FLORIDA DEPARTMENT OF TRANSPORTATION PRODUCTION BIG DATA PLATFORM

RECOMMENDATION AND JUSTIFACTION

Executive Summary:

VHB has been tasked by the Florida Department of Transportation – District Five to design and build a big data platform that will allow the district to store, consume, manage and analyze datasets from various sources. VHB is recommending the production environment to consist of a stack of technologies and architectural methodologies. Together, the technologies and methods will allow the district to have a big data system that is fast, highly-available, secure, flexible and scalable. The district's big data platform will also utilize a combination of virtual and physical hardware. This approach will allow the various components of the big data platform to reside on a platform that will best take advantage of component's job and/or task.

It is recommended that the district's big data platform is composed by various servers that will perform select functions. As an example, the Microsoft SQL server will be utilized in the big data platform to store and manage datasets that tend to be more static in nature and requires a highly structured relational layout. Utilizing SQL server will allow the district to re-use resources that have been already invested into the district's SQL Server cluster. SQL Server will also be used by the district as a required component of ArcGIS Server.

Within the proposed FDOT District 5 big data production platform, will be various servers dedicated to perform task. There will be pre and post processing servers that will be used within the big data production platform as an engine or work horse systems dedicated to ingest, validate, extract, transform and load data, and perform long running tasks.

Another component of the production big data platform is the ArcGIS Server. This server will store, manage and run many of the spatially enabled data of the district.

The production architecture also contains a Web and API servers. The Web and API servers will be designated to manage the web applications and the be the API gateway for the big data environment. The recommended architecture is following the industry best practice principles and each sub system, after careful evaluation, is been put in place to maximize performance of the overall platform. The proposed architecture will have virtualization to make more efficient use of the big data platform infrastructure. Virtualization allows to consolidate and run applications onto fewer physical servers, which drives up server utilization rates and cuts down on operational costs. Additionally, virtualization enables quick provisioning and deployment, improved workload balancing, and enhanced resiliency and availability by giving the ability to dynamically move virtual machines from server to server. Since the environment in question is a production platform, it is recommended that the district utilizes the industry leader and tested virtualization software such as VMWare. With Virtualization, the district will be able to take advantage of their current virtual server investment to build, manage and maintain these server components that are a part of the district's production big data platform.

The District 5 big data component will be managed by Hadoop via the Cloudera distribution and Elasticsearch. Cloudera offers a comprehensive package of the Apache Hadoop system that will allow the district to manage the big data platform from a series of modules and dashboards. These modules and dashboards are used to assist in being able to manage the Hadoop layer of the district's big data platform from one centralized location. Elasticsearch will also serve as another big data managing system in the district's big data platform. Elasticsearch is being recommended to enable integration with the district geospatial data and infrastructure. The district currently utilizes the ESRI products stack and Elasticsearch is the big data platform that ESRI utilizes in their software suite. There are many current and upcoming out of the box features within the ESRI platform that tie directly into an Elasticsearch data source. Elasticsearch will allow the district to utilize their already invested ESRI resources with the districts big data platform while benefiting from the speed of Elasticsearch and the power of Hadoop via the Cloudera distribution.

Why select a big data store environment?

There are many advantages of utilizing a big data store environment. Traditionally, the use of a fully upgraded relational database server(s) was used to manage and host existing data by the district. While this approach has its strengths, the investment to migrate from a SQL environment to a No-SQL environment allows the district to operate within a methodology that is mature, proven, and scalable and will continue to evolve.

Items of concern such as computing ceilings and points of failure are less of an issue when the district's backend is built utilizing a big data platform. Computing ceilings are when a system and the associated databases are reaching and / or exceeding the computational limits of the server. Traditionally, to resolve this issue, the district would upgrade or add another server within the database server cluster to add more computing resources. The downside to this approach is that upgrading a server can only yield a certain amount of additional resource before the server's max configuration is met. This concept is known as '**Scaling-Up'**. Meaning, the district will upgrade a server to meet the requirements needed for the tasks being performed. Another option for the district would be to add another SQL server to the backend SQL cluster. The downside to this approach is that it can be cost prohibitive. The cost of hardware and software licenses can add up quickly.

