

Tuning the Hardware and Hadoop Cluster for Informatica® Big Data Products

Abstract

You can tune the hardware and the Hadoop cluster for better performance of Informatica big data products. This article provides tuning recommendations for Hadoop administrators and system administrators who set up the Hadoop cluster and hardware for Informatica big data products.

Supported Versions

- Big Data Management® 10.2 - 10.2.1
- Enterprise Data Catalog 10.2 - 10.2.1
- Enterprise Data Lake 10.2 - 10.2.1
- Big Data Streaming 10.2 - 10.2.1

Table of Contents

Overview.	2
Tune the Hardware.	2
CPU Frequency.	2
NIC Card Ring Buffer Size.	3
Tune the Hadoop Cluster.	5
Hard Disk Recommendations.	5
Transparent Huge Page.	5
HDFS Block Size Recommendations.	5
HDFS Access Timeout.	5
YARN Settings Recommendations for Parallel Jobs.	6
Documentation Reference.	7

Overview

Tuning the hardware and the Hadoop cluster are crucial for optimal performance of Informatica big data products.

Tune the Hardware

You can tune the following hardware parameters to optimize performance:

- CPU frequency
- NIC card ring buffer size

CPU Frequency

Dynamic frequency scaling allows the processor's frequency to be adjusted dynamically either for power savings or to reduce heat. Ensure that the CPU operates at least at the base frequency.

When CPUs run below the base frequency, the performance degrades by 30% to 40%.

Tuning Tip: Informatica recommends that you work with your IT system administrator to ensure that all the nodes on the cluster are configured to run at least at their supported base frequency.

To tune the CPU frequency for Intel multi-core processors, perform the following steps:

1. Run the `lscpu` command to determine the current CPU frequency, base CPU frequency, and the maximum CPU frequency that is supported for the processor.
2. Ask your system administrator to perform the following tasks:
 - a. Increase the CPU frequency at least to the supported base frequency.
 - b. Change the power management setting to **OS Control** at the BIOS level.
3. Run CPU-intensive tests to monitor the CPU frequency in real time and adjust the frequency for improved performance. On Red Hat Enterprise Linux operating systems, you can install a monitoring tool such as `cpupower`.
4. Work with your IT department to ensure that the CPU frequency and power management settings are persisted even in case of future system reboots.

NIC Card Ring Buffer Size

NIC configuration is a key factor in network performance tuning. When you process large volumes of data, you must tune the Receive (RX) and Transmit (TX) ring buffer size. The ring buffers contain descriptors or pointers to the socket kernel buffers that hold the packet data.

You can run the `ethtool` command to determine the current configuration. For example, run the following command:

```
# ethtool -g eth0
```

The following sections show a sample output:

```
Ring parameters for eth0:
Pre-set maximums:
RX: 2040
RX Mini: 0
RX Jumbo: 8160
TX: 255

Current hardware settings:
RX: 255
RX Mini: 0
RX Jumbo: 0
TX: 255
```

The **Pre-set maximums** section shows the maximum values that you can set for each parameter. The **Current hardware settings** section shows the current configuration details.

A low buffer size leads to low latency. However, low latency comes at the cost of throughput. For greater throughputs, you must configure large buffer ring sizes for RX and TX.

Informatica recommends that you use the `ethtool` command to determine the current hardware settings and the maximum supported values. Then, set the values based on the maximum values that are supported for each operating system. For example, if the maximum supported value for RX is 2040, you can use the `ethtool` command as follows to set the RX value to 2040:

```
# ethtool -G eth0 RX 2040
```

If you set a low ring buffer size for data transfer, packets might get dropped. To find out if packets were dropped, you can use the `netstat` and `ifconfig` commands.

The following image shows a sample output of the netstat command:

```
$ netstat -ip
Kernel Interface table
Iface      MTU Met  RX-OK RX-ERR RX-DRP RX-OVR    TX-OK TX-ERR TX-DRP TX-OVR Flg
em1        1500  0 65054784697 13232  42154  0 10044 59059281218  0  0  0 0 BMRU
em2        1500  0  0  0  0  0  0  0  0  0  0 0 BMU
em3        1500  0  0  0  0  0  0  0  0  0  0 0 BMU
em4        1500  0  0  0  0  0  0  0  0  0  0 0 BMU
lo         65536  0 9502357011  0  0  0  0 9502357011  0  0  0 LRU
```

The RX-DRP column indicates the number of packets dropped. To make sure that the RX-DRP column shows the value as 0, set the RX value accordingly.

