

Pavilion All Memory Array

Highlights

Extremely Powerful

- Shared Storage with the latency of direct-attached SSDs
- 25X More IOPS and 10X More Bandwidth than All-Flash Arrays
- Low Consistent Latency
- Up to One Half PB in 4U
- Symmetric Active-Active Controller Support

Enterprise Ready

- Frictionless Deployment
- Data Resiliency & High Availability
- Guaranteed Performance QOS
- Space-Efficient, Instant Snapshots and Clones
- Enterprise Management & Monitoring
- Thin Provisioning
- Zero Host Footprint
- Non-Disruptive Upgrades
- Standard Ethernet

Unprecedented Value

- Pay As You Grow Scalability and Modularity
- Expand for Capacity or Performance, Independently
- Increase Storage Utilization up to 10X or more

Delivering Maximum Capacity and Performance Density

The Pavilion Memory Array delivers 25X the performance and 10X better latency than typical networked All-Flash-Arrays, enabling the next generation of Real Time Analytics, OLAP and OLTP applications, all in a 4U Package.

Big data analytics applications can now analyze much larger data sets, and therefore deliver more accurate answers. Decisions can be made in real-time at the speed of digital business by analyzing more data quickly.

Ultra-low latency (10s of μ s) allows customers to make provisioning decisions at application run time, as opposed to procurement time, by replacing direct-attached SSDs in database clusters, leading to dramatically increased storage utilization and lower costs in these environments.

A Shared Memory Array Designed Like a Network Switch


Storage system design has not evolved much in the past two decades, even though storage media has become much faster and network speeds have increased dramatically. In order to deliver new levels of performance, storage system design needs to be re-thought in fundamental ways. Pavilion's patented architecture borrows many design ideas from the network switching world. One example is employing modular "IO Line Cards" that provide linear expandability in the areas of network bandwidth and storage controller processing power.

Besides providing maximum performance and scalability, this design also allows customers to pay-as-you-grow by starting with a small number of IO Line Cards initially, and adding more as bandwidth and performance requirements increase over time. Storage capacity can also be expanded in a modular fashion within the chassis, and grow to up to 72 NVMe SSDs in a single 4U system.

Pavilion Clustered Memory Operating System (PCMOS)

Pavilion's Clustered Memory Operating System (PCMOS) is a platform that provides the Pavilion Memory Array's resiliency, manageability, performance and flexibility. It leverages 100s of cores in the system to simultaneously process storage workloads at maximum performance and with the lowest latency. PCMOS allows Pavilion's micro-storage architecture to seamlessly operate like a single system.

PCMOS ensures that there is no single point of failure in the system by orchestrating failover between components. PCMOS can migrate workloads between IO Line Cards to ensure quality of service is maintained, as well as guarantee availability in the event of a port, controller, or line card outage.



Enterprise Features

The Pavilion All Memory Array is the first shared storage system to deliver both record-breaking performance and latency, as well as enterprise management features that today's customers expect in their networked storage arrays, including:

- Symmetric Active-Active Controller support for HA and load-balancing. Parallel access to data volumes through up to 5 storage controllers, and 10 network ports.
- Instant space-efficient snapshots and clones, allowing the effective capacity of the array to grow significantly, and data to be easily backed up and/or leveraged for multiple uses.
- Deliver guaranteed performance for your applications by specifying per-volume QOS policy
- Data recovery in the event of drive failures
- Thin Provisioning, allowing for increased utilization and flexibility

Frictionless Deployment

Pavilion is the only low-latency storage array that doesn't require users to install either custom hardware or software on application servers, or deploy non-standard networks or protocols in order to use the system. The only host components required include:

- Standard RDMA-capable Ethernet
- Standard In-Box NVMe-Over-Fabrics host driver

Using these standard components, logical NVMe volumes can be provisioned on any host.

	RF108-18			RF116-36			RF140-72		
Number of NVMe Drives	18			36			72		
Capacity (TB)	14	28	57	28	57	115	115	230	460
IO Line Cards	2			4			10		
Storage Controllers	4			8			20		
40 Gbe Network Ports	8			16			40		
System Bandwidth (GB/s)	20			60			120		
4K Read IOPS	4,000,000			9,000,000			22,000,000		
Mixed 4K IOPS (70:30)	3,000,000			8,000,000			19,000,000		
Maximum Power Consumption	1100 W			1800 W			3200 W		
Cooling Requirements (BTU/hr)	10,000			10,000			10,000		
System Height (Rack Units)	4								
Expansion Options	Single IO Line Cards, 9-Drive Capacity Packs								
Operating Temperature Range	10 to 35C								
Non-Operating Temperature Range	-40 to 70C								
Humidity Range	8 to 90% (non-condensing)								
Non-Operating Humidity Range	5 to 95% (non-condensing)								

Agency Approvals

Safety	IEC/EN 60950, CB Certificat, UL60950-1 CAN/USA-C22.2 No. 60950-1, CE Mark
Emissions	EN55022/CISPR 22, FCC Part 15 Class A, ICES--003 Issue 5 Class A, VCCI Class A, AS/NZS CISPR22:2009 Class A BSMI CNS 13438 Class A
Immunity	EN55024, CISPR2