


BlueData on Dell PowerEdge Servers

A Quick Reference Configuration Guide

Kris Applegate
Solution Architect
Dell Customer Solution Centers



Summary

This document details the configuration set-up for BlueData's Big-Data-as-a-Service (BDaaS) platform on the PowerEdge Servers. The intended audiences for this document are customers and system architects looking for information on configuring BlueData clusters within their environment for use in performing Big Data analytics. These clusters can host multiple virtual Hadoop and Spark clusters running multiple distributions and software versions simultaneously.

This document will only focus on the Dell hardware configuration; it will not go into detail about information already covered in BlueData's marketing material. For the most current best practices for BlueData deployment, please refer to documentation from their website Dell developed this document to help streamline configuration for BlueData EPIC software on Dell PowerEdge servers.

About BlueData Software, Inc.

BlueData is transforming how enterprises deploy their Big Data applications and infrastructure. The BlueData EPIC™ software platform uses container technology to make it easier, faster, and more cost-effective for enterprises of all sizes to leverage Big Data – enabling Big-Data-as-a-Service in an on-premises deployment model. With BlueData, they can spin up virtual Hadoop or Spark clusters within minutes, providing data scientists with on-demand access to the applications, data and infrastructure they need. Based in Santa Clara, California, BlueData was founded by VMware veterans and its investors including Amplify Partners, Atlantic Bridge, Ignition Partners, and Intel Capital. To learn more about BlueData, visit <http://www.bluedata.com> or follow @bluedatainc.

THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.

© 2016 Dell Inc. All rights reserved. Reproduction of this material in any manner whatsoever without the express written permission of Dell Inc. is strictly forbidden. For more information, contact Dell. Dell, the DELL logo, and the DELL badge are trademarks of Dell Inc. Intel, the Intel logo, Xeon, and Xeon Inside are trademarks or registered trademarks of Intel Corporation in the U.S. and/or other countries. Red Hat is a registered trademark of Red Hat Inc. Linux is a registered trademark of Linus Torvalds. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell Inc. disclaims any proprietary interest in trademarks and trade names other than its own.

February 2016



1 Reference Configuration Purpose

In an effort to better assist our customers looking to adopt the latest disruptive technologies, Dell produces whitepapers that detail known working hardware configurations that have been validated to pass basic functional and performance criteria. In addition, we make some prescriptive recommendations for additional configurations that may be leveraged to meet different platform, performance, or density requirements. When selecting the best platform to run BlueData on, please make sure to consult with your BlueData and Dell account team's technical specialists so that you can incorporate the latest and greatest best practices into your configuration.

The recommendations laid out in this document are just that, recommendations. They should not be relied upon to provide the perfect configuration for every use-case. Please use them as a conversation starter and optimize from there.



2 Dell Customer Solution Centers

The Dell Customer Solution Centers are a global network of connected labs that allow Dell to help customers architect, validate and build solutions. With footprints in every region, they can help you whether through an informal 30-60 minute briefing, a longer half-day workshop, or even on to a proof-of-concept that would allow you to kick the tires of a solution prior to signing on the dotted line. Simply engage with your account team and have them submit a request to get started today.



3 About BlueData Software

Two of the most disruptive trends in recent IT memory are Cloud and Big Data. These two technologies have been running straight at each other for years. While many public cloud implementations of Hadoop use virtualized platforms to provide their services, customers with on premise requirements have been looking for something that offers the flexibility of virtualization but with the performance of bare-metal.

BlueData's EPIC software provides a simple on premise platform for delivering Big-Data-as-a-Service to your enterprise. BlueData's ability to seamlessly deliver a single shared platform for multiple distributions and versions of Hadoop, Spark, and other BI / Analytics tools is very useful in the modern reality of the new Future-Ready Enterprise. Whether it's the need to support separate business units disparate Hadoop distribution requirements (e.g. Cloudera versus Hortonworks) or to support two different versions of Hadoop for two different BI tool-chains, the BlueData EPIC software platform can pool all these resources on the same bare-metal hardware stack.

