IRON Big Data Appliance Platform for Hadoop

IRON Networks Big Data Appliance HDPOD is a comprehensive Hadoop Big Data platform, engineered for reliable data processing with a low overall total cost of ownership. The HDPOD combines network, server and storage hardware, and the comprehensive Hortonworks Data Platform (HDP) software suite into an optimized turnkey configuration. This creates a complete single SKU solution with integrated hardware, software and support services that is simple to acquire and deploy.

The HDPOD Appliance for Hadoop with Hortonworks Data Platform is built using industry-standard commodity hardware and offers low cost, fast performance with the latest generation of Intel technology combined with a fabric-based hardware configuration for enterprise-class data storage and management. The HDPOD appliance platform helps enterprises derive value almost instantly with low project risk and cost by providing a single platform for multi-workload data processing across an array of processing methods, from batch through interactive to real-time - all supported with solutions for governance, integration, security and operations.

Hortonworks, Hadoop Data Platform (HDP): Enabling your modern data architecture with enterprise Apache Hadoop

Apache Hadoop is a core component of the modern data architecture, integrating with and complementing your existing systems to create a highly efficient, highly scalable way to manage all your enterprise data. HDP's modern data architecture provides the foundation for your own enterprise Data Lake, which is an integral part of your business strategy for unleashing analytic insights and innovations that help you compete and win in today's marketplace. HDPOD platforms provides analysts the ability to leverage valuable customer insights (across all channels; web, mobile, social media, email, contact center, database, and storefront, and other type of customer interaction data) in their analysis while using familiar tools.

- Information Infrastructure Optimization, for faster and flexible application delivery and better business results.
- Information Management Standardization, eliminate islands of incompatible infrastructures.
- Information Insight Acceleration, faster intelligent business decisions.

Better Together: Hortonworks Hadoop Data Platform (HDP), IRON HDPOD Hardware and IRONClad Support Services

Accelerate the journey to big data using simple, efficient, flexible and open architecture

IRON Networks offers enterprise and service provider system designs that are reliable and energy-efficient, coupled with simplified serviceability and deployment services. Our solutions with power efficient design, high density configurations and leading-edge management software help organizations manage the modern data center and benefit from cloud and big data computing-faster and easier.

HDPOD is pre-configured and pre-integrated to achieve better results:

- Hadoop Standardized Architecture: A 100 percent open source Apache Hadoop-based software stack designed to store, process and analyze data in any format at any scale.
- Deep Integration: Purpose-built, low cost big data appliance that integrates hundreds of software packages including database, server, storage and big data technologies into a single, easy to manage, industry standard x86 system without vendor lock-ins.
- Simplicity: Automated management and provisioning, shorter POC deployment cycles and faster time to value for analytic initiatives, allows your experts to focus on innovation.
- Highly Scalable: Massively parallel analytic system allows customers to crunch through multi-terabyte data volumes.
- Simplified Experience: Optimized Single SKU configurations to accelerate purchase, deployment and support

Single SKU Big Data Infrastructure

- **Fully Integrated Solution:** Turnkey, Single vendor hardware and software solution for easy deployment
- **Low Cost POC Deployment:** Out-of-the-Box integration cuts down setup time, helps kicks start Big Data experience in less than a day
- **Commodity Hardware, Reduced TCO:** Up to 50% saving over traditional tier-1 Big Data Infrastructure solutions
- **High Bandwidth Networking:** Scalable to multiple racks for multi-petabytes of storage expansion
- **24x7 Support Services:** Expedited deployment, helpdesk support & on-site HW replacement services

HDPOD System 200 POC Solution Overview

Hadoop deployments require customers to acquire hardware to specific configurations, racking and configuring the hardware, operating system deployments and configurations, and installation of a huge set of Hadoop software packages with hundreds of pre-requisite packages. These deployments typically require several weeks to months of development cycle.

IRON HDPOD Big Data Appliance is designed to be a Hortonworks Data Platform (HDP) powered Hadoop Proof-of-Concept (POC) platform for running enterprise data analytics environments. It is a complete Hadoop Big Data environment which helps validate big data for target business and boot straps project for a small price tag, sold as a single SKU. It is simple to install as a PC for a very small cost-> plug into the corporate network directly, power up and it is ready to use immediately.

