vNode performance testing

Made by Fabian on 26/8/2020

Environment:

I've created 2 VMs with 8 cores and 16Gb RAM.

CPU: Intel Xeon CPU-d-1528 @ 1,90 GHz

Created 20k tags in the Remote.

Linked to the Central.

No Store&Forward.

At the Remote:

Connect timeout: 30s

reconnect delay: 30s

Keep Alive: 60s

At the Central:

Timeout: 45s

NOTE: Result for the Store&Forward at the Repsol project:

The recovery rate after a disconnection is 2.000 values/second

TEST 1

Reading from the Simulation Server every 10 seconds.

It works well with changes higher than 10seconds.

In the log can be seen that 20k tags need around 10 seconds to arrive.

About 2 values per ms. --> 20.000 values in 10.000 ms --> 2.000 values/second --> 120.000 values/minute

If the rate is changed to 20k tags every 8 seconds, the Central starts getting delayed.

After some time, the link breaks up and reconnects again to restart the cycle.

TEST 2

2 Remote vNodes with 20k tags each.

Reading from the Simulation Server every 60 seconds.

40.000 values/minute

Works OK.

TEST 3

2 Remote vNodes with 20k tags each.

Reading from the Simulation Server every 40 seconds.

Total: 30k values/minute x 2 = 60k values/minute

Works in the limit.

TEST 4

2 Remote vNodes with 20k tags each.

Reading from the Simulation Server every 30 seconds.

Total: 40k values/minute x 2 = 80k values/minute

Works in the limit.

Sometimes the Central is delayed for 2 minutes and sometimes the delay is just 30 seconds.

The delay does not increase.

TEST 5

2 Remote vNodes with 20k tags each.

Reading from the Simulation Server every 20 seconds.

Total: 60k values/minute x 2 = 120k values/minute

Works in the limit.

Sometimes the Central is delayed for 2 minutes and sometimes the delay is just 30 seconds.

The delay does not increase.

TEST 6

2 Remote vNodes with 20k tags each.

Reading from the Simulation Server every 15 seconds.

Total: 80k values/minute x 2 = 160k values/minute

The Central gest delayed for about 10 minutes until the delay is about 6 minutes.

Then the link breaks down.

TEST 7

2 Remote vNodes with 20k tags each.

Reading from the Simulation Server every 10 seconds.

Total: 120k values/minute x 2 = 240k values/minute

The link gest broken periodically.

TEST 8

2 Remote vNodes with 20k tags each.

Reading from the Simulation Server

5k tags every 10 seconds.

5k tags every 30 seconds.

10k tags every 60 seconds.

Total: 50k values/minute x 2 = 100k values/minute

The link gets broken every 2 minutes.

TEST 9

2 Remote vNodes with 20k tags each.

Reading from the Simulation Server

5k tags every 15 seconds.

5k tags every 30 seconds.

10k tags every 60 seconds.

Total: 40k values/minute x 2 = 80k values/minute

Works well.

TEST 10

2 Remote vNodes with 20k tags each.

Reading from the Simulation Server

2k tags every 10 seconds.

5k tags every 30 seconds.

13k tags every 60 seconds.

Total: 35k values/min x 2 = 70k values/minute

Works well.

TEST 11

2 Remote vNodes with 20k tags each.

Reading from the Simulation Server

2k tags every 5 seconds.

8k tags every 30 seconds.

10k tags every 60 seconds.

Total: 50k values/min x 2 = 100k values/minute

Works well.

TEST 12

2 Remote vNodes with 3k tags each.

Reading from the Simulation Server

3k tags every 1 seconds.

Total: 180k values/min x 2 = 360k values/minute

Data in Central gest delayed until link gets broken.

TEST 13

2 Remote vNodes with 3k tags each.

Reading from the Simulation Server

1,5k tags every 1 seconds.

1,5k tags every 4 seconds.

Total: 82.500 values/min x 2 = 165k values/minute

Data in Central gest delayed until link gets broken.

TEST 14

2 Remote vNodes with 3k tags each.

Reading from the Simulation Server

1500 tags every 1 seconds.