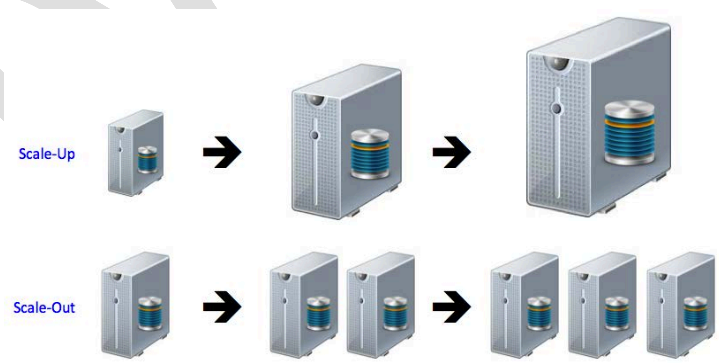

Figure 1 Source Image: Google Images

Utilizing a big data platform can reduce these issues by allowing the backend cluster to '**Scale-Out'**. To 'Scale-Out' means that instead of upgrading the server to increase computing resources, another commodity server is added to the cluster to add additional resources. Commodity servers are lower cost hardware that can be used to scale the district's backend instead of using highly upgraded and high cost servers. These commodity servers should be physical servers instead of virtualized servers.

Utilizing virtualized servers to expand the district's big data platform will eventually lead to the same problems that utilizing a highly upgraded server would have. At some point, the server hosting the virtual servers will reach its computational resource limit, and at that point the district will have to setup another high cost server that can host more virtual machines.

This approach also means that the virtual host server will serve as a single point of failure for the districts big data platform. If the virtual host server starts to exhibit problems and goes offline, any associated servers that were a part of the host server will also go offline and can cause the district's big data platform to be offline until the virtual host server is restored. By utilizing less expensive and physical commodity servers in place of virtual servers, the district will have the ability to build a platform that is more stable and will not be affected by a single point of failure. If one of the Big data servers starts to have problems and goes offline, the district will only have to replace the failed hardware with a new commodity server, and the Big data platform will reconfigure the server back into the cluster automatically.

The district's big data platform will continue to grow and evolve as more datasets are introduced to the platform. The rate of growth of the district's platform can grow quickly and thus a plan to expand storage and servers as needed will play a vital role to keep the district's big data platform operational and performing at its best. Various skillsets will be required by the district in order to maintain the big data platform. Skillsets such as Windows and Linux Server administration, Hadoop, Cloudera and Elasticsearch skillsets are some of the vital skills required for the management of the district's big data platform.

Do All the Servers in The District's Big Data Platform Need to Be Physical Servers?

While best practices would suggest that the servers within the data node cluster should be physical hardware commodity servers, not every server that will be part of the district's big data platform will need to be a physical server. Utilizing virtualized software such as VMWare to make servers available outside the data node cluster **can** and **should be** used. Utilizing virtualized technology allows for easy and repeatable setups of various server types. Due to the scalability features of the data nodes within the big data platform, the data nodes should always be physical hardware servers. This allows the district to repair and scale out the data node layer as needed using commodity servers. The other servers within the Big data platform will also benefit from being in a virtualized environment. The district can create virtual image files of the various server types and add them to the virtual host server as needed. Servers such as the Web / API Server, SQL Server Instances, Pre and Post Processing Servers, and ArcGIS Servers are server types that can be virtualized. The district will want to setup and maintain running, replicated setups of the virtual host server at all times. Maintaining replicated setups of the virtual host server will allow the district to maintain a high availability

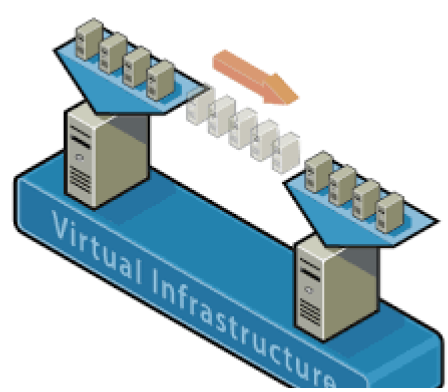

Figure 2 Source Image: Google Images

of the servers in case the main virtual host goes offline. If the main virtual host does go offline, the district will roll over to the replicated virtual host thus minimizing any downtime.