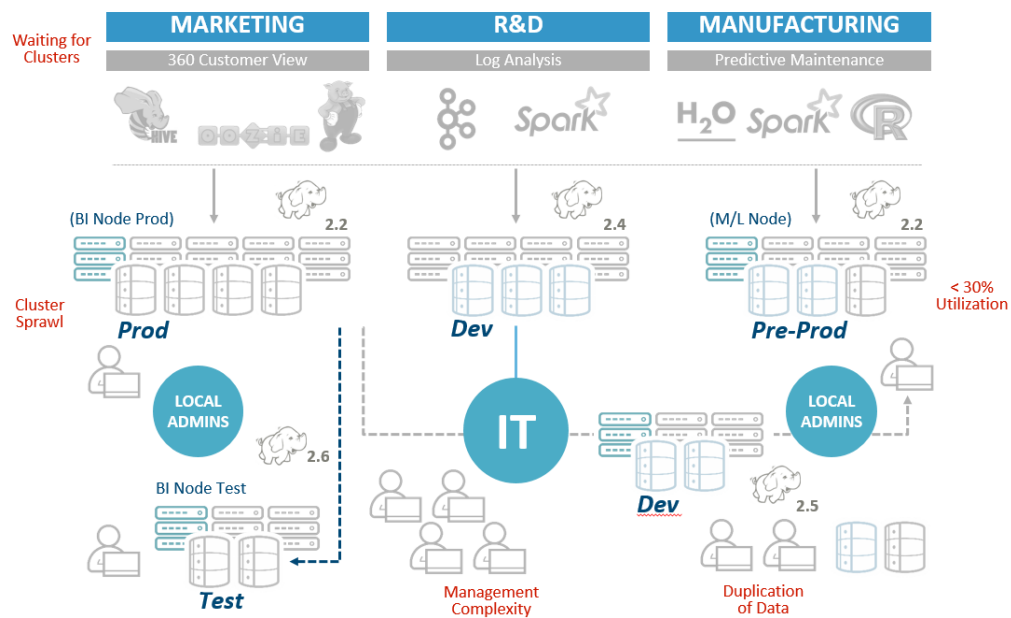


Figure 1. Multiple big data clusters across the enterprise

Using the latest technologies like containerization and software-defined networking, this solution can modernize your data platform, lower the complexity, and reduce the resources needed to deliver to your stake-holder's business requirements. Standing up a cluster is as simple as pointing your users to the self-service web page. From there, they can stand up a cluster using known-tested versions of multiple vendors distributions. They operate inside the quotas and resource pools that you, as an administrator, establish for them. Inside a couple minutes they can have a fully functional Hadoop or Spark cluster ready to run their jobs, all running on a shared pool of resources alongside their co-workers.



blue data

- Dashboard
- Jobs
- Clusters
- DataTaps
- Nodes
- Users

Create New Cluster

Cluster Name

Select Cluster Type

Cluster FS Path
Optional

Distribution

MR Type

Master Node Flavor

Worker Count

Worker Node Flavor

YARN High Availability ☐

Pig, Hive, Oozie, Impala & Hue ☐

Hunk ☐

[Advanced Settings](#)

Figure 2. Creating a Hadoop Cluster

The ability of BlueData's Datatap technology provides the ability to seamlessly abstract the underlying big data storage layer away from the computational workload, gives you the flexibility to leverage data sources from all over your enterprise. Whether you want to execute a job against data on the BlueData data nodes themselves, data from another Hadoop cluster, or even data from object or NFS storage, you can do that in a way that makes it painless for your users.