With ~100 (300 max) terabytes of storage, and 10 (2+8) nodes cluster, the HDPOD is powerful enough to be a full-fledged data analytics platform for any enterprise; customers can realize an ROI with the first few jobs. HDPOD can be used for testing, training, and later for production as-is or migrated to larger hardware platforms as desired.

Hortonworks Data Platform (HDP) Software – Key Features:

- **Fully Integrated Software Package:** HDPOD includes stable versions of all the critical Apache Hadoop components in an integrated and tested package
- **Easy Cluster Installation and Management:** HDPOD includes an installation and provisioning tool with a modern, intuitive user interface
- **Data Integration Services:** HDPOD includes Talend Open Studio for Big Data, the leading open source integration tool for easily connecting Hadoop to hundreds of data systems without having to write code
- **Metadata Cluster Services:** HDPOD includes Apache HCatalog, which simplifies data sharing between Hadoop applications and between Hadoop and other data systems
- **Cluster Management and Monitoring Services**: HDPOD includes intuitive dashboards for monitoring clusters and creating alerts
- **Operating Systems Flexibility:**: Supports Linux & Windows Server 2012

IRON HDPOD Hardware Design –Key Features:

- **End-End Integrated Platform:** Purchase as a single SKU
- **Out-of-Box, Turnkey Key Solution:** Hadoop infrastructure, self-contained within a 3 foot high rack. Entire Hortonwork Hadoop implementation pre-installed, custom-configured, QA'd and performance tested as complete environment at the factory
- **Zero IT department impact:** Requires a simple connection to the corporate network with a single IP address (for DHCP boot). Hadoop environment and all hardware in rack is isolated from corporate network
- **Enterprise features suitable for corporate datacenters:** Clustered in highly-available configurations, hot-swap hard-drives, power-supplies and compute blades for easy field service
- **Small Footprint:** 14U, Easy to be installed in departmental lab
- **Enterprise-Grade Packaging:**: Sized for Hadoop POC engagements; suitable as complete solution for small to mid-sized companies

Network Fabric

Single Switch, Dual (dual path) optional

- Data & IPMI Management Switch: Single, 48 Port 1GbE RJ45

IRON Cluster Management Server Appliance

Management Node, Advance Monitoring, Administration

- Out of Box Experience: Simple Setup, Automated bare metal cluster provisioning using IRON Management SW
- Local Hadoop Repository: Complete Hadoop Software integration, Linux, EPEL, various dependent packages
- Hardware Health and Scale Up/Out management

Name Nodes Cluster (x2 Intel x86 Server Nodes)

Optimized for data node orchestration and HA

- Two, Twin Server blade nodes (2U) for high availability, RAID Boot
- Dual Intel 4-Core Sandy Bridge Xeon Processors @ 2.4Ghz, 64GB of memory, and ten 1TB SAS 7.2K HDDs

Data Nodes Cluster (x8 Intel x86 Server Nodes)

Optimized for data storage and staging at large scale

- Total 8 nodes using 2x Blade Chassis (4 nodes per Chassis)
- Dual Intel 6-Core Sandy Bridge Xeon Processors @ 2.0Ghz, 128GB of memory and six 2TB (max 4TB or 6TB HDD size) HDDs; Total 48x2=96TB capacity- default configuration
- Dual redundant Power Supply; All components Hot Swappable for easy field replacement

HDPOD System 300 Solution Design Architecture Overview

IRON HDPOD is designed to be a Hortonworks Hardoop HDFS powered Big Data storage platform for running enterprise data analytics environments, the design offers:

- **Rack-Level Resiliency:** Resiliency is built into the design at multiple levels to deal with hardware and software failures; all components have N+1. The degree of resiliency is adaptable for mitigating power, network, storage and compute outages.
- **Modular, Scalable Storage Architecture:** Balanced compute CPU, memory, network throughput and storage IOPS and capacity per rack. Multiple racks can be networked together to build scale-out modular storage infrastructure
- **Flexible Data Node Configurations:** For Hadoop workloads, the storage to compute ratios vary from application to application; HDPOD offers multiple building block choices in 2U, 1U and half-U form factors, as well as Low Power consumption CPU options.
- **High Return-on-Investment (ROI):** Best-in-class standardized components to provide maximum performance leveraging synergies between components that are based on open standards, reliable and scalable, allowing organizations to save on cost and help grow with business needs.

