REPORT REPRINT

# Cray comes into the HPE fold with first combined product lines

**AUGUST 20 2020**

**By John Abbott**

The combination of Cray's Shasta technologies with HPE's Apollo clusters could broaden the high-end HPC business into hyperscalers and enterprises. HPE sees more data-centric, data-intensive workloads, with big-data analytics and AI folding into HPC simulation. The integration of Cray also brings HPE squarely into the race to exascale computing.



## Introduction

HPE is now sharing its plans for high-performance computing (HPC), post last year's acquisition of Cray, which was completed in September 2019. It's introducing the first Cray-branded line under the HPE umbrella, in two versions: the high-density, direct-liquid-cooled HPE Cray EX supercomputer, and the air-cooled HPE Cray supercomputer (same name but without the 'EX'), which comes in standard 19" racks. The new releases mark the gradual introduction of Cray's next-generation Shasta architecture, intended to be much more flexible and extensible than Cray's previous flagship XC systems, supporting a wide range of processors (AMD, Intel, NVIDIA), a new interconnect (HPE Slingshot) optimized for data-intensive workloads, and with a whole new software stack with dynamic, cloud-like capabilities and hybrid system operation. HPE sees an expanding role for HPC that encompasses more data-centric, data-intensive workloads, with big-data analytics and AI folding into HPC simulation, as well as wider cloud hyperscale and enterprise use. The integration of Cray also brings HPE squarely into the race toward exascale computing.

## 451 TAKE

Cray struggled for years to broaden its business beyond the high-end, high-profile 'grand challenge' contracts it is best known for, which can be hard to monetize, and which can result in very lumpy quarter-on-quarter financial results when projects slip. It never really succeeded in getting a good hold on more commercial and enterprise sales, but its development efforts were heading in that direction well before the HPE acquisition. The problem was perhaps that to fully commercialize something like the Shasta interconnect, a lot of other pieces are required. HPE has a lot of those pieces. The combination of Cray's Shasta technologies with HPE's Apollo clustering business and the SuperDome Flex, where large in-memory computing is a requirement, immediately broadens the prospects for Cray. And there's likely to be some interesting innovations ahead that will prove relevant to users beyond pure HPC – especially the hyperscalers. Those include the flexible runtime software environments and libraries that can dynamically adapt for different workloads, and the combination of the HPE Slingshot interconnect and its congestion control mechanisms, both with current Ethernet networks, and the forthcoming silicon photonics works in progress within HPE Labs.

## Context

HPE agreed to acquire Cray in May 2019 for $35 per share ($1.4bn in cash), completing the deal on September 25. Cray's president and CEO, Peter Ungaro, was appointed head of HPC and AI within the Hybrid IT business group at HPE, which also includes hybrid IT infrastructure, software-defined and cloud. Ungaro takes control not only of the Cray assets and the Apollo dense clustering line, but also mission-critical systems, SuperDome Flex and Moonshot edge servers. On HPE's Q2 earnings call in May, CEO Antonio Neri said that the Cray integration remains on track to deliver on its fiscal 2020 revenue targets, with 'triple-digit run rate synergies' anticipated by fiscal year 2021.

## Technology

Cray's Shasta architecture, partway through its market introduction, is based on two physical infrastructures: Apollo (from HPE) and Olympus (from Cray). Apollo is air-cooled, uses 19" standard racks and can take advantage of a wide range of available compute and storage options (including from third parties). Olympus (now HPE Cray EX) is warm-water-cooled and built for high density and scalability. To do this effectively, it supports a smaller number of custom blades than Apollo, and they are optimized for key computational tasks. The Apollo-based and Cray EX-based infrastructure share the same interconnect for expansion and the same software environment. Cray EX can scale up to 64 compute blades, 128 CPUs and 512 GPUs per cabinet. The blades will be upgradable to support future CPUs, GPUs and interconnects, including optical. Each cabinet can support up to 64 Slingshot switches, with scaling from one to 'hundreds' of cabinets on offer. High levels of density can be achieved through the direct-liquid cooling (including processors, memory, optical transceivers, etc.), with single-cabinet power capability currently topping 300kW, but expected to be increased to 400kW over time. (The comparable limit for air-cooled cabinets is roughly 35kW.) This headroom will be important as processors reach beyond 1kW power consumption.

Slingshot is Cray's eighth generation of interconnect, but it is Ethernet-based rather than proprietary – in other words, standard Ethernet connectivity is supported at the edges, but with a newly designed HPC fabric inside. The switch has 64 ports at 200Gbps and has a network diameter of three hops for up to 100,000 nodes, enabling really large systems to be built. It adopts adaptive routing and also quality of service, the latter a surprising first for Cray, but a necessary feature to better support hybrid workloads. To improve workload performance isolation even further, Cray has added congestion control. It adds up to low latency, especially tail latency and uniform latency, which is important for HPC synchronization but also for hyperscalers managing service-level agreements that rely on sequencing multiple thousands of remote procedure calls.