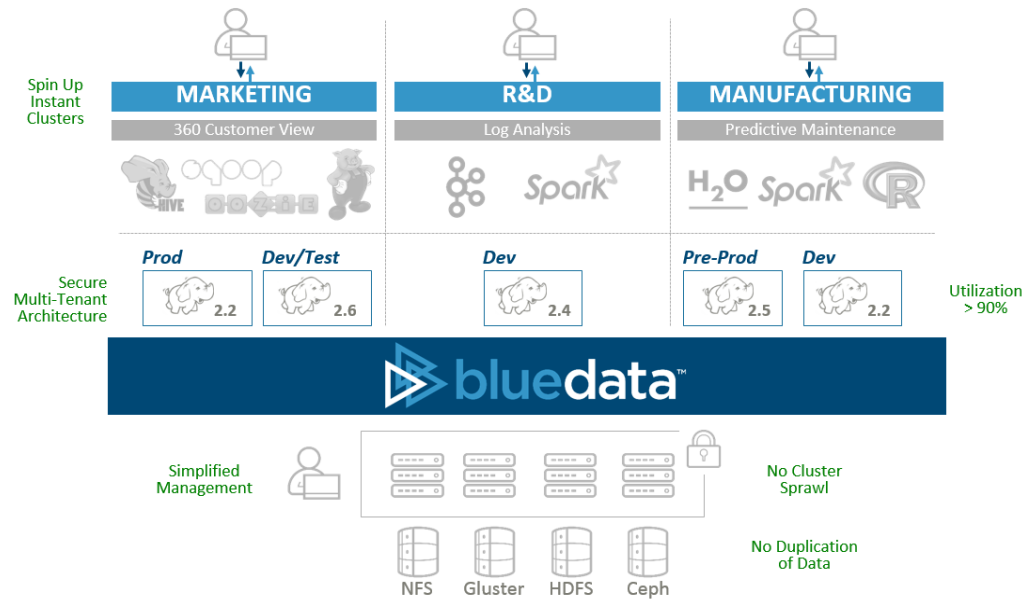


Figure 3. BlueData-Enabled Big Data Deployment

The screenshot shows the **Tenant Management** interface in the BlueData admin console. The interface includes a sidebar with navigation links: Dashboard, Tenants, Users, Hosts, Nodes, Installation, App Store, Settings, and Support. The main content area displays a table of tenants with columns for Tenant Name, Tenant Description, Tenant Quota, and Actions. The table lists six tenants: Site Admin, Marketing, Production, DELL, Dev/Test, and Demo Tenant. Each tenant has a specific quota for Cores and Memory. The interface also includes a search bar, a 'Create' button, and a 'Delete' button.

Tenant Name	Tenant Description	Tenant Quota	Actions
Site Admin	Site Admin Tenant for BlueData clusters		Info
Marketing	Marketing Team	Cores: 0 / 400 Memory: 0 / 545 GB	Edit Delete Info
Production	Production Team	Cores: 0 / 400 Memory: 0 / 545 GB	Edit Delete Info
DELL	Dell	Cores: 0 / 40 Memory: 0 / 54 GB	Edit Delete Info
Dev/Test	Development and Testing	Cores: 0 / 400 Memory: 0 / 545 GB	Edit Delete Info
Demo Tenant	Demo Tenant for BlueData Clusters	Cores: 1060 (No Quota) Memory: 1318 GB (No Quota)	Edit Delete Info

Showing 1 to 6 of 6 entries

Figure 4. Multiple Tenant Administration



4 About Dell Hardware

Dell PowerEdge Servers

Dell's 13th generation of servers bring a greater degree of Big Data customizations. In the previous generation of 2U disk-centric servers there were only a handful of configuration options. Dell is showing its commitment to the Big Data space by offering 10 different configuration options available on this one platform alone. Whether your needs are capacity or performance driven, there is a platform for you.

Choice of Dell platforms isn't limited to 2U or to standard rack chassis. If you want an ultra-dense rack server configuration, you could easily use a 1U Dell PowerEdge R630 as a data node. It shares many of the same internals as the 2U counter-part, so you can have a high level of confidence in compatibility. There is also the more modular approach of the Dell PowerEdge FX2 chassis. This configuration could allow you twice as many compute nodes as an R730XD based approach. If density is really a huge concern, you could even go with 4x FC430 + 2x FD332 sleds for 4 dual-socket high-end servers in 2U.



Figure 5. Dell PowerEdge R730XD Models (12x3.5", 24x2.5", and 16x1.8"+8x3.5")



Figure 6. Dell PowerEdge FX2 with (2) FC630 Compute Sleds and (2) FD332 Disk Sleds



Dell Networking Ethernet Switches

Dell's robust portfolio of switches offers many choices for the management, access, and aggregation layers. In an industry first, Dell also offers switches with your choice of switch OS. Whether it's the proven Dell Networking OS or an alternate Linux-based OS, your freedom of choice allows this to be integrated into the latest/greatest networking topologies.