SQL vs No-SQL and The District's Big Data Platform.

The district's big data platform will consist of a mixture of data storing and management technologies. Together, this mix of technologies will allow the district to utilize various datasets in a technology platform that will allow the data to be used in the most efficient way. Microsoft SQL Server will be one of the platforms used within the districts big data platform. SQL Server utilizes SQL and SQL methodologies to maintain, query and retain data. SQL Server is also known as a RDBMS (**R**elational **D**ata**b**ase **M**anagement **S**ystem). SQL type data is data that can be broken down into tabular form. This data consists of tables and within these tables are rows of data. Each row of data is formatted exactly alike and will adhere to what is known as a schema. A schema is simply a predefined format in which the data will be stored. Each column in a row is associated with its column header / name. SQL datasets can reside within various tables in the database. These tables relate to one another through what are known as keys. Utilizing a language known as SQL, the district can construct SQL queries to join, select, update and delete data from the various tables. RDBMS systems such as SQL Server has its advantages. Concepts and features such a Relationships, Structure, and Transactional Logging are benefits within a RDBMS system. But sometimes the same benefits that makes RDBMS systems great can also be a con. Being structured and relational can be a drawback in a data system that is ever evolving. If a query starts to utilized many 'Joins' aggregate its dataset, the joins will start to become a heavy computational expense to the district's system and can slow performance. If a dataset needs to expand in attributes, then it is not as simple as adding the needed attributes into the dataset. The database schema will need to be updated to reflect the changes before the district can utilize them.

Unlike a RDBMS system that relies on a structured schema to manage its data, a NoSQL database platform or unstructured data, does not require a bound schema and can allow the database to evolve as needed. Hadoop and Elasticsearch are what many will call a 'NoSQL' platform. Hadoop via the Cloudera distribution and Elasticsearch will be the other data managing platforms that will reside within the FDOT District 5 Big Data production platform. NoSQL platforms such as Hadoop and Elasticsearch aren't limited to schemas and are more flexible in regards to managing and handling ever evolving datasets. Data is not stored in tables and columns. Data is stored into what are known as JSON (JavaScript Object Notation) documents. JSON documents are enriched with the data needed for the dataset into one document. Unlike a RDBMS database, the records / documents within a NoSQL collection do not have to look or be structured alike. The concept of joining data across various tables into one dataset does not exist with a NoSQL platform.

The ability to scale a data managing platform is vital in a big data environment. Both SQL and NoSQL platforms are scalable, but each platform scales in different ways. Organizations will 'Scale-Up' the SQL server to add additional computing resource, but eventually, the server will hit its max configuration. Once a server has hit its max resource capability, the district will have to add an additional SQL server to the SQL Server Cluster to integrate additional resources. The common challenge with this approach is the cost of expansion. A highly upgradable server as well as the associated SQL Server licensing cost can be cost prohibiting. The volume of data coming into the district's platform can quickly outpace the budget and ability to scale the SQL environment to keep up with the data demands needed. NoSQL env

