



NVME-OF STORAGE FOR GREENPLUM

Features

Greenplum Benefits

- Massively parallel NVMe storage array ideally aligned to Greenplum
- Consistent Read/Write bandwidth up to 120GB/sec. with end-to-end latencies < 40µs
- >14TB/hour parallel ELT
- Simple deployment of Greenplum on Kubernetes
- Multiple storage zones for Master, Segment, Capacity and ELT
- 10/100GBE standard interconnect (up to 80 ports)
- Built-in RAID-6

Pavilion Benefits

- Up to 1PB in 4U
- Snapshots easily integrate with backup tools eliminating network traffic
- Thin provisioning for unplanned non-disruptive Master and Segment expansion
- 20 Active-Active controllers with 4 100GBE ports each
- **OPENCHOICE** Storage of NVMe drives for bandwidth, endurance or extremely low latency inside or across chassis

Deploy Disaggregated Storage to Greenplum environments.

With improved transaction processing capability and support for streaming ingest, Greenplum can address workloads across a spectrum of analytic and operational contexts, from traditional business intelligence to deep learning. Greenplum is designed to run anywhere—on-premises, in public and private clouds, and in modern containerized environments—for easier installation, operation, and upgrades.

Greenplum features a shared-nothing architecture that automates parallel processing of data and queries and petabyte-scale data ingestion. However, today's architectures present challenges for infrastructure teams when supporting these large scale data warehouses.

New Storage Challenges

Greenplum recommends Direct-Attached Storage (DAS) to deploy distributed resources in a scale-out fashion. However, architecting for performance and capacity can be easily compromised. With the advent of NVMe-oF, it is now possible to disaggregate Greenplum compute from storage and achieve superior performance with low latency and advanced data management features. Pavilion Data's NVMe-oF Storage Platform is ideally suited for massively parallel Greenplum environments.

- With >1PB of capacity in 4U, Up to 20 Active-Active Controllers and 80 100GBE interconnect ports, the system can be deployed as an all-in-one storage array for Master, Segment and ELT nodes, or as a stand-alone array assigned to a specific number of Segment nodes for the Greenplum data warehouse.
- The system supports an OPENCHOICE of NVMe SSDs. Choose 3D XPoint(TM) for maximum performance with Masters, leverage high endurance NVMe for ELT, deploy high-bandwidth drives for Segment nodes and use the latest 30TB NVMe drives for capacity, snapshots/clones and backups.
- Simple and proven deployment model delivered through Dell Extended Technologies.

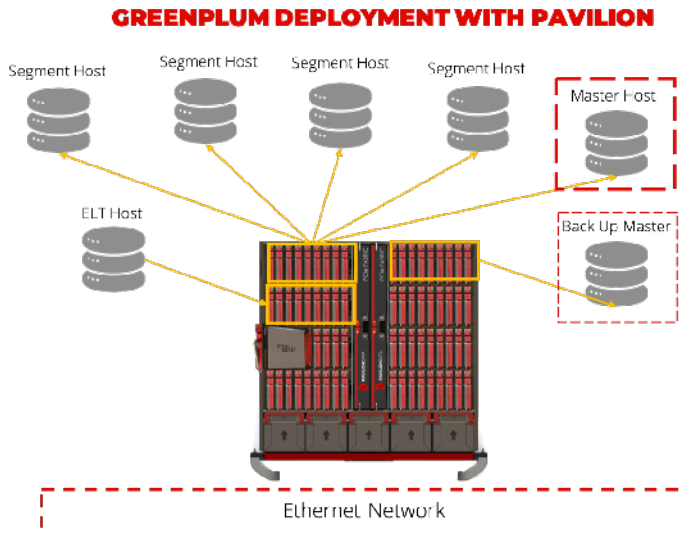
Pavilion's NVMe-oF Storage Array

We provide an ideal companion to Greenplum's Massively Parallel Compute (MPC) architecture that delivers up to 120GBs of bandwidth and up to 1PB of capacity in 4RU. The Array requires no custom software to be installed on application servers and includes important data management and availability features, including thin provisioning, instant zero-space snapshots and clones, and no single point of failure.

Deliver Disaggregated NVMe-oF Storage for Greenplum Deployments:

Until now, Greenplum clusters were deployed on DAS SSDs because of the performance and fault isolation requirements. The absolute high bandwidth requirement, in particular, drove the need for DAS.

However, with new high-speed RDMA-capable networking and efficient storage protocols like NVMe-oF, it is now possible to get the same performance advantages with shared storage.



Pavilion's Platform splits into pre-configured zones. A zone could be the whole platform or just a subset of SSDs like the example shown on the left.

- Master Host - High performance NVMe access for data loading, connection handling, and query planning. Seamlessly configure volumes if extra space is required for landing load files and backup files, especially in production environments.
- ELT Host - We enable customers to maximize parallelism and load bandwidth for their external tables features. Production systems can be deployed by designating ELT zones within the platform for data loading purposes.
- Segment Host - Provide high performance to query processing across segments in the Greenplum Database system and greatly reduce network traffic used for segment instances.

With **OPENCHOICE**, customers deploy different classes of SSDs for different zones. For example, Segment Hosts would require bandwidth for higher query processing rates, whereas back ups might require a more capacity based SSD.

Deploy up to 4X+ less flash deployed:

By leveraging thin-provisioned logical flash storage, our platform delivers the required needs of Master or Segment nodes. You can decide at deployment time how much storage to provision to any given node and are no longer are constrained by the size of the SSDs that were purchased and installed in any given node. Thin Provisioning allows the Greenplum to use the required amount of storage at any given time, regardless of how much capacity has been advertised to that specific database node. This greatly reduces the amount of raw flash storage deployed in large data warehouse environments.

Simplify data protection and reduce server overhead:

Integrate instant snapshots and clones with Greenplum parallel backup utility to allow an entire database to be backed up or copied for recovery purposes, on the fly, without any performance impact. Our platform provides no single point of failure, ensuring maximum application uptime and data availability. This significantly reduces East-West network traffic and shrinks backup windows. Clone the snapshot inside the same Pavilion Platform and mount the image for new Test/Dev activities.

Increased Compute Density per Rack by deploying Disk-less server nodes:

By provisioning high-speed logical flash storage volumes to each server in a rack, you no longer need to purchase servers that accommodate SSDs. This provides the ability to increase the compute density of a rack by leveraging 1U servers instead of 2U servers with front-loading drive bays. Pavilion also requires no custom software to be installed on database nodes, allowing Greenplum to take full advantage of the application host processing resources as well as simplifying deployment complexity.

The result? The power, simplicity, and density offered by the Pavilion NVMe-oF provides the first flexible service for scale-out Greenplum deployments, increasing agility and flexibility and lowering TCO in the process.