The management network is simply providing out-of-band access to DRACs, CMCs, and switches. These are just standard run-of-the-mill managed 1GbE switches like the Dell Networking S3048-ON.

The data network (Top-of-Rack /access/leaf network) is the principle data transit path and thus, needs to be robust. We recommend a switch with plenty of non-blocking bandwidth and deep per-port or shared packet buffers to tolerate the storm of activity Hadoop can produce. Dell Networking's S4048-ON is the perfect switch for this role. If needed, you can aggregate these with a 40GbE upstream switch like a Dell Networking S6000-ON.






Figure 7. Networking Diagram

5 Recommended Configurations


The configurations outlined below range from the “as-tested” configuration we validated in the lab on through a couple popular configuration recommendations for the variety of use-cases we see on a daily basis. These configurations should not be perceived as rigid inflexible templates, but rather as conversation starters. Debating over which of the configurations you are targeting can often times produce healthy dialogue over priorities and expectations. Please engage your Dell and BlueData account teams to help in this discourse since they’ll have the latest information available around changes to the below recommendations. You also have the Dell Customer Solutions Center’s subject-matter experts that can provide real-world guidance as to what we’ve seen work for other customers in similar situations.

As-Tested Configuration

Component Management Nodes (1-3)		As-Tested Configuration
Model		Dell PowerEdge R730XD (16x 3.5”+2x 2.5” Flexbay)
Processor		(2) Intel E5-2680v3 2.5GHz 12-core Xeon Processors
Memory		256 GB RAM
OS Disks		(2) 600GB 10K RPM SAS Drives (Flexbay) RAID1
Controller		Dell PERC 730P
Data Disks		(16) 4TB 7.2K RPM SATA NON-RAID Mode
Network Card		Intel I350 1GbE + X520 10GbE Network Daughter Card
Data Nodes (3+)		
Model		Dell PowerEdge R730XD (16x 3.5”+2x 2.5” Flexbay)
Processor		(2) Intel E5-2680v3 2.5GHz 12-core Xeon Processors
Memory		256 GB RAM
OS Disks		(2) 600GB 10K RPM SAS Drives (Flexbay) RAID1
Controller		Dell PERC 730P
Data Disks		(16) 4TB 7.2K RPM SATA NON-RAID Mode
Network Card		Intel I350 1GbE + X520 10GbE Network Daughter Card
Networking		
Management Switch		(2) Dell Networking S3048-ON
Data Switches		(2) Dell Networking S40480-ON

Table 1. As-Tested Configuration

Potential Configuration Recommendations




Component	Recommended Capacity-Optimized Configuration
-----------	--



Management Nodes (1-3)	
Model	Dell PowerEdge R730XD (12x 3.5"+2x 2.5" Flexbay)
Processor	(2) Intel E5-2680v3 2.5GHz 12-core Xeon Processors
Memory	256 GB RAM
OS Disks	(2) 600GB 10K RPM SAS Drives (Flexbay) RAID1
Controller	Dell PERC 730P
Data Disks	(12) 4TB 7.2K RPM SATA NON-RAID Mode
Network Card	Intel I350 1GbE + X520 10GbE Network Daughter Card
Data Nodes (3+)	
Model	Dell PowerEdge R730XD (12x 3.5"+2x 2.5" Flexbay)
Processor	(2) Intel E5-2680v3 2.5GHz 12-core Xeon Processors
Memory	256 GB RAM
OS Disks	(2) 600GB 10K RPM SAS Drives (Flexbay) RAID1
Controller	Dell PERC 730P
Data Disks	(12) 4TB 7.2K RPM SATA NON-RAID Mode
Network Card	Intel I350 1GbE + X520 10GbE Network Daughter Card
Networking	
Management Switch	(2) Dell Networking S3048-ON
Data Switches	(2) Dell Networking S40480-ON


