

Flume NG Performance Measurements

Flume NG Performance Measurements

Table of Contents

- Summary
- Tests
 - Syslog Performance Test 2012-04-30
 - Syslog Stress Test 2012-04-28

Summary

The goal of this page is to collect the latest performance metrics achieved by those in the Apache Flume (incubating) community when using Flume NG.

Test results should be provided in reverse chronological order (newest first) on included sub-pages. Please separate tests by horizontal rules.

Load generation and validation scripts are available on <https://github.com/mpercy/flume-load-gen>

Tests

Syslog Performance Test 2012-04-30

Who ran the test: Mike Percy <mpercy at cloudera dot com>

Test Setup

Overview

The Flume NG agent was run on its own physical machine in a single JVM. A separate client machine generated load against the Flume box in syslog format. Flume stored data onto a 9-node HDFS cluster configured on its own separate hardware. No virtual machines were used in this test.

Hardware specs

CPU: Intel Xeon L5630 2 x quad-core with Hyper-Threading @ 2133MHz (8 physical cores)

Memory: 48GB

OS: SLES 11sp1 (SuSE Linux 64-bit)

Flume configuration

Java version: 1.6.0u26 (Server Hotspot VM)

Java heap size: 2GB

Num. agents: 1

Num. parallel flows: varies (see results)

Source: SyslogTcpSource

Channel: MemoryChannel

Sink: HDFSEventSink with avro_event serialization and snappy serializer compression

Fragment of flume.conf config file

```
# Number of sources, channels, and sinks varied depending on tests.
# In each case, they are independent flows and therefore do not share threads, data,
or resources.
# This example only shows 3 flows. The # of flows were varied from 6 to 16.
agent.sources = svc_0_src svc_1_src svc_2_src
agent.channels = svc_0_chan svc_1_chan svc_2_chan
agent.sinks = svc_0_sink svc_1_sink svc_2_sink

# example of one flow is below, i.e. "flow 0"
agent.channels.svc_0_chan.type = memory
agent.channels.svc_0_chan.capacity = 100000
agent.channels.svc_0_chan.transactionCapacity = 1000

agent.sources.svc_0_src.type = org.apache.flume.source.SyslogTcpSource
agent.sources.svc_0_src.port = 10001
agent.sources.svc_0_src.channels = svc_0_chan

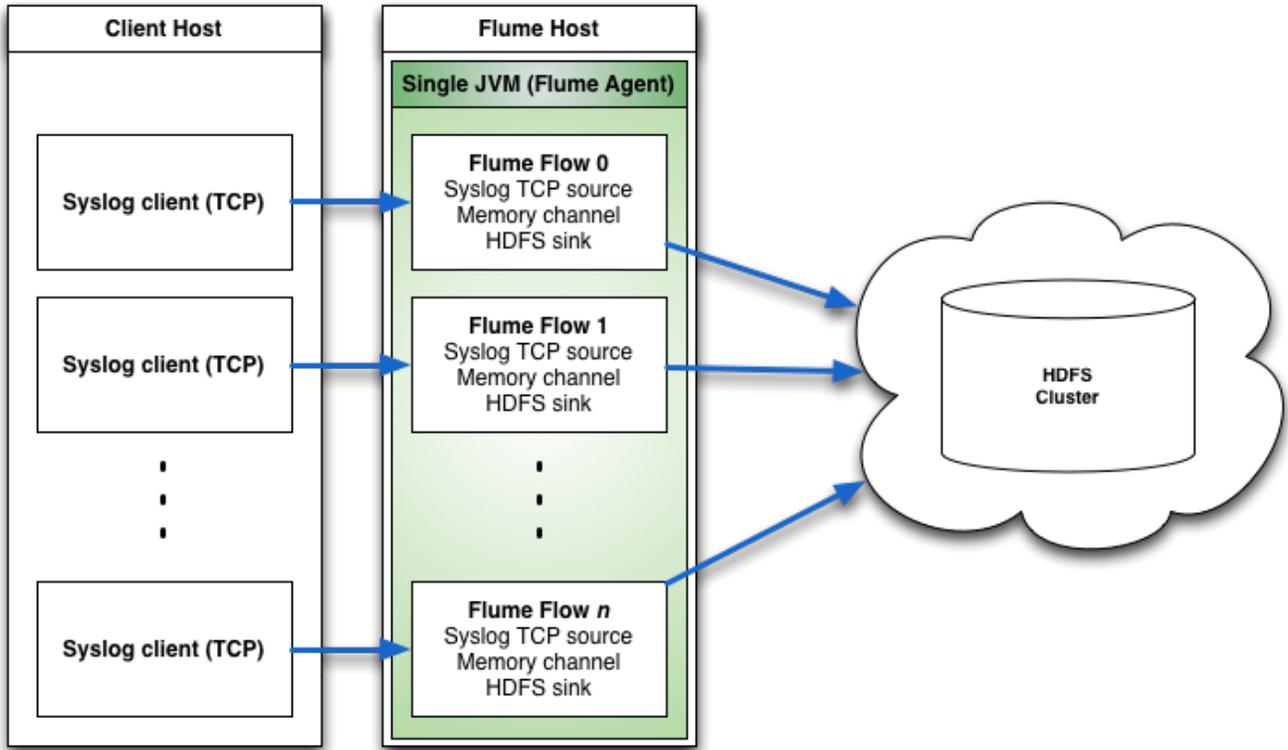
agent.sinks.svc_0_sink.type = hdfs
agent.sinks.svc_0_sink.hdfs.path = hdfs://xxxxxx.cloudera.com/service/20120430/flow0
agent.sinks.svc_0_sink.hdfs.fileType = DataStream
agent.sinks.svc_0_sink.hdfs.rollInterval = 300
agent.sinks.svc_0_sink.hdfs.rollSize = 0
agent.sinks.svc_0_sink.hdfs.rollCount = 0
agent.sinks.svc_0_sink.hdfs.batchSize = 1000
agent.sinks.svc_0_sink.hdfs.txnEventMax = 1000
agent.sinks.svc_0_sink.hdfs.kerberosPrincipal = flume/_HOST@CLOUDERA.COM
agent.sinks.svc_0_sink.hdfs.kerberosKeytab = /etc/flume-ng/conf/flume-xxxxxx.keytab
agent.sinks.svc_0_sink.serializer = avro_event
agent.sinks.svc_0_sink.serializer.compressionCodec = snappy
agent.sinks.svc_0_sink.channel = svc_0_chan

# ... define flow 1 ...
```

