUPERCOMPUTING – PURPOSE BUILT FOR HPC

A rebalanced system architecture and very fast CPU deliver breakthrough performance for memory bandwidth bound, floating point intensive workloads seen in scientific and engineering simulation.

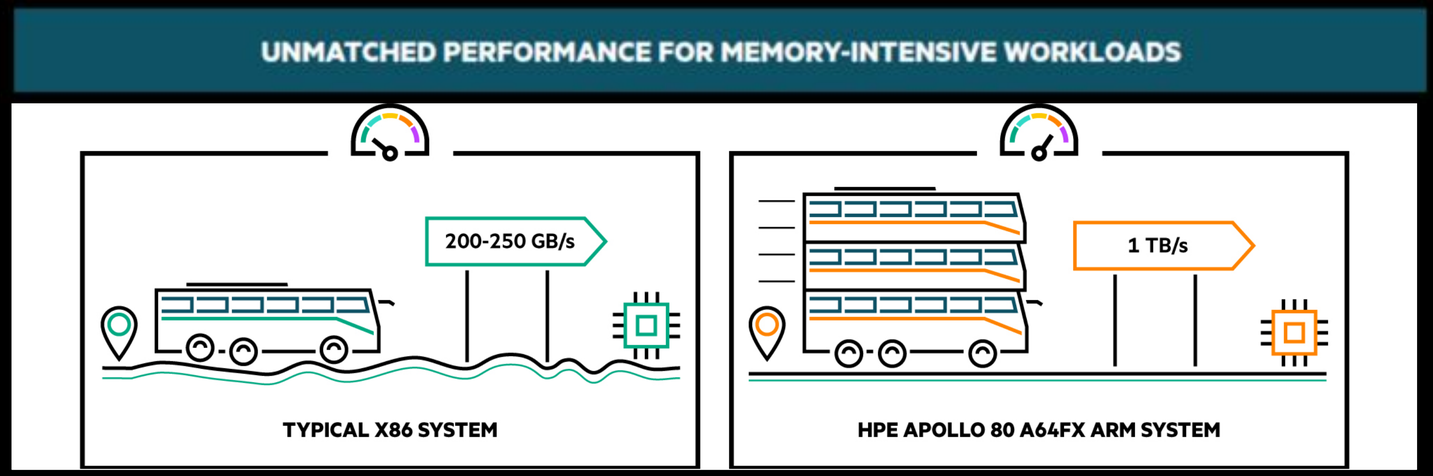

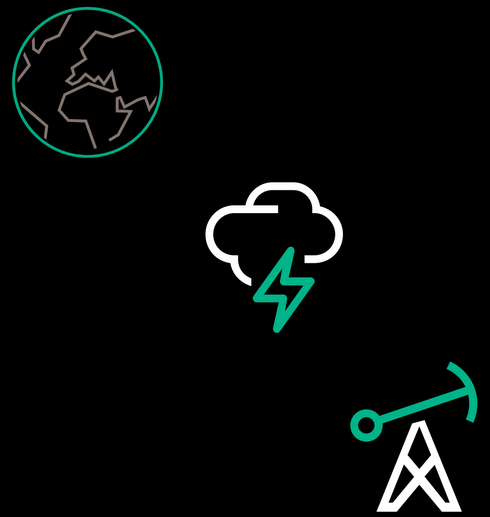

Removes bottlenecks and rebalances memory

Provides a rapid path for deploying and tuning apps

Utilizes tightly integrated processor for exceptional performance

Incorporates evolved power efficiency for best price/performance ratio

arm SystemReady

Vision

Software Can Just Work on Arm-based Devices
Generic off-the-shelf OSes can install and run with no modifications

