OnX Hadoop Appliance Services

OVERVIEW

OnX Hadoop Appliances are pre-engineered Hadoop environments. OnX enables customers to make choices about the hardware, workload form factor, and cluster size through a customized configuration tool. Additionally, the customer is optionally allowed to add services for installation and configuration, administration, training, and professional services for development.

OnX Hadoop Appliances are vetted and proven pre-engineered solutions for customers that want to run Hortonworks on HPE infrastructure.

Background

When Hadoop Distribution vendors (Hortonworks) sell initial Hadoop distros to customers, the customer is largely left on their own to make a complex hardware purchase for the Hadoop infrastructure. When hardware vendors (HPE) sell Hadoop to customers, they tend to recommend infrastructure specifications that could be more complex and expensive than what the customer actually requires. OnX Hadoop Appliances take the risk out of initial Hadoop projects by providing vetted pre-engineered systems. This allows our customers to focus on delivering their application rather than the environment for it.

OnX offers expedited delivery on our Hadoop appliances. The lead time for Hadoop Distro vendors and OEMs to deliver the necessary infrastructure for a Hadoop cluster typically takes between 12 to 16 weeks. This delay in infrastructure delivery can cause significant project risks and may result in critical project resources being diverted to other running projects. This delay will be prolonged should incorrect components be ordered and need to be exchanged.

The advantages of OnX Hadoop Appliances are:

- Configured for specific workloads (Compute, Storage, General Purpose)
- Vetted system architecture with Hadoop Distro providers, OEM vendors, as well as OnX customers
- Expedited delivery
- Fixed price add-on service options to help customers be successful in project delivery

OnX Hadoop Appliances are delivered preassembled. OnX helps connect the appliance to your network and we provide 2-3 weeks OnX Quick-Start Hadoop administration services. All included in the Hadoop Appliance.

Use Case

Hadoop is a complicated environment to setup. There are many reference architectures from the multiple hardware and Hadoop vendors which can cause confusion.

When Hadoop is purchased, the right infrastructure needs to be deployed to run applications at optimal performance. Infrastructure choices that need to be made include selecting the right disk, cache, memory, CPU, network cards, top of rack switches, rack wiring, and much more. These must be selected with knowledge of correct memory to CPU to disk ratios to ensure the best performance for the application being delivered. Next is infrastructure configuration from the disk pool configuration to redundant network rack wiring. Afterward, the Hadoop distro must be configured including services to run, queuing, compression, replication, job scheduling, redundancy & failover, security, and much more.

Hit the easy button and **let OnX do it for you!** Leverage OnX's significant experience in setup and management of large 1000-node environments at fortune 500 service provide and financial service clients.

Benefits

The OnX Hadoop Appliance adds value to our customers, Distribution partners, and OEM vendors by streamlining the delivery of pre-configured, well-engineered Hadoop environments:

- Accelerate deployments and time to market
- Eliminate the complexities of hardware selection
- Reduce project risk by using vetted /proven configurations
- Don't lose project momentum by ensuring faster delivery than other channels

ONX APPROACH

Contact OnX

Simply send an email to *Hadoop@OnX.com*, include your contact information including Name, Phone and Title. An OnX configuration specialist will contact you within 1 business day.

Select a Workload

Choose the Hadoop infrastructure that best meets your workload requirements. OnX Hadoop Appliances are engineered for running Hortonworks on HPE infrastructure to meet the demands of specific deployment workloads (form factors). Select from:

- **Process Intensive** – for CPU-bound data processing workloads such as Spark. Other examples of CPU-Bound workloads include:
  - Clustering/Classification
  - Complex text mining
  - Natural-language processing
  - Feature extraction
- **Efficient Storage** – for IO-Bound workloads requiring intense reading and writing from disk such as Data Warehouse Offload where Hadoop is used to support a tiered storage strategy.
  - Sorting
  - Indexing
  - Grouping
  - Data importing and exporting
  - Data movement and transformation
- **General Purpose** – for environments that must support mixed workloads.

