BOOST PERFORMANCE OF HADOOP WITH PROACTIVE MONITORING

What is Proactive Monitoring?

Proactive Monitoring is used by Devops, IT Admin Team to identify & address issues before they occur. Proactive Monitoring in Hadoop & Big Data environment covers workloads, ad-hoc jobs, services, resource management framework and cluster. With Proactive Monitoring, Devops & IT admin team can quickly understand application behavior, service performance, cluster utilization, network bottlenecks, host or daemon malfunctioning and can proactively apply action plans to eradicate cause to avoid potential problem.

Hadoop Data Lake - Proactive Monitoring

Proactive Monitoring in resource shared environments like Hadoop Data Lake requires deep & cumulative analysis of offered capabilities from underlying infrastructure, configured capacity of yarn, live availability of yarn containers at each node manager, utilization patterns of queues/LOB/Pools - all combined with resource utilization by submitted application on the execution engine (MapReduce, Tez, Spark, etc.).

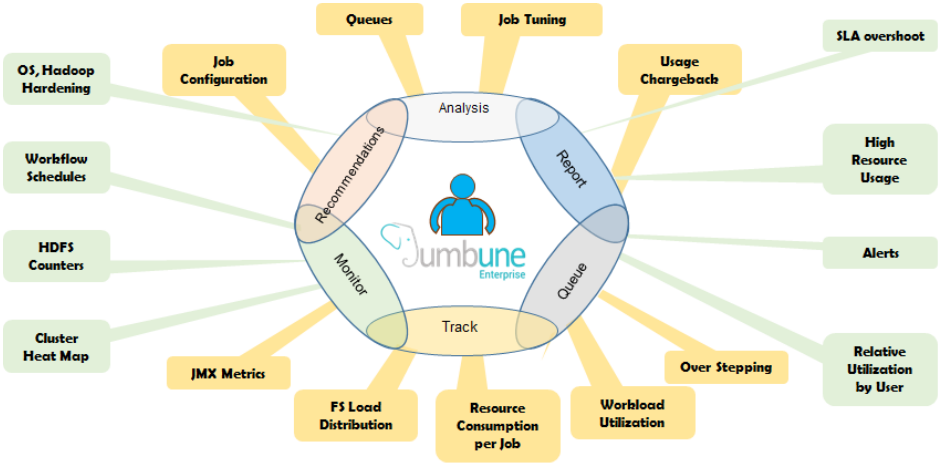

Jumbune – Cluster Monitoring & Analysis

Jumbune offers monitoring & analysis feature for Devops and IT Admin Teams whom are presented with a complementary dashboard which is derived from continuous deep analysis of each job, services and cluster. Every single execution is tracked thoroughly for its demanded capacity, resource utilization and underlying YARN based services like Schedulers, Queues and worker Managers to get precise recommendations and summaries.

Jumbune

Rapidly identify performance degradation cause

Hundreds of recommendations

Live capability & utilization analysis

Custom notification escalation

Queue/LOB deep analysis

http://jumbune.com

contact-at-jumbune.com

+1 408 252 7111

Impetus Technologies, Inc.

720 University Avenue, Suite 130
Los Gatos, CA 95032, USA

Highlighting features which embraces IT Teams for enabling proactive monitoring are,

- *Rapidly pinpoint root bottlenecks of workload/application performance issue*
- *Recommendations for cluster performance, workload, operations for each application, each node*
- *Real time offered capability of the underlying cluster*
- *Live utilization analysis for each job, queue/LOB, container*
- *Customize workflows of notification handling*

More than hundreds of recommendations are produced by Jumbune which are related to submitted workloads, configured capabilities of queues, maturing patterns of workload submissions on cluster, configured Yarn services & Hadoop parameters, operating system tuning, Service level agreements (SLA) overshoots, engine level optimization. Below is a list of high level categories of recommendations offered by Jumbune,

- *Optimal Workflow Schedules*
- *Optimal Configurations for Spark Jobs*
- *Resource over commitment/underutilization for analytical jobs*
- *Optimal configurations for cluster, node manager, queue capacities*
- *Offline Tuning for Hive, MapReduce, Pig Jobs*
- *Cost Based Optimization for MapReduce Jobs*

Jumbune analyzes resource utilization of scheduled Data Ingestion ETL workloads for use cases, scheduled MapReduce, Hive, PIG, Spark analytical workloads to derive and suggest optimal patterns which can assist in more uniform utilization and performance efficient execution of those workloads.

Upgrade maturity level of defined Hadoop LOBs

Moreover, Jumbune continuously observes and analyzes defined Line of Business (LOBs), workload queues for performing analysis of - how efficiently defined queues are being utilized by the job submissions happening on respective queues.

Another set of important analysis and observation presented to the IT Teams is - how much quickly the container allocation happening on each of the queue, which users are over utilizing the offered resources of the respective queue compared to other submissive users.

Optimal configurations for workloads

Another prominent offered feature in the space of proactive monitoring is live configuration recommendations for each executed MapReduce, Hive, PIG jobs. The recommendations are helpful to proactively suggest optimal configurations to the use case teams so that every on-boarded job perform efficient and don't clog the offered capabilities of the cluster by over utilizing them.

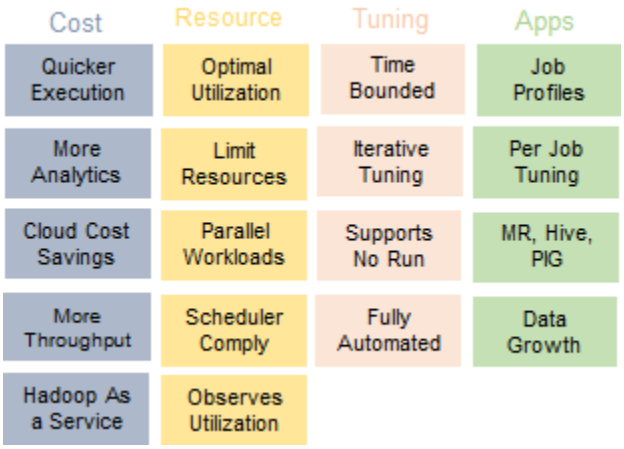

Customizable escalations of notifications

All above deep analysis of performance degradation factors, inefficient job submissions, mistuned cluster/queues/services can lead to severe operational overhead and may lead to serious business loss. Jumbune offers a customizable notification & alerting feature which can be easily plugged in email notifications, SNMP traps, enterprise ticketing system. The whole notification escalation handling is configurable to route alerts occurring since longer time to any higher escalated delivery path.

Jumbune is completely interoperable with your existing enterprise Hadoop distribution running in Cloud, On Premise or Hybrid Infrastructure. All features offered by Jumbune works complementary with distribution based existing monitoring tools and promises to deliver extended capabilities.