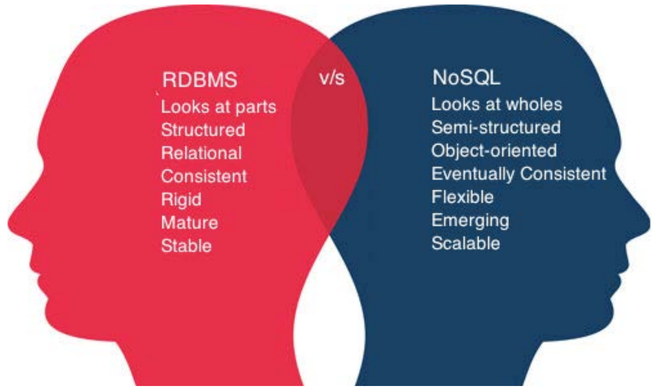

Figure 3 Source Image: Google Images

With NoSQL platforms such as Hadoop and Elasticsearch, the district will not worry about 'Scaling-Up' servers. Instead, the district will have the option to add commodity servers to the data cluster as needed to manage the current data demands needed. Commodity servers as well as the associated software are meant to be resources that can be added quickly to a big data environment without requiring significant investment to enable a system expansion. The goal with a commodity server is that if one of the data nodes goes offline, the district should be able to quickly replace the offline server with a new one and with minimal setup and having the new resource apart of the big data production platform. This scenario illustrates well the best practice of having commodity servers as physical hardware servers. The ability to quickly add and remove hardware from the big data environment will be a lot easier to manage with physical hardware servers compared to virtualized servers.

Operating Systems

The FDOT District 5 Big Data production platform will run in an environment that will support a mix of technologies such as SQL Server, Hadoop via the Cloudera distribution, Elasticsearch, and a variety of operating systems. The district's big data platform will consist of two operating systems; Windows Server and Ubuntu Server (Linux). Windows Server will serve as the operating system for the servers that are outside of the big data node servers. The Web / API server, SQL Server, pre and post processing servers and the ArcGIS servers will all sit on servers running the Windows Server operating system. Utilizing Windows Server for these various servers allows the district to reuse many of their software licenses and dev op procedures. Ubuntu Server will be the operating system that will reside on the district's data nodes within the big data platform. Ubuntu Server is a distribution of Linux. Ubuntu's light footprint, mature server technology, and free cost makes Ubuntu an excellent operating system for the commodity servers that will run the data nodes in the districts big data platform.

Data Storage Capacity

The FDOT District 5 Big Data production platform is designed to have a scalable architecture. Big Data is often defined through the V's of Big Data as volume, velocity, variety and others depending on the definition. It is a crucial to understand the big data platform's capacity, volume needs, rate of growth and future expansion requirements. A key aspect to the success of a big data store implementation is the ability to adapt and grow. The nature of big data is that as additional information is introduced, collected, and assimilated adequate space must be available in order the system to continue running at a high level of efficiency.

As data in the big data platform is designed to be continuously collected, stored, and analyzed and the platform architecture must account for the rate growth of the system. This includes not only the key component of storage space availability but also its future expansion. A big data store must have enough system resources in place for incoming information to be collected, stored, and analyzed while not compromising speed or efficiency of data and requests that are being accessed through data store. It is important to not only understand the rate of growth of the incoming datasets but to properly account for storage requirements of the big data cluster environment. Big data storage utilizes processes such as "*sharding*" or distributing multiple copies of the dataset across a cluster environment. It is this procedure which enables the big data store ability to interact with these massive datasets and quickly return results. Adequate storage must be available for these procedures to function properly. Using these procedures can establish the rate of growth that will be required by the big data store for future expansion, allowing data system architects the opportunity to plan accordingly based on the true growth needs of the system.
(As an example: The Seminole County Public Works department has sensor data that grows at a rate of about 7GB a day (this number will increase as they bring more sensors online and collecting data). At this rate, just with Seminole County Public Works, the district can potentially project the growth of just this data set to be around 3 (three) TB of data storage needed a year.)

Big data platforms require a lot of space. Each data node within the big data platform should be configured to what is known as JBOD configuration. JBOD by definition means 'Just A Bunch of Disks or Drives'. This array of disks will not be configured as an array of redundant drives (That configuration is known as RAID configuration). The big data platform and its data nodes are setup to be scalable. JBOD configuration allows the district to add additional drives to data nodes as needed with little configuration needed. The various data nodes within the big data platform do not need to mimic each other in disk size and / or number of drives.