1500 tags every 5 seconds.

Total: 78k values/min x 2 = 156k values/minute

Data in Central gest delayed until link gets broken.

TEST 15

1 Remote vNode with 3k tags.

Reading from the Simulation Server

3000 tags every 1,153 seconds.

Total: 156k values/minute

Data in Central gest delayed at some moments.

Havent's seen the link broken.

Looks like is in the limit.

TEST 16

1 Remote vNode with 3k tags.

Reading from the Simulation Server

1500 tags every 1 seconds.

1500 tags every 2 seconds.

Total: 135k values/min

Works fine

TEST 17

2 Remote vNodes with 3k tags each.

Reading from the Simulation Server

1,5k tags every 2 seconds.

1,5k tags every 4 seconds.

Total: 67.500 values/min x 2 = 135k values/minute

Works fine

TEST 18

5 Remote vNodes with 5k tags each.

Reading from the Simulation Server every 15 seconds.

Total: 20.000 values/min x 5 = 100.000 values/min

Works fine

TEST 19

5 Remote vNodes with 5k tags each.

Reading from the Simulation Server every 10 seconds.

Total: 30.000 values/min x 5 = 150.000 values/min

Data in Central gest delayed until link gets broken.

TEST 20

5 Remote vNodes with 5k tags each.

Reading from the Simulation Server every 12 seconds.

Total: 25.000 values/min x 5 = 125.000 values/min

Data in Central gest delayed at some moments.

Havent's seen the link broken.

Looks like is in the limit.

TEST 21

Central vNode is in AWS → RAM 8Gb, CPU 2.80 GHz, 4 CPU Cores

5 Remote vNodes with 5k tags each.

Reading from the Simulation Server every 10 seconds.

Total: 30.000 values/min x 5 = 150.000 values/min

Works fine

TEST 22

Central vNode is in AWS → RAM 8Gb, CPU 2.80 GHz, 4 CPU Cores

6 Remote vNodes with 5k tags each.

Reading from the Simulation Server every 10 seconds.

Total: 30.000 values/min x 6 = 180.000 values/min

CPU load: 40% Works fine

TEST 23

Central vNode is in AWS → RAM 8Gb, CPU 2.80 GHz, 4 CPU Cores

7 Remote vNodes with 5k tags each.

Reading from the Simulation Server every 10 seconds.

Total: 30.000 values/min x 7 = 210.000 values/min

CPU load: 45% Works fine

TEST 24

Central vNode is in AWS → RAM 8Gb, CPU 2.80 GHz, 4 CPU Cores

8 Remote vNodes with 5k tags each.

Reading from the Simulation Server every 10 seconds.

Total: 30.000 values/min x 8 = 240.000 values/min

CPU load: 50%, Bandwith Receive: 12 Mbps

Works fine but it is in the limit.

TEST 25

Central vNode is in AWS → RAM 8Gb, CPU 2.80 GHz, 4 CPU Cores

1 Remote vNode with 5k tags each.

Reading from the Simulation Server every 2 seconds.

Total: 150.000 values/min

CPU load: 25%, Bandwith Receive: 4 Mbps

Works fine

TEST 26

Central vNode is in AWS → RAM 8Gb, CPU 2.80 GHz, 4 CPU Cores

1 Remote vNode with 5k tags each.

Reading from the Simulation Server every 1,5 seconds.

Total: 200.000 values/min

CPU load: 35%, Bandwith Receive: 6 Mbps

Works fine

TEST 27

Central vNode is in AWS → RAM 8Gb, CPU 2.80 GHz, 4 CPU Cores

1 Remote vNode with 5k tags each.

Reading from the Simulation Server every 1,25 seconds.

Total: 240.000 values/min

CPU load: 40%, Bandwith Receive: 6 Mbps

Data gets delayed until link gets broken.

Looks like the capacity for a Central to process data is:

- For a 1,9 GHz CPU → between 60.000 values/minute AND 135.000 values/minute depending on the amount of tags to process.

- For a 2,8 GHz CPU → Up to 240.000 values/minute depending on the amount of tags to process.