The following image shows a sample output of the ifconfig command:

```
$ ifconfig
em1      Link encap:Ethernet  HWaddr 14:18:77:37:5C:A5
         inet addr:10.1.43.8  Bcast:10.1.43.255  Mask:255.255.252.0
         inet6 addr: fe80::1618:77ff:fe37:5ca5/64 Scope:Link
         UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
         RX packets:65054791888 errors:13232 dropped:0 overruns:10044 frame:3188
         TX packets:59059286151 errors:0 dropped:0 overruns:0 carrier:0
         collisions:0 txqueuelen:1000
         RX bytes:84969007070942 (77.2 TiB)  TX bytes:83254421015805 (75.7 TiB)
         Interrupt:40 Memory:95000000-957fffff

em2      Link encap:Ethernet  HWaddr 14:18:77:37:5C:A7
         UP BROADCAST MULTICAST  MTU:1500  Metric:1
         RX packets:0 errors:0 dropped:42154 overruns:0 frame:0
         TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
         collisions:0 txqueuelen:1000
         RX bytes:0 (0.0 b)  TX bytes:0 (0.0 b)
         Interrupt:44 Memory:94000000-947fffff

em3      Link encap:Ethernet  HWaddr 14:18:77:37:5C:A9
         UP BROADCAST MULTICAST  MTU:1500  Metric:1
         RX packets:0 errors:0 dropped:39877 overruns:0 frame:0
         TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
         collisions:0 txqueuelen:1000
         RX bytes:0 (0.0 b)  TX bytes:0 (0.0 b)
         Interrupt:44 Memory:93000000-937fffff

em4      Link encap:Ethernet  HWaddr 14:18:77:37:5C:AB
         UP BROADCAST MULTICAST  MTU:1500  Metric:1
         RX packets:0 errors:0 dropped:0 overruns:0 frame:0
         TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
         collisions:0 txqueuelen:1000
         RX bytes:0 (0.0 b)  TX bytes:0 (0.0 b)
         Interrupt:45 Memory:92000000-927fffff

lo       Link encap:Local Loopback
         inet addr:127.0.0.1  Mask:255.0.0.0
         inet6 addr: ::1/128 Scope:Host
         UP LOOPBACK RUNNING  MTU:65536  Metric:1
         RX packets:9502357904 errors:0 dropped:0 overruns:0 frame:0
         TX packets:9502357904 errors:0 dropped:0 overruns:0 carrier:0
         collisions:0 txqueuelen:0
         RX bytes:76681251808094 (69.7 TiB)  TX bytes:76681251808094 (69.7 TiB)
```

The status messages indicate the number of packets that were dropped.

Tune the Hadoop Cluster

You can tune the following Hadoop cluster level areas to optimize performance:

- Hard disk
- Transparent huge page
- HDFS block size
- HDFS access timeout
- YARN settings for parallel jobs

Hard Disk Recommendations

Hadoop workloads are always composite where there is a demand for multiple resources like CPU, Memory, Disk IO and Network IO. Disk performance plays a critical role in the overall Hadoop job's performance.

Consider the following factors to improve the performance of the Hadoop job:

- Use EXT4 or XFS file systems for the directories used for the cluster.
- Use SAS disks with 15K RPM for best performance.

Transparent Huge Page

Linux has transparent huge page compaction that impacts the performance of Hadoop workloads. Informatica recommends that you disable transparent huge page compaction.

For more information about disabling the transparent huge page compaction feature, see [KB article: 147609](#).

HDFS Block Size Recommendations

Set the HDFS block size based on your requirements.

The `dfs.block.size` parameter is the file system block size parameter for the data stored in the `hdfs-site.xml` file. The default block size is 128 MB. An increase or decrease in block size impacts parallelism and resource contention when you run MapReduce tasks. You can set the block size to 256 MB on a medium sized cluster with up to 40 nodes and a smaller value for a larger cluster. Tune the `dfs.block.size` value after experimenting on the basis of your requirements.