Hadoop configuration

The HDFS sink was connected to a 9-node Hadoop cluster running CDH3u3 with MIT Kerberos v5 security enabled.

Visualization of test setup



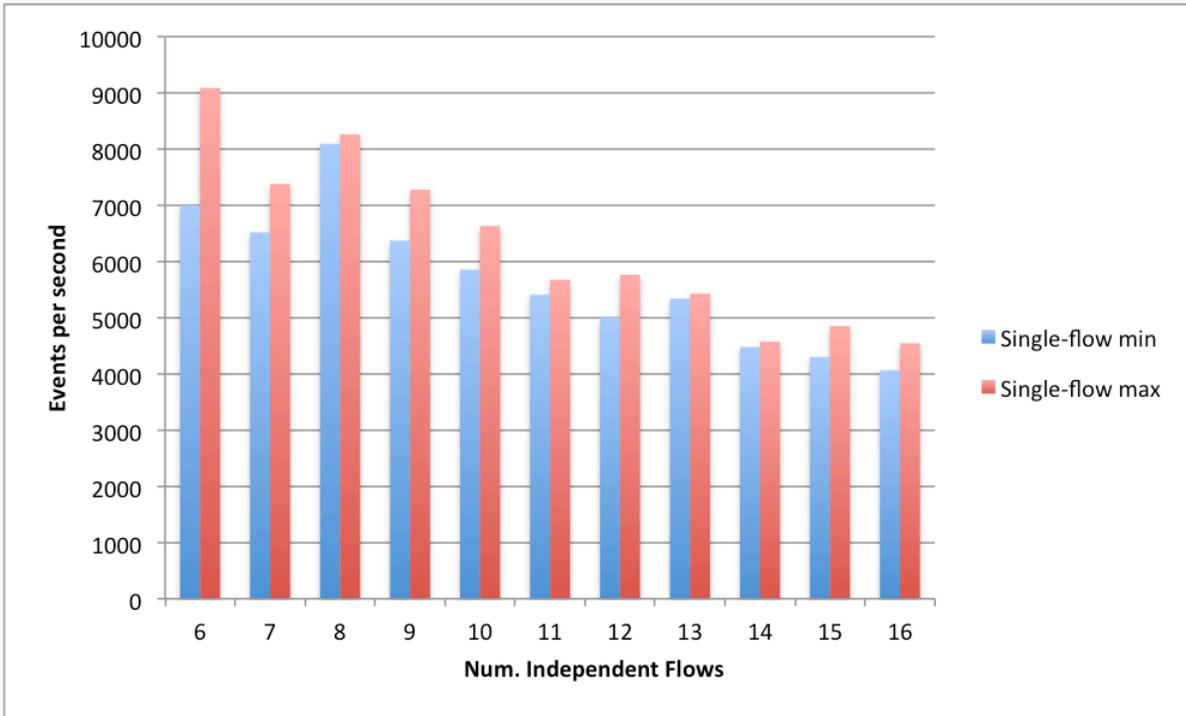
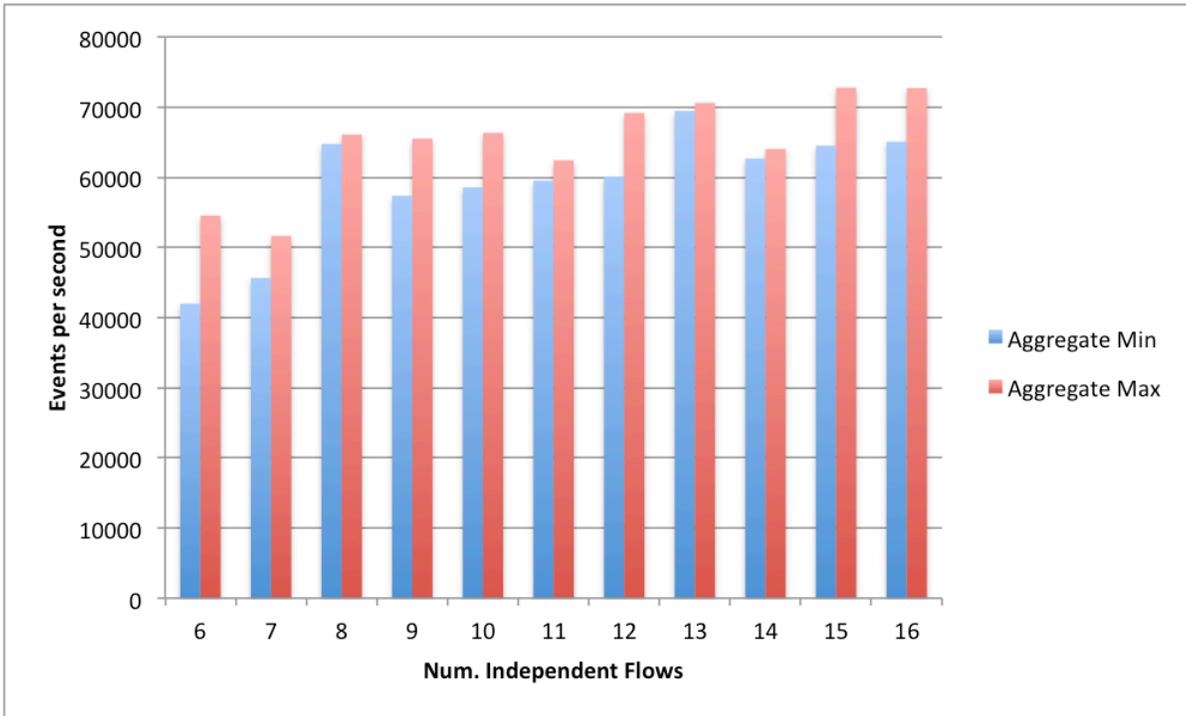
Data description

Syslog entries containing sequentially increasing integers plus padding
 Event size: 300 bytes

Results

Throughput summary

Num flows	Min aggregate events/sec	Max aggregate events/sec	Min avg. single-flow events/sec	Max avg. single-flow events/sec
6	41982.34	54538.92	6997.06	9089.82
7	45639.21	51646.33	6519.89	7378.05
8	64748.63	66095.53	8093.58	8261.94
9	57358.73	65506.95	6373.19	7278.55
10	58557.15	66324.04	5855.72	6632.40
11	59519.33	62419.89	5410.85	5674.54
12	60105.21	69164.94	5008.77	5763.74
13	69450.87	70590.71	5342.37	5430.05
14	62674.97	64030.08	4476.78	4573.58
15	64499.65	72783.06	4303.64	4852.20
16	65064.07	72714.94	4066.50	4544.68



Conclusions

1. Flume appears to be capable of achieving approx. 70,000 events/sec on a single machine at the time of the test with no data loss
2. The optimal number of parallel flows is nearly achieved by creating one flow per CPU core. Additional flows may be added with marginal benefit, likely up to 2x the number of physical cores available on the system, if hyper-threading is available.