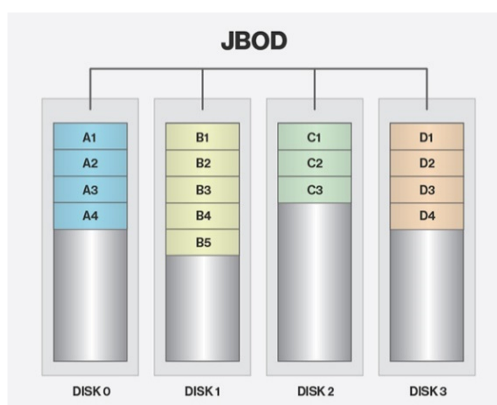

Utilizing SAN storage is another storage method that should be avoided. SAN storage can be cost prohibitive as the big data platform grows. SAN storage also utilize a single resource methodology. This means that the district will have a single pool (or two) of SAN storage that can be divided amongst various systems and servers. This approach will negate the purpose and one of the leading features of a big data platform, which is horizontal scalability. The ability to add, update and replace node servers and their drives is vital in a big data platform.

Following the architecture diagrams included in this document, the FDOT District 5 Big Data production platform was designed with the consideration to grow horizontally and scale as required. The flexibility of this architecture design enables future expansion to be seamless and additional resources to be applied without requiring massive reconfigurations of the entire big data platform and process.

Network Speed

Requests and interactions from users, applications, and other machines rely on the speed across the network to deliver the user experience in the time desired. The ability to handle the velocity of both incoming data, data requests and the returns from within the confines of the big data platform rely heavily on the access speed of the surrounding network. Maximizing the efficiency of the big data platform without providing adequate connection speed will result in a powerful big store that is heavily stifled by network traffic, a classic weakest link scenario.

The district's big data platform will need to transfer large datasets across its network. The districts network's ability to transmit and network speed will be vital in keeping the platform's performance optimal. Best practices show that investments in 10GbE, 40Gbe and 100Gbe networks will play a vital role in the success of the district's big data platform.
"Many organizations are in the early stages of running their [1 GbE or 10 GbE] networks, and even if they know they will need a certain amount of bandwidth and certain latency requirements, they are not ready to expend a ton of money on them. But businesses that are serious about Big Data and are investing heavily in 10 GbE may want to bypass 40 and go directly to 100. It depends on whether you want to be in the vanguard."
(Source: http://www.enterprisenetworkingplanet.com/datacenter/network-infrastructure-concerns-for-big-data.html)

Best practices show that starting with at least a 10GbE network will give the district a great starting point and then be able to evaluate if and what upgrades will be needed as data sets grows and data consumption needs increase.

Assumptions - FDOT Roles and Responsibility

Even though VHB is recommending the architecture for the FDOT District 5 Big Data production platform it is important to note that the proposed platform must live in an IT environment where the regular maintenance procedures of back up policies, data disaster recovery software and practices and disk management, just to name a few are established and will be extended and enhanced to absorb the new big data environment. It is the VHB assumption that the maintenance of physical and logical infrastructure will be the responsibility and handled by the district.

Managing a big data environment will require various skillsets that the district will need to account for. Skillsets such as network and server security, Windows Server administration, Linux / Ubuntu administration, Hadoop, Cloudera, and Elasticsearch administration will be required to maintain the operation of the district's big data platform.

FDOT D5 Production Big Data Platform Architecture

Platform Diagram:

The initial platform configuration shown below serves as a starting point for the districts big data platform. As datasets are added and dataset sizes increase, additional resources will be needed to allow the big data platform to grow and scale as needed. The diagram and server component sections below will show the architecture of the layout of the big data platform as well as server level information concerning the software, storage and ram requirements.