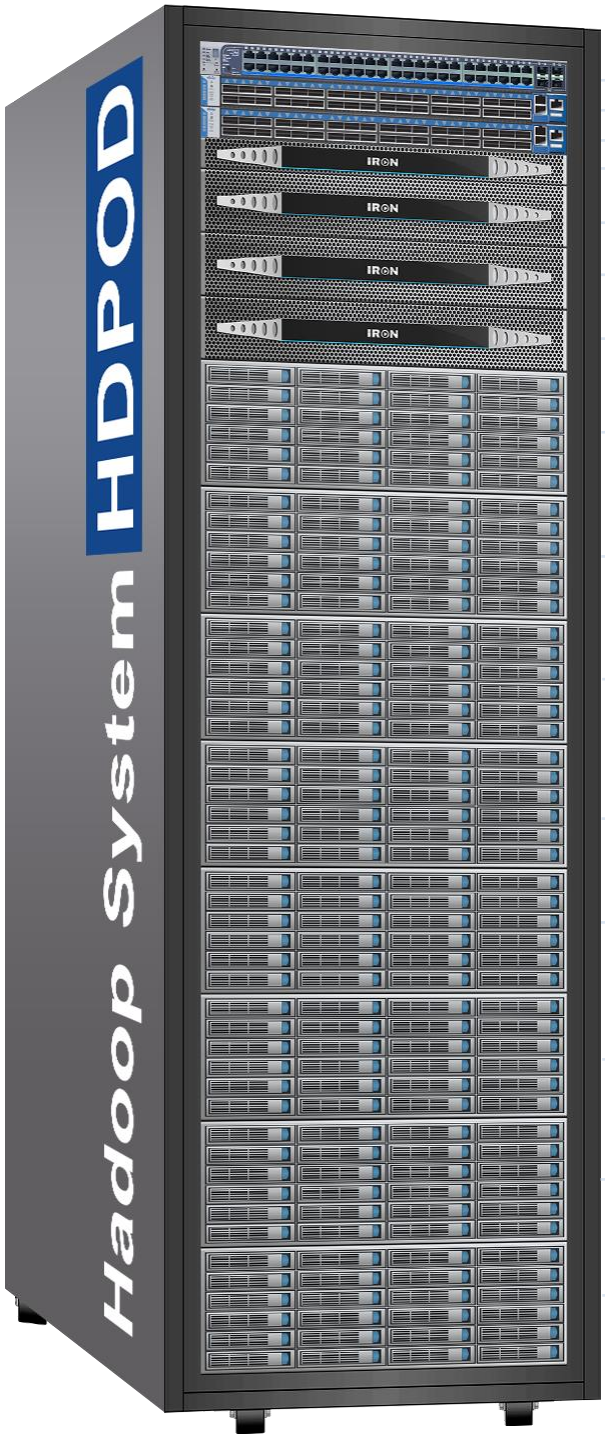

High Speed Network Fabric

Supports Dual path for high availability

- High Speed 40GbE across the entire HDPOD 300 series; this fast, dual path interconnect can also be used to connect to external Analytics infrastructure
- Data Switch: Single or Dual*, 36 Port 40GbE QSF (64 10GbE SFP+, uses 4x SFP+ breakout cables) , Managed; 10GbE SFP+ connectivity to servers and 40/56GbE uplink to next level of aggregation or east-west switching
- Network Management Switch: Single, 48 Port 1GbE RJ45 connected to IPMI Server ports
- Scalable to multiple racks for multi-petabytes expansion

IRON Cluster Management Server Appliance

Management Node-Advance Monitoring, administration, and analysis of Hadoop HW & SW

- Out of Box Experience: Simple Setup, Automated bare metal cluster provisioning using IRON Management SW
- Local Hadoop Repository: Complete Hadoop Software integration, Linux, EPEL, various dependent packages
- Hardware Health and Scale Up/Out management

