

Accelerating Apache Hadoop with Pluribus Netvisor[®]

Big Data Processing

Hadoop Needs Optimized Networking

The business value hidden in the large quantity of data that organizations produce daily and accumulate over time is becoming more and more evident. As a consequence, Big Data processing using infrastructure such as Apache Hadoop is becoming a mainstream activity in the typical enterprise data center. In most cases, the execution time of a Hadoop job is directly related to the business outcome, and Big Data processing is expected to happen in real-time or quasi real-time.

Unfortunately, the Hadoop job results are available only when all its numerous tasks are completed, making Hadoop very sensitive to any unbalance in parallel task execution. This places unique demands on the network infrastructure, which must offer visibility into finely grained performance parameters and an almost surgical degree of control.

Plus, overall execution time of a Hadoop job depends on multiple phases, each that place huge demand on the network. These include the need to exchange, monitor and control huge flows of data for mapping, reduction, shuffling and ultimately output processing. Data ingestion and distribution also place significant load on the network.

Pluribus Netvisor[®] Operating System

To support the need to exchange, monitor and control huge flows of data, Pluribus has developed the Netvisor[®] OS (operating system) to run on bare metal switches. Without the need for an external controller, the Pluribus switches running Netvisor federate into a fabric, offering unprecedented insight, agility and security. Pluribus Netvisor switches enhance the deployment of Apache Hadoop above and beyond traditional network operating systems.

Pluribus Networks advances network virtualization and software-defined networking (SDN) through Netvisor, the industry's most programmable, open source-based network operating system. Netvisor is based upon the Pluribus Virtualization-Centric Fabric (VCF[™]) architecture, a proven approach to understanding flow, rapidly responding to business needs and securing your data.

Pluribus Netvisor combines the benefits of switch clustering for networking with a controller-less fabric. A traditional CLI

(Command Line Interface) is paired with fabric-wide programmability (C, RESTful API) and DevOps tools (e.g. Ansible) for agility and automation via a single point of management. Granular visibility and control is through a fabric-wide directory that contains endpoint information (vPorts) as well as allows for granular flow filtering and control (vFlow).

In combination with the Pluribus switches, Netvisor provides best-in-class switching economics. The deployment flexibility is guaranteed by Netvisor full L2/L3 stack providing complete interoperability with legacy networking infrastructure for easy insertion into brownfield deployments.

Factors that affect Hadoop processing speed:

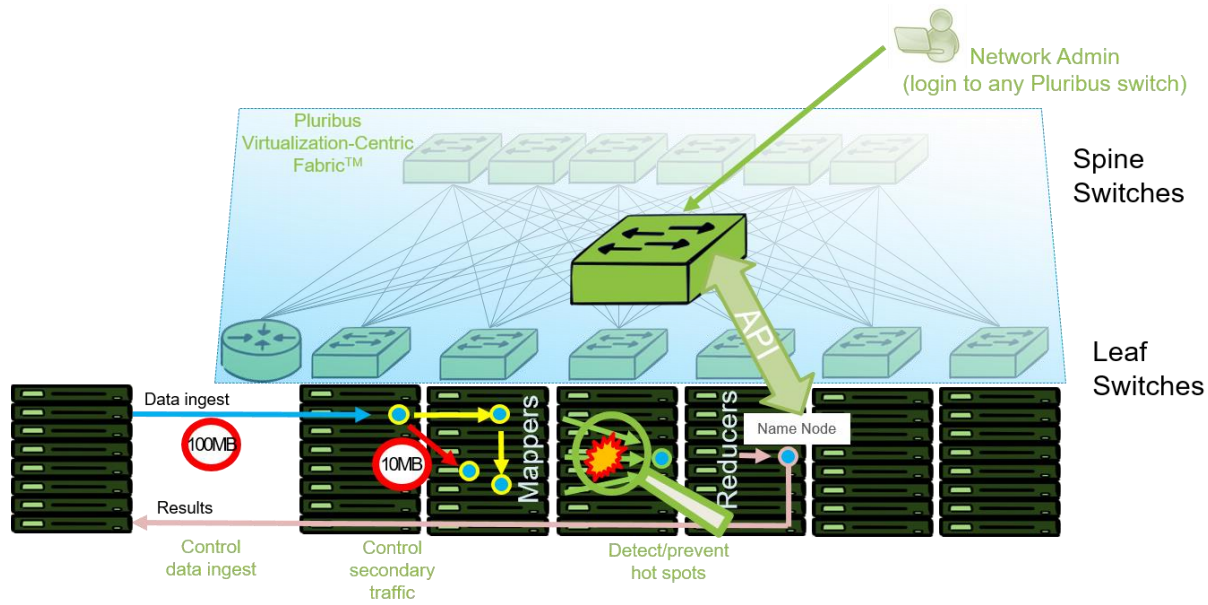
- Data ingestion
- Mapper data access
- Intermediate result shuffling from mappers to reducers
- Reducers data access
- Results output
- HDFS initial replication
- HDFS replica recovery

Benefits

- Simplified joint performance analysis and troubleshooting with vPort data (e.g. node names, data, roles, etc.).
- Improved HDFS primary and secondary traffic control via intelligent bandwidth allocation via vFlow commands.
- Troubleshoot HDFS replica recovery data conflicts with Mapper/Reducer data access with detailed telemetry.
- Identification of congestion/hot spots during results shuffling (especially from many mappers to few reducers).
- Spotting of excessive traffic between Task Trackers and Data Nodes due to suboptimal data locality.

Features

- Feature-rich L2/L3 and multicast to enable flexible deployment options.
- Pluribus switches can join into a distributed controller fabric and be managed as a single switch via CLI and/or API.
- Integrated tap-less visibility for data capture + post-analysis.
- vPort table, a fabric-wide endpoint "directory" accessible from any node for comprehensive endpoint and VM lifecycle tracking across the fabric.
- vFlow for granular visibility and control of every flow across the fabric.



Apache Hadoop deployment with Pluribus Netvisor

Hadoop Aspect	Challenge	Pluribus Solution
Time is Money	A job is as fast as the slowest task Many jobs in parallel can delay each other	Visibility into individual jobs/tasks performances
Big data	Massive amounts of data in/out, elephant flows are the norm	Bandwidth allocation to elephant flow to keep the jobs efficient (vFlow)
Map/reduce algorithm	Data ingest and result shuffling create time sensitive east/west traffic	Fabric-wide congestion monitor application
Distributed file system housekeeping	Large secondary traffic (replicas) contend for resources with jobs	Monitor and control individual client-server connections for data replication
Data locality	Stale topology information can affect task placement and keep a job waiting for last task	Database to access endpoints information from a fabric-wide API (vPort)

About Pluribus Networks

Pluribus Networks provides data center solutions that allow your business to run unconstrained. Our software-defined, open networking, fabric-based solutions transform existing network infrastructures into flexible and strategic assets fully aligned with today's digital business needs. Our Virtualization-Centric Fabric (VCF™) architecture provides unprecedented insight, agility and security to customers seeking to simplify operations, run more cost effectively and bring new applications online faster.

Learn more at www.pluribusnetworks.com and [@pluribusnet](https://twitter.com/pluribusnet)