With the Shasta architecture, Cray has pulled storage directly into the high-speed network rather than use an external storage area network. ClusterStor E1000 can mix and match flash nodes, hybrid or all hard-drive-based nodes into the Slingshot fabric. It's cheaper, less complex, has lower latency, and is beneficial for small IOPs-type operations. The final new element of Shasta is the Linux-based software stack, which includes a systems management platform using open and documented RESTful APIs, orchestration and containerized services (alongside HPC batch job management), a monitoring framework and, on the developer side, a set of parallel performance libraries, runtimes and tools, including an AI analytics suite.

## Product

The air-cooled HPE Cray Supercomputer is the first available system family offering to combine HPE and Cray Shasta technologies. It brings together the Apollo 2000 Gen10 Plus System (a four-server, 2U high-density-optimized system based on Intel Xeon CPUs), the Apollo 6500 Gen10 (an eight GPU server), Slingshot, and the HPE Cray software stack. HPE Cray Supercomputer EX, currently based around four AMD EPYC 7002 dual-CPU nodes, is derived from the Cray side of the house. In future, other CPUs and blades will be supported and other HPE products and services integrated, such as GreenLake flexible capacity and HPC platform as a service (including the recently launched Ezmeral platform for running containerized, AI, machine learning and analytics workloads). Partners will offer managed HPC as a service, and key technology components and runtimes, including ClusterStor storage, will be available in the Microsoft Azure cloud.

Rather to one side of all of this is the HPE Apollo 80, a new compact cluster with eight single-processor servers in a standard 2U rack chassis, supported by the HPE Performance Cluster Manager. It utilizes Fujitsu's Arm-based A64FX CPU and is a continuation of a system line that's been adopted by Los Alamos National Labs, Oak Ridge Labs and the University of Bristol. Apollo 80 systems will be able to run the HPE Cray Programming Environment. As for other HPC systems previously in the portfolio: Cray's CS commodity clusters (derived from its Appro acquisition) will be superseded by Apollo; XC systems are nearing their end-of-life; and SuperDome Flex (using the UV technology from HPE's previous HPC acquisition Silicon Graphics) will continue in its role as a platform for high-capacity in-memory computing.

## Strategy

The original justification for buying Cray was the stimulus that data-driven applications such as AI, machine learning and analytics have provided to the high-performance computing sector – of which Cray has long been the dominant player at the very high end. On that note, the race toward exascale computing was also mentioned as a motivator, with Cray awarded three of the earliest exascale projects at the US Department of Energy: 'Aurora' (Argonne National Labs, in partnership with Intel); 'Frontier' (Oak Ridge National Labs, with AMD); and most recently 'El Capitan' (Lawrence Livermore Labs, also with AMD). But beyond those high-profile and extreme technology examples, HPE is looking at a wider opportunity – the extension and integration of AI and analytics into classical HPC shops, and (from the other direction) the increasing demand for HPC-like capabilities from the hyperscalers, again influenced by AI and analytics. A lot of the technical decisions outlined above were made with this transition in mind, particularly the Ethernet compatibility and congestion control built into the HPE Slingshot interface. The common software stack, aimed at supporting hybrid workloads, will also play an important role here.

## Competition

Cray was the last significant independent supercomputing company until HPE acquired it last year. Combined, HPE and Cray have 76 systems on the latest (June 2020) Top 500 supercomputing list, putting them third after Lenovo (180 systems) and China's Sugon (68 systems). When measured by performance, the HPE-Cray combination would rank second, not far behind Lenovo. ATOS/Bull, Fujitsu and NEC are among the few others (aside from Cray) working on fundamental HPC architectures that might also be applied to hyperscale, though we are seeing a rapid ramp-up in homegrown innovation from China and the rest of Asia, partly driven by emerging geo-political pressures. Arm-based supercomputers are in the ascendant, and AMD has made significant strides (including on CPU-GPU integration), but Intel is still by far the dominant CPU architecture in HPC. IBM continues to develop the Power architecture and is gradually opening up the licensing. NVIDIA's ownership of Mellanox gives it a strong position in both GPUs and high-performance networking, signifying its increasing engagement in the HPC sector.

## SWOT Analysis

### STRENGTHS

Cray's fundamental rethinking of HPC architectures with Shasta was well underway before the HPE acquisition. HPE has the deeper pockets required to get it into the market.

### WEAKNESSES

HPC users are probably the most demanding in the industry and very set in their ways. If HPE-Cray is to make the technology more general-purpose, it mustn't lose its core customer base in the process.

### OPPORTUNITIES

HPE brings to Cray the wider commercial portfolio and enterprise sales channels, as well as related services like GreenLake, that could finally get it noticed by the cloud hyperscalers.

### THREATS

The hyperscale industry is moving toward HPC support from the other direction. Hyperscalers have the biggest budgets in the industry and have established huge datacenter networks they can bring to bear.