



INTERSEC BENCHMARK

High Performance for Fast Data & Real-Time
Analytics
Part II: Vs Spark



BENCHMARK VS SPARK ON REALTIME DECISIONING

Intersec solution in a Red Hat Openstack NFV framework proves 250% faster than Spark Streaming technology

The combination of Intersec solution with Red Hat Openstack NFV framework provides you with a flexible Fast Data platform that allows you to analyse data as it flows, and trigger automated rules when detecting specific patterns. This document presents the tests led by Red Hat and Intersec during the summer 2015, evaluating the performance of this combined solution with Spark Streaming technology, when applied in a context of customer base management.

I. INTERSEC INTEGRATION WITH RED HAT OPENSTACK NFV FRAMEWORK

Most service providers operate heterogenous networks with multiple, diverse data sources that prove challenging for legacy data management systems to analyse in realtime and trigger a specific action or rule when detecting a specific pattern. This is particularly the case in customer base management, where policy control or marketing strategy sometimes require an instant responsiveness based on customer's usage.

Intersec exclusive technology combined with Red Hat Openstack NFV framework provides an efficient realtime engine adapted to high volumes of streamed data, whatever their types, volumes and sources. This combined solution is ideal for high velocity businesses across all industries: telecommunications, but also financial services, public utilities, online retail or manufacturing who require instant responsiveness based on specific event detection.

We compared Intersec performance to Spark Streaming technology, one of the most promising solution for realtime streaming analytics. Even if other technologies (not based on MapReduce) may be more adapted to realtime triggering, Spark delivers among the the best performances in terms of streaming analytics, using inmemory processing and distributed architecture.

II SCOPE

The test was run over a customer base management simplified usecase, where a service provider triggers an offer as soon as a customer reaches a threshold of 50 SMS sent.

We compared Intersec solution and Spark Streaming in terms of maximum throughputs each configuration could ingest.



Data coming from a live production environment was simplified, keeping only a few anonymized fields with additional noise over dates and events. It amounted to around 45 million events in CSV files (4.3 GB), with the following fields:

- Timestamp
- Customer identifier
- Transaction type (Topup, call, SMS, data session)
- Volume (duration of a call, data volume, etc)
- Transaction price

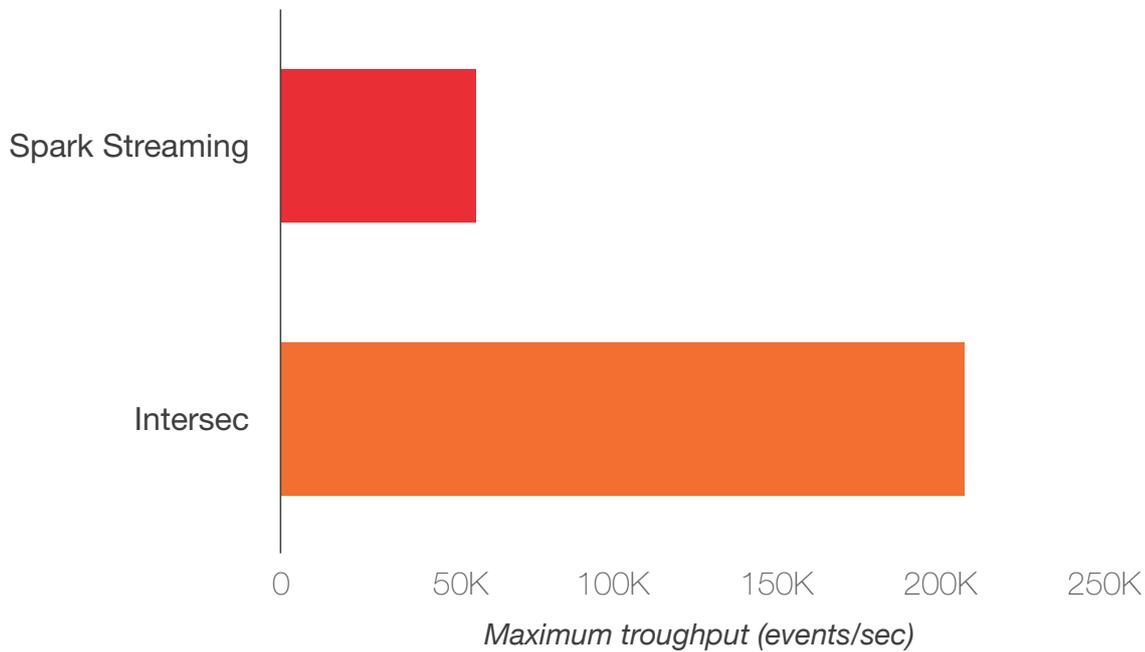
The storage was distributed over four different hosts.



The two solutions were deployed on four virtual machines, each with the following configuration:

- 10 vCPUs
- 60GB of RAM
- 500GB of disk storage.

III RESULTS



Maximum processing speed when running the scenario, measured in thousands events per second.

This performance may be explained by the fact that, even if both solutions use inmemory processing and distributed architecture, Intersec benefits from a realtime rule engine, more efficient than the microbatches run by Spark.

The advantage of Intersec technology is probably greater than the 250% measured increment, as we compared a simple development made on Spark with our complete product, including those Enterprise-grade features that have certainly affected its performance:

- Finegrain monitoring at all stages of the process
- Flexible configuration by the enduser to enable plentiful of use cases

- Users' right management
- Historization of all events for future analysis
- etc.

INTERSEC
SOLUTION
PROVED MORE
THAN
3.5x FASTER

IV CONCLUSION

Our tests showed Intersec technology implemented on Red Hat Openstack NFV framework provided a 250% improvement over Spark Streaming technology for a simple use case of customer base management. This higher level of responsiveness is critical in many cases where realtime decisioning is required, to ensure bestinclass customer experience, online control policy or fraud detection.