N.B. These various performance tests were somewhat brief and not completely automated, so these min/max numbers were only across 3 runs per flow combination, and the min/max was a snapshot in time after a minute or two of allowing the traffic to "settle" (and Hotspot to kick in). Because of the low number of samples, some significant variance is to be expected. Nevertheless, some general trends are apparent in the data.

Who ran the test: Mike Percy <mpercy at cloudera dot com>

Test setup

Overview

The Flume NG agent was run on its own physical machine in a single JVM. A separate client machine generated load against the Flume box in syslog format. Flume stored data onto a 9-node HDFS cluster configured on its own separate hardware. No virtual machines were used in this test.

Hardware specs

CPU: Intel Xeon L5630 2 x quad-core with Hyper-Threading @ 2133MHz (8 physical cores)
Memory: 48GB
OS: SLES 11sp1 (SuSE Linux 64-bit)

Flume configuration

Java version: 1.6.0u26 (Server Hotspot VM)
Java heap size: 2GB
Num. agents: 1
Num. parallel flows: 10
Source: SyslogTcpSource
Channel: MemoryChannel
Sink: HDFSEventSink with avro_event serialization and snappy serializer compression

Single-flow config

```
agent.channels.svc_0_chan.type = memory
agent.channels.svc_0_chan.capacity = 100000
agent.channels.svc_0_chan.transactionCapacity = 1000

agent.sources.svc_0_src.type = org.apache.flume.source.SyslogTcpSource
agent.sources.svc_0_src.port = 10001
agent.sources.svc_0_src.channels = svc_0_chan

agent.sinks.svc_0_sink.type = hdfs
agent.sinks.svc_0_sink.hdfs.path = hdfs://xxxxxx.cloudera.com/service/20120428/flow0
agent.sinks.svc_0_sink.hdfs.fileType = DataStream
agent.sinks.svc_0_sink.hdfs.rollInterval = 300
agent.sinks.svc_0_sink.hdfs.rollSize = 0
agent.sinks.svc_0_sink.hdfs.rollCount = 0
agent.sinks.svc_0_sink.hdfs.batchSize = 1000
agent.sinks.svc_0_sink.hdfs.txnEventMax = 1000
agent.sinks.svc_0_sink.hdfs.kerberosPrincipal = flume/_HOST@CLOUDERA.COM
agent.sinks.svc_0_sink.hdfs.kerberosKeytab = /etc/flume-ng/conf/flume-xxxxxx.keytab
agent.sinks.svc_0_sink.serializer = avro_event
agent.sinks.svc_0_sink.serializer.compressionCodec = snappy
agent.sinks.svc_0_sink.channel = svc_0_chan
```

Hadoop configuration

The HDFS sink was connected to a 9-node Hadoop cluster running CDH3u3 with MIT Kerberos v5 security enabled.

Data description

Syslog entries containing sequentially increasing integers plus padding
Event size: 300 bytes

Results

Summary analysis

Load: 58,582 events/sec aggregate == approx. 5,850 events/sec per flow on average x 10 flows.
Event size: 300 bytes/event.
Duration: The load test ran for 23 hours and 20 minutes.
Result: Total events sent: 4,920,930,988; No lost events, only 7,000 duplicates (7 retried transactions).

Automated data integrity report

```
WARNING: xxxxxx.cloudera.com/hammer-0: dups: 1000
INFO: xxxxxx.cloudera.com/hammer-0: total=493252903 is OK
WARNING: xxxxxx.cloudera.com/hammer-1: dups: 1000
INFO: xxxxxx.cloudera.com/hammer-1: total=491792986 is OK
INFO: xxxxxx.cloudera.com/hammer-2: total=494361441 is OK
WARNING: xxxxxx.cloudera.com/hammer-3: dups: 1000
INFO: xxxxxx.cloudera.com/hammer-3: total=491668462 is OK
WARNING: xxxxxx.cloudera.com/hammer-4: dups: 1000
INFO: xxxxxx.cloudera.com/hammer-4: total=493222420 is OK
WARNING: xxxxxx.cloudera.com/hammer-5: dups: 1000
INFO: xxxxxx.cloudera.com/hammer-5: total=492295186 is OK
WARNING: xxxxxx.cloudera.com/hammer-6: dups: 1000
INFO: xxxxxx.cloudera.com/hammer-6: total=490704238 is OK
WARNING: xxxxxx.cloudera.com/hammer-7: dups: 1000
INFO: xxxxxx.cloudera.com/hammer-7: total=491209003 is OK
INFO: xxxxxx.cloudera.com/hammer-8: total=493007823 is OK
INFO: xxxxxx.cloudera.com/hammer-9: total=489416526 is OK
INFO: grand total=4920930988
```