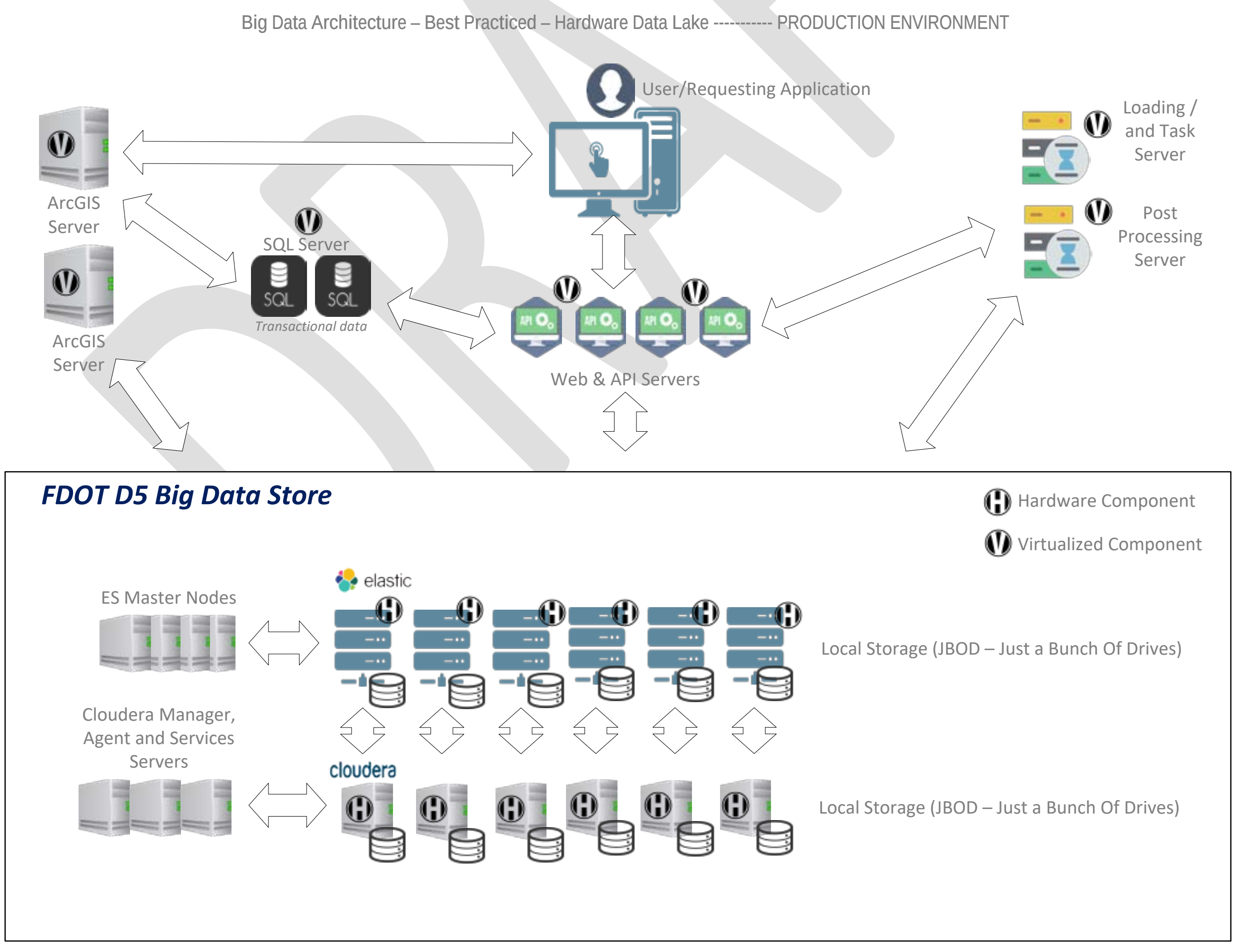

Platform Components and Functions:

ArcGIS Servers:

Function: Handle GIS Processing and Tasks Within the Big Data Platform.

Software Required:	• Windows Server 2012 R2 64bit - Load Balance Servers • ArcGIS Server • ArcGIS Desktop (Advanced License) • Geo-Processing Tools for Hadoop • ESRI Software
Local Hard Disk Size:	• 1 TB
Ram Allocation:	• 32gb minimal

SQL SERVER:

Function: To store and manage relational data that doesn't reside in the Big Data platform.