Select a Cluster Size

Select one of our pre-configured appliance sizes of 8, 12 or 20 servers. Non-Standard cluster sizes can be ordered but are considered custom builds.

The ONX Hadoop Appliance rack e designed to support up to 32 servers and associated network infrastructure.

OnX is able to provide expedited delivery of our pre-engineered Hadoop Appliances. Customization is available but could impact delivery schedule.

Select a Hadoop Distribution

Select either Hortonworks Enterprise or Hortonworks Enterprise Plus. Most companies use Enterprise Plus for Storm, Spark, Kafka, and Ranger support.

OnX Hadoop Appliance

OnX has created 9 configurations/SKUs that are based on usage form factor and Hadoop Appliance size. The following table provides a high-level overview of the pre-engineered configurations (Memory, Processor, and Storage); as well as, an estimate of the associated cluster storage capacity.

	Node Types	Process Intensive	General Purpose	Efficient Storage
8	2-Master 6-Data-Util	**Memory**: > Master: 256GB > Data-Util: 256GB **2.3GHz Processor**: 14 core,35MB Cache **Storage**: 10-1TB SAS Drives	**Memory**: > Master: 256GB > Data-Util: 128GB **2.66GHz Processor**: 10 core, 25MB Cache **Storage**: 10-2TB SAS Drives	**Memory**: > Master: 256GB > Data-Util: 128GB **2.66GHz Processor**: 10 core, 25MB Cache **Storage**: 4-6TB SAS Drives
12	3-Master 9-Data-Util			
20	3-Master 17-Data-Util			

- Order ONX racks hold up to 32 servers per rack
- Non-Standard cluster sizes are considered custom builds

Cluster Capacity Table*				
8	6-Data-Util	60 TB raw 20 TB w/ replication	120 TB raw 40 TB w/ replication	144 TB raw 48 TB w/ replication
12	9-Data-Util	90 TB raw 30 TB w/ replication	180 TB raw 60 TB w/ replication	216 TB raw 72 TB w/ replication
20	17-Data-Util	153 TB raw 51 TB w/ replication	340 TB raw 133 TB w/ replication	408 TB raw 136 TB w/ replication

* Estimated capacities

Onsite network hookup and OnX Hadoop Quick-Start services are included as part of the OnX Hadoop Appliance.

Hadoop Appliance Rack

- Standard 42U - 19inch Rack
- 2 Workload switches
- 1 Management switch
- Cabling

Hadoop Appliance Standard Server Configuration

- HPE ProLiant DL360 Gen9 servers
- 10-GiGE dual port NIC
- 128G Micro SD for OS
- Redhat OS

Hadoop Appliance Standard Services

- Rack and Stack plus Linux OS, Volume Configuration and MySQL (onsite, integration center)
- Hadoop Distro Install and Config (onsite, integration center)

Select Hadoop Appliance Options

OnX is experienced in delivery Hadoop clusters for some of the largest US fortune 500 companies. We use that that experience to drive flexibility in OnX's offerings. OnX Hadoop Appliances consist of a rack, servers, networking, wiring, Hadoop Software and Hadoop Appliance configuration. The Hadoop Appliance is everything you need to get started with your Hadoop project.

OnX provides our customers the option to add Hadoop services that they may require to make the project successful. We offer Hadoop Administration, Training, Development and Data Science Services to help our customers be successful.

Hadoop Administration Services

Hadoop Training

- Administration
- Development
- Data Science

Hadoop Development

Business Analytics and Data Science Services

OnX Hadoop Appliance Delivery

We capture your requirements in a build document created during the order process. OnX builds the Hadoop Appliance to your requirements within our integration center capabilities. We configure network, volumes and other key functions at the factory, as much as possible. Once the Appliance is built, we install the Hadoop projects per your requirements. We start the appliance prior to shipping to ensure quality. Once the appliance is delivered to your site, OnX will connect it to your network and perform final configure the Hadoop appliance to your specification per our Quick-Start service and do knowledge transfer to your staff.