Name Nodes Cluster (x3 Intel x86 Server Nodes)

Optimized for data node orchestration and HA

- Three, 2U server nodes for high availability, RAID boot
- Dual Intel 12-Core Ivy Bridge Xeon Processors @ 2.6Ghz, 256GB of memory, and ten 900GB SAS 10K HDDs
- Dual 10GbE SFP+ Network paths, 1x RJ45 IPMI LOM Port
- Dual redundant Power Supply

Data Nodes Cluster (x32 Intel x86 Server Nodes)

Optimized for data storage and staging at large scale

- Total 32 nodes using 8x Blade Chassis (4 nodes per Chassis)
- Dual Intel 6-Core Ivy Bridge Xeon Processors @ 2.0Ghz, 128GB of memory, and six 4TB (multiple HDD to node ratio config. available) HDDs per node.
- Dual 10GbE SFP+ Network paths, 1x RJ45 IPMI LOM Port
- Dual redundant Power Supply; All components hot swappable for easy field replacement

IRON HDPOD Commodity Systems: Hardware Configurations Options

Model: H220
(192TB Max. Storage)

Model: H320SH
(384TB Storage)

Model: H320SF
(768TB Storage)

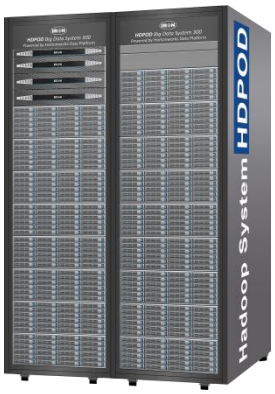

Model: H320S-P2
(1.63PB Storage)

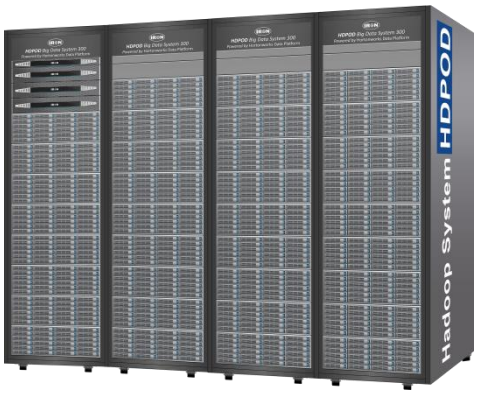

Model: H330S-P4
(3.36PB Storage)

Hardware Components & Specifications		H220SH	H320SF	H320S-P2	H320-P4
Network Fabric					
• Aggregate/ Data Storage Switch:	1U, 36x 40/56GbE QSFP Ports Ethernet Switch	2	2	4	8
• Management Switch:	1U, 44x 1GbE RJ45, 4x SFP Ports Ethernet Switch	1	1	2	4
Cluster Management Server Appliance					
• Management Node Server:	1U, IRON iServer Single Node, 2x 10GbE & 1-IPMI Ports	1	1	1	1
Name Nodes for Hadoop					
• Nodes (Master Server)	2U, IRON iServer Single Node, 2x 10GbE & 1-IPMI Ports	3	3	3	3
• CPU:	Intel Dual Socket, 12 Core, 2.6Ghz CPU (Total 24 Cores/Node)				
• Memory (GB):	96, 128, 192, 256, 384 & 512 Options (256GB default)				
• Disk Drives	Number of Disks – 10x; 900GB 10K RPM (per Node)				
Three nodes: 1x Primary named node, 1x Secondary named nodes (+HBase Master, Hive Server), 1x Job Tracker node					
Data Storage Nodes for Hadoop HDFS					
• Blade Enclosures:	4U, IRON iServer, Quad Node Blade Chassis (4 blades)	4	8	17	24
• Nodes (Server Blades):	4 Server Nodes/Chassis, 2x 40Gbe LAN & 1-IPMI Ports	16	32	68	140
• Processors (CPU/Cores):	Dual Socket, 6 Core, 2.0Ghz CPU (Total 12 Cores/Blade)	32/192	64/384	136/816	280/1,680
• Memory (GB):	96, 128, 192, 256, 384 & 512 Options (128GB default)	2,048	4,096	8,704	17,920
• Disk Drives	Number of Disks – 6,8 or 12 per Node (6 default)	96	192	408	840
Total Disk Capacity (Raw*)	1TB, 2TB, 3TB & 6TB SAS (4TB default)	384	768	1,632	3,360
**Capacities listed uncompressed (3x compression assumed typical) without Compression*					
Total Disk Capacity (Usable)	Usable mirrored capacity (appr.)	192	384	816	1,680
Total Disk Capacity (Usable)	Usable Triple mirrored capacity (appr.)	128	256	644	1,120
Additional Hardware and Software					
• Rack	IRON 42U Rack, Redundant PDUs	1	1	2	4
• KVM	IRON Keyboard/Mouse/Video Display Console, Optional				
• Software (System)	Linux Operating System, 64bit				
	Microsoft Windows Server 2012 R2 Optional				
• Software (Add-On), Optional	Data Protection for Volume & MapReduce layers				
	Dynamic Data Set management, metadata, and data lineage				