Software Required:	• SQL Server 2012 • Windows Server 2012 R2 64bit • Latest .NET Frameworks and Windows Updates
Local Hard Disk Size:	• 1 – 5tb
Ram Allocation:	• 32gb minimal

WEB / API SERVER (PER SERVER):

Function: To house and host web applications and API's that the Big Data Platform may utilize.

4 SERVERS – LOAD BALANCE ON SEPARATE BLADES

Software Required:	• Windows Server 2012 R2 64bit • Latest .NET Frameworks and Windows Updates • NODEJS • Python 2 • Python 3 • JAVA (Latest JRE) • Scala • GIT
Local Hard Disk Size:	• 2 TB per server
Ram Allocation:	• 32gb minimal

LOADING, TASK AND POST PROCESSING SERVER:

Function: The loading, task and post processing servers handles automated tasks, loading, and post processing of data that goes in and out the Big Data platform. These servers also handle long running tasks and processing of data needed by other applications.

Software Required:	• Windows Server 2012 R2 64bit • Latest .NET Frameworks and Windows Updates • NODEJS • Python 2 • Python 3 • JAVA (Latest JRE) • Scala • GIT
Local Hard Disk Size:	• 2 TB
Ram Allocation:	• 32gb minimal

ELASTIC SEARCH NODE / SERVER (PER DATA NODE):

Function: Data Nodes are servers within the Big Data platform that stores and manages the data documents that reside within the data lake. These nodes are also used to retrieve data from queries quickly to send back to calling applications.

Software Required:	• Ubuntu Server 16.04+ • JAVA • Spatial Framework for Hadoop (GIS Tools for Hadoop) • ESRI Geometry API for Java • Elastic Search
Local Hard Disk Size:	• Needs to be local array of hard drive (JBOD) (NOTE:) Each Node Should Have a Storage Capacity of 30 - 50tb
Ram Allocation:	• 32gb minimal, Ideally 64GB per server

ELASTIC SEARCH NODE / SERVER (MASTER NODES):

Function: Master Nodes are used to maintain, and keep track of all the data nodes and monitor its data and health of the cluster. It's important to have enough master nodes in your cluster so that the maintenance of your cluster is handled properly.

Software Required:	• Ubuntu Server 16.04+ • JAVA • Elasticsearch
Local Hard Disk Size:	• Each Node Should Have a Storage Capacity of 30 - 50tb - JBOD
Ram Allocation:	• at least 32 Gb per server ideally 64 Gb per server

(*Elastic has a formula to dictate how many Master Nodes a cluster should have. For 'N' = Number of Data Nodes in Cluster. **Master Nodes Formula:** (Number of Master Nodes 'MN' = ('N' / 2) + 1))

CLOUDERA NODE / SERVER (PER DATA NODE):

Function: Data Nodes are servers within the Big Data platform that stores and manages the data documents that reside within the data lake. These nodes are also used to retrieve data from queries quickly to send back to calling applications.

Software Required:	• Ubuntu Server 16.04+ • JAVA • Cloudera Platform • Additional tools will be installed after FDOT Procures Cloudera licenses with Cloudera's team
Local Hard Disk Size:	• Each Node Should Have a Storage Capacity of 30 - 50tb - JBOD
Ram Allocation:	• at least 32 Gb per server ideally 64 Gb per server

CLOUDERA MANAGER, AGENT, AND SERVICES SERVERS

Function: These servers exist to manager, examine and maintain the health of the Cloudera cluster. Many of the services that Cloudera utilizes will exist on these servers and are utilized when needed by the cluster.

Software Required:	• Ubuntu Server 16.04+ • JAVA • Cloudera Platform • Additional tools will be installed after FDOT Procures Cloudera licenses with Cloudera's team
Local Hard Disk Size:	• 1 – 5tb
Ram Allocation:	• at least 32 Gb per server ideally 64 Gb per server