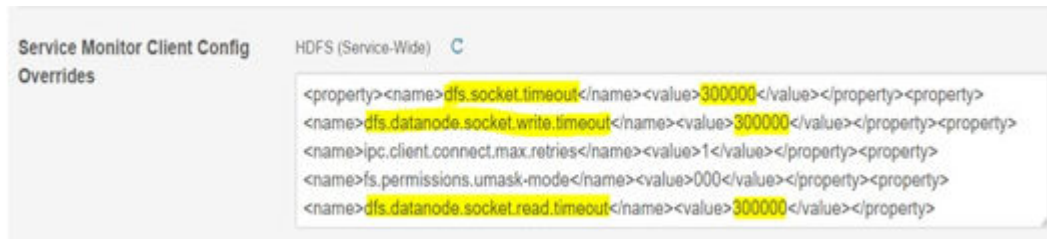
HDFS Access Timeout

You can increase the HDFS access timeout properties on the Hadoop cluster.

When processing a large data volume from HDFS, the mapping execution might fail due to access timeout errors. To avoid access timeout errors, the Hadoop administrator must set the following HDFS timeout properties on the Hadoop cluster to sufficiently large values:

- `dfs.client.socket.timeout`
- `dfs.datanode.socket.write.timeout`
- `dfs.datanode.socket.read.timeout`

The following image shows the sample values for the HDFS timeout properties:



YARN Settings Recommendations for Parallel Jobs

You can change or tune the Yet-Another-Resource-Negotiator (YARN) settings to improve the performance of Informatica big data products.

YARN's functions include splitting the operations of resource management and job scheduling/monitoring into separate services.

The number of containers that YARN node manager can run depends on the memory size, number of CPU cores, number of physical disks, and type of tasks. Avoid letting the number of parallel containers go beyond the minimum of four times the number of physical disks and the number of physical cores.

You can change or tune the following parameters to make sure that the Hadoop node can allocate that many parallel containers:

Parameter	Description
yarn.nodemanager.resource.memory-mb	The amount of physical memory in MB allocated for containers. Informatica recommends that you reserve some memory for other processes running on a node.
yarn.nodemanager.resource.cpu-vcores	The number of CPU cores allocated for containers. Informatica recommends that you set this value to the number of physical cores available on the node.
yarn.nodemanager.vmem-check-enabled	The virtual memory check is set to <code>false</code> by default. Retain the default value.
yarn.scheduler.minimum-allocation-mb	The minimum allocation, in MBs, for every container request at the resource manager. Memory requests lower than this will not take effect, and the specified value will get allocated at minimum. Set this value to 1024.
yarn.scheduler.minimum-allocation-vcores	The minimum allocation for every container request at the resource manager in terms of virtual CPU cores. Requests lower than this will not take effect, and the specified value will get allocated the minimum. Set this value to 1.

Note: Consult your Hadoop administrator before you change these settings. These recommendations are based on internal tests and might differ from the Hadoop vendor's recommendations.

Documentation Reference

The following table lists performance-related How-To Library articles for Informatica big data products:

Article	Description
Informatica big data products	
Tuning the Hardware and Hadoop Clusters for Informatica Big Data Products	Provides tuning recommendations for the hardware and the Hadoop cluster for better performance of Informatica big data products.
Big Data Management	
Performance Tuning and Sizing Guidelines for Big Data Management 10.2	Provides sizing recommendations for the Hadoop cluster and the Informatica domain, tuning recommendations for various Big Data Management components, best practices to design efficient mappings, troubleshooting tips, and case studies.
Tuning the Hive Engine for Big Data Management	Provides tuning recommendations to run mappings on the Hive engine, best practices to design efficient mappings, and case studies.
Strategies for Incremental Updates on Hive	Describes alternative solutions to the Update Strategy transformation for updating Hive tables to support incremental loads.
Intelligent Data Lake	
Performance Tuning and Sizing Guidelines for Intelligent Data Lake 10.2	Provides sizing recommendations and tuning guidelines for ingesting data, previewing data assets, adding data assets to projects, managing projects, publishing projects, searching for data assets, exporting data assets, and profiling data.
Intelligent Streaming	
Performance Tuning and Sizing Guidelines for Informatica Intelligent Streaming 10.2	Provides sizing recommendations and tuning guidelines for Informatica Intelligent Streaming.

Authors

Vishal Kamath

Anand Sridharan

Indra Sivakumar