IRONClad Single-Call Support Services

Installation and Support Services are included to help accelerate your big data solutions
Our highly experienced and expertly trained sales professionals and partners deliver a converged infrastructure precisely the way you need it. IRON also offers a variety of services and tools to help our customers get started with confidence at their own pace.

Hortonworks Powered converged infrastructure reference design guide offers detailed configurations, best practice and lessons learned. These offering includes solution planning for data management, data access, integration, security and operations for automated provisioning and management.

IRON Big Data Infrastructure Services can help you design, finance, implement, and support a converged infrastructure. These services include a visioning workshop, planning services, design and implementation service and proof of concepts.

Efficient Architecture Big Data Workshop allows you to spend a day with an experienced big data infrastructure expert to help learn fundamental modern data infrastructure architecture and integration topics with best practices using Hortonworks Data Platform.

Open Standards and Open Integration building blocks allows IRON to provide the most complete, simplest, and fastest way to deploy the right solution or solutions to meet your varying requirements-all built to industry standards. This gives you the ability to change components in and out as you like. You're not locked in. This is a very important part of our architecture. It's not about lock-in. It's about building to industry standards and giving our customers choice so they can accelerate IT to deliver better business value.

Onsite installation and implementation is included with every HDPOD deployment, complete with an orientation training session and configuration support. Iron support provides the foundation for secure and reliable high-availability infrastructures with enhanced hardware support and software technical support for problem resolution.

The HDPOD Infrastructure family of solutions demonstrates how open and flexible architectures powered by core Hortonworks hadoop technologies combined with IRON Services create practical, innovative solutions designed to reduce complexity and costs while improving productivity. IRON Services can help deliver elastic IT, whether you are building a platform for private cloud computing or for better availability and management of applications.

About Iron Networks

Based in Fremont, CA, Iron Networks is a Microsoft OEM partner that delivers turnkey networking and infrastructure platforms for Microsoft-based Software Defined Data Centers. With over 1000 successful Microsoft infrastructure appliance implementations around the globe, Iron Networks has a depth of experience building, deploying and supporting Windows Server based solutions. Utilizing the product and market experience with Microsoft-technologies and the core competencies as a leader in turnkey rack level hardware manufacturing, OEM appliance design, integration and supply chain management solutions, Iron Networks is positioned to rapidly deploy and support enterprise grade cloud and datacenter solutions.

For more information, visit www.ironnetworks.com or contact an authorized HDPOD reseller.

Iron Networks, Inc.
980 Mission Court, Fremont, CA 94539, USA
Phone: 408-895-5000 (Local), +(1) 877-895-6277 (US-Toll Free), +(1) 408-895-5000 (International)
Fax: +(1) 408-943-8222/8101 Email: info@ironnetworks.com Website: www.ironnetworks.com

For further information, please contact:
Iron Networks, Inc.
980 Mission Court, Fremont, CA 94539, USA
Phone: +(1) 408-895-5000 (Local) +(1) 877-895-6277 (US-Toll Free)
Fax: +(1) 408-943-8101
Email: info@ironnetworks.com