Use Case

When a customer want to purchase Hadoop, they should send an email to Hadoop@OnX.com with their contact information, to engage OnX. An OnX sales representative will be in contact within 1 business day. We will take time to understand the specific requirements of the environment and will build a customized pre-engineered system to meet those needs. OnX will ensure delivery of the assembled system with the Hadoop distro installed and configured. Should extra assistance be needed, OnX can help. OnX offers several Hadoop add-on services which can be optionally added.

Custom configuration support is available on request

OnX understands that business requirements may require some customers to need a custom configuration. OnX can work with your team to build a custom Hadoop cluster architecture that is purpose built to meet your specific needs. Custom option include:

- HPE Proliant MoonShot Architecture
- HPE Apollo high performance computing
- Special HPE ProLiant Gen9 server requirements

RELATED SERVICES

- OnX HDP Administration Training
- Onx HDP Developer: Pig & Hive Training
- OnX HDP Data Science Training
- OnX Hadoop Quick Start
- OnX Hadoop Administration
- OnX Hadoop Development
- OnX Business Analytics and Data Science

OnX provides a full range of capabilities to our customers including Data Integration, Data Management, Data Warehousing, Big Data, Data Governance, Business Intelligence and Data Science.

PROJECT MANAGEMENT

OnX includes project management leadership to manage the overall project resources, create and maintain the detailed project plan, communicate status on a recurring basis and facilitate escalations as needed. This is crucial to minimize risks and ensure a timely and successful service delivery. The OnX project engagement activity consists of a four phase project lifecycle specifically tailored to deliver comprehensive, quality capabilities. OnX maintains a knowledgebase of "lessons learned" comprised of feedback from all service deliveries to help prevent unforeseen delays and other impact on the project. We also leverage world-class knowledge centers to ensure our past successes translate into a consistent process contributing to your project's success.

WHY ONX?

Organizations can procure technology products from a number of different sources, but it's the services, technical expertise and proven methodology that makes a difference toward reaching business goals. OnX is highly certified and has in-depth knowledge of "Best-In-Class" OEMs and ISVs, giving us the flexibility to specify the best technology to drive business initiatives. We provide assessment, project management, implementation and residency services that support our delivery of IT as a Service, Next Gen Data Center, Information Management & Analytics, and End User Experience & Mobility solutions.

- ONX specializes in providing platform engineering of Hadoop environments.
- OnX has experience working with large fortune 500 Service Provider and Financial Service companies deploying their 1000+ node Hadoop clusters
- OnX has a seasoned and knowledgeable engineering team that can help you successfully deploy your Hadoop cluster
- OnX can provide Staff Augmentation and Project Delivery services to assist you after your cluster is deployed
- OnX is one of HPE's largest Enterprise Solutions partners in North America. We are an HPE Platinum Converged Infrastructure partner, the highest designation available.
- OnX provides our customers world-class Industry, Architecture and Project Management expertise in designing and integrating enterprise solutions.
- OnX has a proven track record of successfully delivering approximately 1,100 projects annually to our clients.
  - OnX project management methodology is applied to all our projects.
  - OnX has Project Managers on staff certified in ITIL, Agile and PMP practices.
  - OnX can deliver projects using the approach which the customer or project warrants whether waterfall, agile, or scrum.
- OnX applies Industry, OEM, and OnX Best Practices and lessons learned to reduce project risk and ensure a successful deployment.
- OnX's maintains certifications across a broad selection of best-in-class technology which allow OnX to provide an agnostic view point across industry technologies.
- OnX's relationship with market and technology leading vendors and their solutions allow us to maintain a high exposure to a wide varied of technology innovation across information, technology, process, organization, and enablement beyond our competitors.