



A POWERFUL COMPOSABLE APPROACH FOR YOUR BIG DATA ANALYTICS





Data insights are at the core of any digital transformation, as they are the key for organizations to increase competitiveness, operational efficiency, and improved customer services.

Traditional converged or even hyperconverged infrastructure with fixed ratios of storage, networking, and compute resources do not provide enough agility to address these challenges to meet the needs of the Big Data applications.

Next-generation Big Data analytics applications that are critical to derive insights are very dynamic and range from streaming analytics to interactive analytics, to artificial intelligence with a variety of requirements. Following are the current infrastructure challenges that the organizations are facing today:

- Complexity
 - Accommodating heterogeneous deployment models such as bare-metal, VMs, containers and so on, in a single enclosure
 - Upgrading enterprise-grade tested firmware on all the nodes simultaneously
- Lack of flexibility in provisioning infrastructure
 - Not able to independently select and scale resources (storage, compute, and network) based on workload demands.
 - Lack of elasticity to release the resources when they are no longer needed by the applications.
- Longer time to market
 - Taking too long to get the system up and running for production

The combination of HPE Elastic Platform for Analytics (EPA) architecture with HPE Synergy composable infrastructure helps organizations overcome these challenges and rapidly provisions Big Data analytics infrastructure on-premises with security and scalability.

KEY VALUES OF OUR SOLUTION

HPE EPA with composability is designed to be a modular infrastructure foundation that can be scaled to support different workloads and requirements. Why is this elastic design model with composability so important? Building a platform based on an elastic architecture design provides the following benefits:

- **Breakthrough economics:** Significantly better density, cost, and power through workload-optimized components; accelerate time from information to analysis to action/decision
- **Elasticity and flexibility:** Scale compute and storage independently; rapidly provision compute without affecting storage
- **Efficiency with performance:** Fast deployment and firmware update; multiple next-generation Big Data workloads can be hosted on a single platform while still maintaining performance





Three use cases successfully deployed at the customer sites

- Test/Dev environment
- Data-intensive analytics
- Streaming analytics

USE CASES TESTED AND DEPLOYED SUCCESSFULLY AT THE CUSTOMER SITES

Scenario I: HPE EPA with single HPE Synergy frame

Many enterprises need a test/dev environment to solve complex next-generation Big Data analytics problems before deploying the solution to the production environment. These test/dev environments need to be self-contained and flexible with both hardware resources and software tools, and with faster deployment and redeployment options. HPE EPA with composable infrastructure provides the required hardware resources with high processing power within a single HPE Synergy frame through template-driven workload composition.

Scenario II: HPE EPA with multiple HPE Synergy frames

This configuration is best suited for use cases requiring faster data access (SSDs for storage tier). In this data-driven era, data is generated from everywhere and from multiple sources in multiple formats and structures. It's difficult to get actionable insights from data due to its volume, velocity, and variety. Customers are looking for analysis from complex data set with low latency on partial data (terabyte data) to gain quick analysis to get results faster. Data-intensive analytics like interactive/iterative analytics and descriptive analytics using Hadoop, Impala, Kudu, Hive, and so on can help here. But they require huge compute power and memory capacity to run queries faster with low latency on terabytes of data and deliver the results in real-time (in this context, real-time is defined as a user can wait for results coming out) or to process the same kind of queries several times.

HPE EPA with composable infrastructure is designed to support such cases. It is a modular infrastructure to provide high-density compute and high-capacity memory workloads with optimized performance over HPE Synergy.

HPE Synergy can easily repurpose the infrastructure to run different workloads via HPE OneView configuration. It supports software-defined networking, reduces the complexity (network sprawl) and cost by connecting blades internally within the frame with huge cable reduction.



Solution overview

Modular architecture for multi tenant on-premises deployment

- The solution provides building blocks based on density, capacity, and performance for workload optimization
- HPE Synergy composable infrastructure scales storage and compute independently to reduce total cost of ownership and data center footprint while optimizing performance Big Data workloads

Improved efficiency and time-to-production

- The solution enables configuring and deploying Big Data analytics workloads rapidly
- HPE Synergy Composer and HPE OneView enables high-level, automated operations for template-driven workload composition, reducing complexity and manual tasks, and resources can be dynamically allocated to meet the needs of any Big Data applications.

Scenario III: HPE EPA with HPE Synergy and HPE Apollo 4200 Storage Server

This configuration is best suited when the customer opting for streaming analytics use case requires higher-density HDD drives. Streaming analytics helps organizations make quicker insights on the real- or near-time data and helps to absorb huge chunk of data into the data lake. Since streaming analytics is real time/near real time, the frameworks and analytics applications need to be designed as part of a data pipeline (for example, Kafka, Spark, NoSQL DBs and so on). For these real-time/near real-time applications, the infrastructure has to be flexible, scalable, easy to deploy, and optimized to deliver the required performance.

HPE EPA with HPE Synergy's composable infrastructure and HPE Apollo 4200 storage provides an ideal modular infrastructure to address all these requirements. It helps organizations to easily analyze the stream data and make decision at real-time/near real-time. A scalable, multitenant platform that enables independent scaling of compute and storage. These infrastructure building blocks are optimized for density and capable of running disparate workloads. HPE Synergy also allows to repurpose the hardware resources based on the workload needs and dynamically provision the system to help maximize the performance.

THE PLATFORM YOU CAN COUNT ON

This modular infrastructure platform consists of HPE EPA with HPE Synergy supporting a wide range of Big Data and analytics applications. It benefits you with faster deployment, flexibility, elasticity, scalability, and optimized performance that's required for your next-generation Big Data applications.

LEARN MORE AT

hpe.com/info/synergy

hpe.com/info/ra

hpe.com/info/oneview

hpe.com/info/apollo

Make the right purchase decision.
Contact our presales specialists.



Chat



Email



Call



Get updates

**Hewlett Packard
Enterprise**

© Copyright 2020 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

a50001228ENW, May 2020