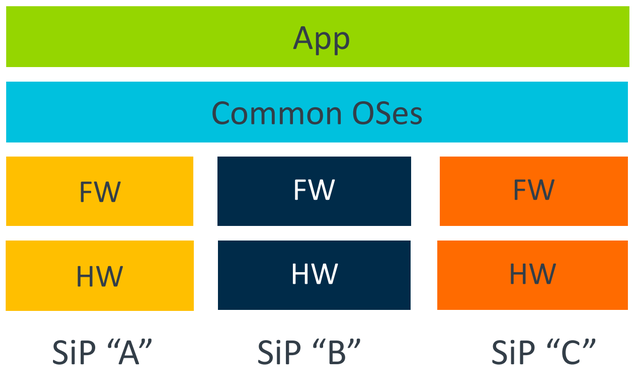

Specifications

BSA: baseline minimal set of CPU and System architecture necessary for an OS to boot and run across all market segments
SBSA: server market segment specific hardware requirements
BBR: baseline minimal set of Firmware interface requirements with SBBR, EBBR, LBBR Recipes targeting different OSes
BBSR: security boot and update interface requirements

ACS (Architecture Compliance Suite)

A test suite checking the device’s spec compliance
Server and Workstation
- LS: BSA, SBSA and LBBR
- SR: BSA, SBSA and SBBR

Edge and IoT
- ES: BSA and SBBR
- IR: BSA and EBBR

Security Extension: BBSR

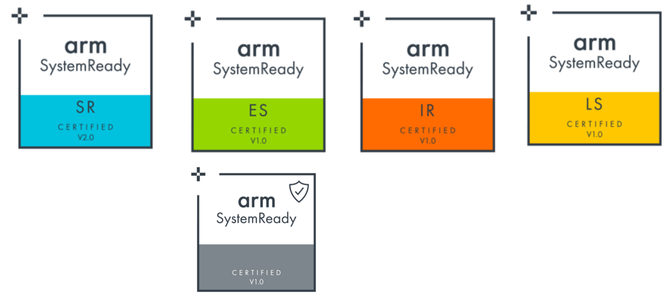

SVE: Future-Proof, Accessible, Highly Efficient

Superior auto-vectorization brings vector performance to complex codes

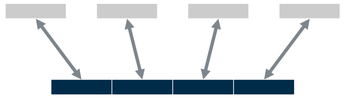

Gather-load and scatter-store

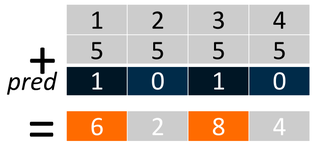

Per-lane predication

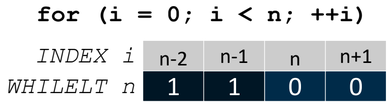

Predicate-driven loop control and management

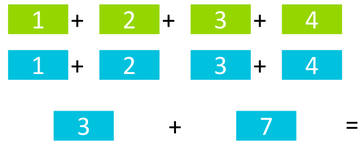

Extended floating point horizontal reductions

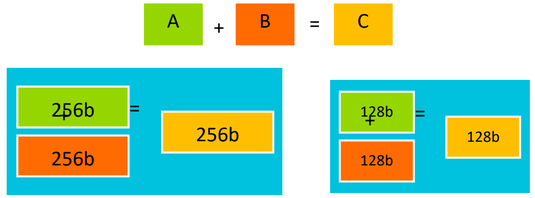

Vector-length agnostic software

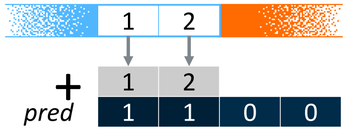

Vector partitioning and software-managed speculation

DELIVER TIMELY AND ACTIONABLE RESULTS FROM SIMULATION

A rebalanced system architecture and very fast CPU deliver breakthrough performance for memory bandwidth bound, floating point intensive workloads seen in scientific and engineering simulation.

- Our simulation workloads are supercomputer class HPC
- Need to bring top end capabilities of exascale and pre-exascale compute systems
  - Can't shrink the problem resolution, 64-bit accuracy, high fidelity required
    - –Verification and validation
  - Capable of running workloads like machine learning (BFLOAT 16) for HPC+AI workloads
  - Vectorization returns to the forefront – Arm Scalable Vector Extensions (SVE) make it easy

Pack cache lines and use the bandwidth to keep the CPU fed	Easy to work with when porting, developing, tuning application code	Accelerated performance with CPU flexibility and programming models.	Run simulations and process the data with the same system

OIL & GAS APPLICATION PERFORMANCE STUDIES ON A64FX