Table 2. Recommended Capacity-Optimized Configuration



Component	Recommended Performance-Optimized Configuration
Management Nodes (1-3)	
Model	Dell PowerEdge R730XD (24x 2.5"+2x 2.5" Flexbay)
Processor	(2) Intel E5-2680v3 2.5GHz 12-core Xeon Processors
Memory	256 GB RAM
OS Disks	(2) 600GB 10K RPM SAS Drives (Flexbay) RAID1
Controller	Dell PERC 730P
Data Disks	(24) 1.2TB 10K RPM SAS NON-RAID Mode
Network Card	Intel I350 1GbE + X520 10GbE Network Daughter Card
Data Nodes (3+)	
Model	Dell PowerEdge R730XD (24x 2.5"+2x 2.5" Flexbay)
Processor	(2) Intel E5-2680v3 2.5GHz 12-core Xeon Processors
Memory	256 GB RAM
OS Disks	(2) 600GB 10K RPM SAS Drives (Flexbay) RAID1
Controller	Dell PERC 730P
Data Disks	(24) 1.2TB 10K RPM SAS NON-RAID Mode
Network Card	Intel I350 1GbE + X520 10GbE Network Daughter Card
Networking	
Management Switch	(2) Dell Networking S3048-ON
Data Switches	(2) Dell Networking S40480-ON



Table 3. Recommended Performance-Optimized Configuration



Component Management / Data Chassis (2+ with 2x nodes in each Chassis)		Recommended Modular Configuration	
Chassis		Dell PowerEdge FX2	
Chassis IO Module		10-GbE Pass-through Module	
Model		(2) Dell PowerEdge FC630	
Processor		(2) Intel E5-2680v3 2.5GHz 12-core Xeon Processors	
Memory		256 GB RAM	
OS Disks		(2) 600GB 10K RPM SAS Drives (Flexbay) RAID1	
Controller		Dell PERC 730P	
Network Card		Intel I350 1GbE + X520 10GbE Network Daughter Card	
Disk Sled		(2) Dell PowerEdge FD332 NON-RAID Mode	
Disk Sled Drives		(16) 1.2 TB 10K RPM SAS	
Networking			
Management Switch		(2) Dell Networking S3048-ON	
Data Switches		(2) Dell Networking S40480-ON	

Table 4. Recommended Modular Configuration



6 Configuration Notes

Deployment

Deployment can be performed with a variety of methods. The requirement to start the BlueData installation process is to stand up the box running a supported Linux configuration. We see customers either use a tool like Foreman, Dell Active System Manager, Dell's Deployment Toolkit for Linux, or even homebrew Kickstart scripts to get the initial bare metal configuration and OS installed.

Disk Configuration

OS Disks should be in a RAID1 mirror for protection. BlueData recommends having at least 300GB free on the root (/) partition in order to begin the installation. Data drives should all be flagged as NON-RAID in the PERC controller and remain unformatted and unmounted.

Operating System

The "As-Tested" configuration used CentOS 6.7 x64. You should check with BlueData for the latest supported OS requirements.

Network Bonding

We recommend you bond the 10GbE ports together either in a software-bond (ALB mode 5) or in an LACP bond (with the appropriate configuration at the switch).

Transparent Huge Page (THP) Compaction

Red Hat Enterprise Linux Server and derivatives attempt to reduce the number of huge pages in use by defragmenting the used memory blocks. There is a performance cost to this operation. Dell recommends that this functionality be turned off for a Hadoop cluster by executing the following command, and adding it to the rc.local file.

```
# echo never > /sys/kernel/mm/redhat_transparent_hugepage/defrag
```

Swap Settings

Dell recommends that vm.swappiness be set based on the Linux kernel version.

To check the kernel version, run:

```
# uname -a
```

To check the vm.swappiness parameter setting, run:

```
# cat /proc/sys/vm/swappiness
```

To set the vm.swappiness parameter for kernel versions earlier than 2.6.32-303:

```
# sysctl -w vm.swappiness=0
```

To set the vm.swappiness parameter for later kernel versions:

```
# sysctl -w vm.swappiness=1
```



7

Resources

BlueData's Homepage (<http://www.BlueData.com>)

Dell's Customer Solution Centers (<http://Dell.com/SolutionCenter>)

