

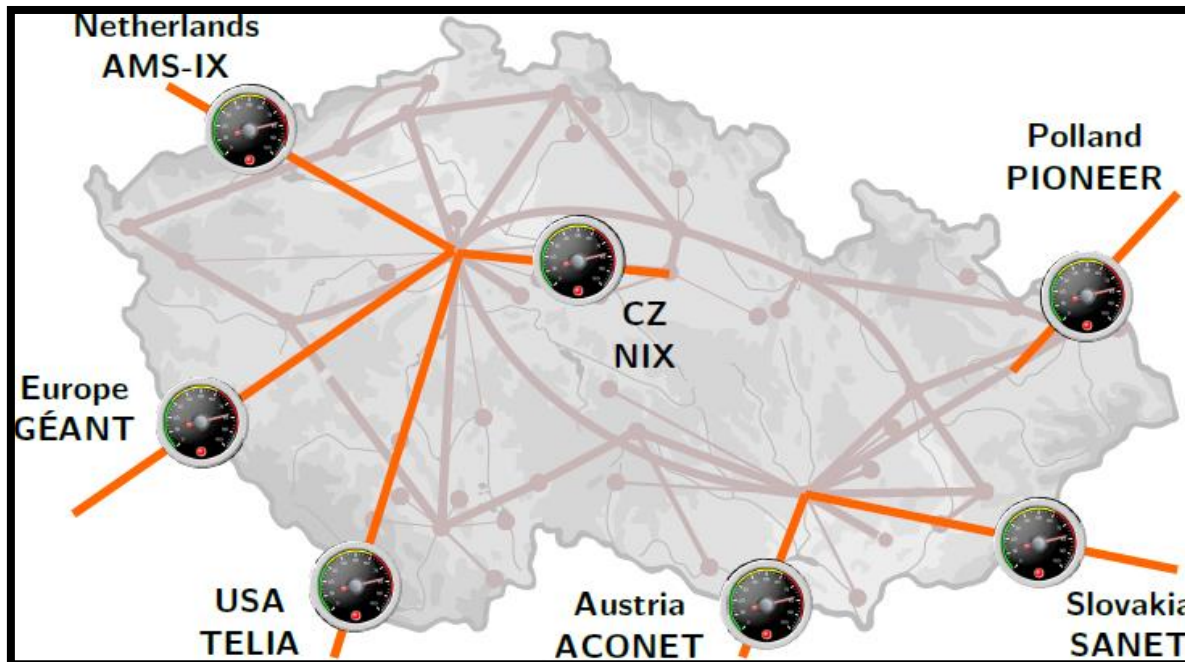


# Flow data storage and retrieval utilizing big data aproach

CESNET, INVEA-TECH, MU

# Motivation

- Network flow monitoring generates large amount of data – 250 GB per day
- Interactive work with data is an issue



# Intro

---

- There are several open-source platforms enabling big data processing
  - Hadoop, (native, Hive, Pig, nfdist) – MapReduce
  - ElasticSearch
  - Vertica
  - Proprietary implementation

# Queries

---

- Query 1: Total number of flows, packets, bytes
- Query 2: Number of flows with port 53 and proto TCP
- Query 3: Print flows with destination port 53.
- Query 4: Print IP addresses sorted by bytes with flows, packets and bytes

# Data

---

- One 10Gbps line
- 24 hours
- 877 million of flow records
- Records simplified to NetFlow v5 equivalent
- CSV as well as binary data representation

# Hadoop cluster

---

- 24 slave + 3 master nodes
- Intel Xeon CPU E5-2630 v3 @ 2.40GHz
- 128 GB RAM each node
- Total disk capacity: 1 PB


# Hadoop, Hive, Pig

---

- Hadoop configuration can be customized, e.g. replication factor, heartbeat
- Queries in Hadoop are written in Java as MapReduce operations, text and binary format
- Hive is an SQL interface into Hadoop, data are uploaded into Hive representation
- Pig is a functional interface into Hadoop, data are stored in CSV format

# Results


Query 2 - "count of flows, port 53" I.






# Results

### Query 4 - "srcIP counts, order by bytes, TCP" I.



# Results: Hive parallel

## Hive - Query 2



# Hadoop summary

---

- No significant differences between native Hadoop query implementation and Hive
- Pig is worse and fails arbitrarily
- Hadoop utilize heartbeat messages not only to liveness detection but also to distribute jobs and collect results – this cause long latencies before retrieving first data (around 20s)

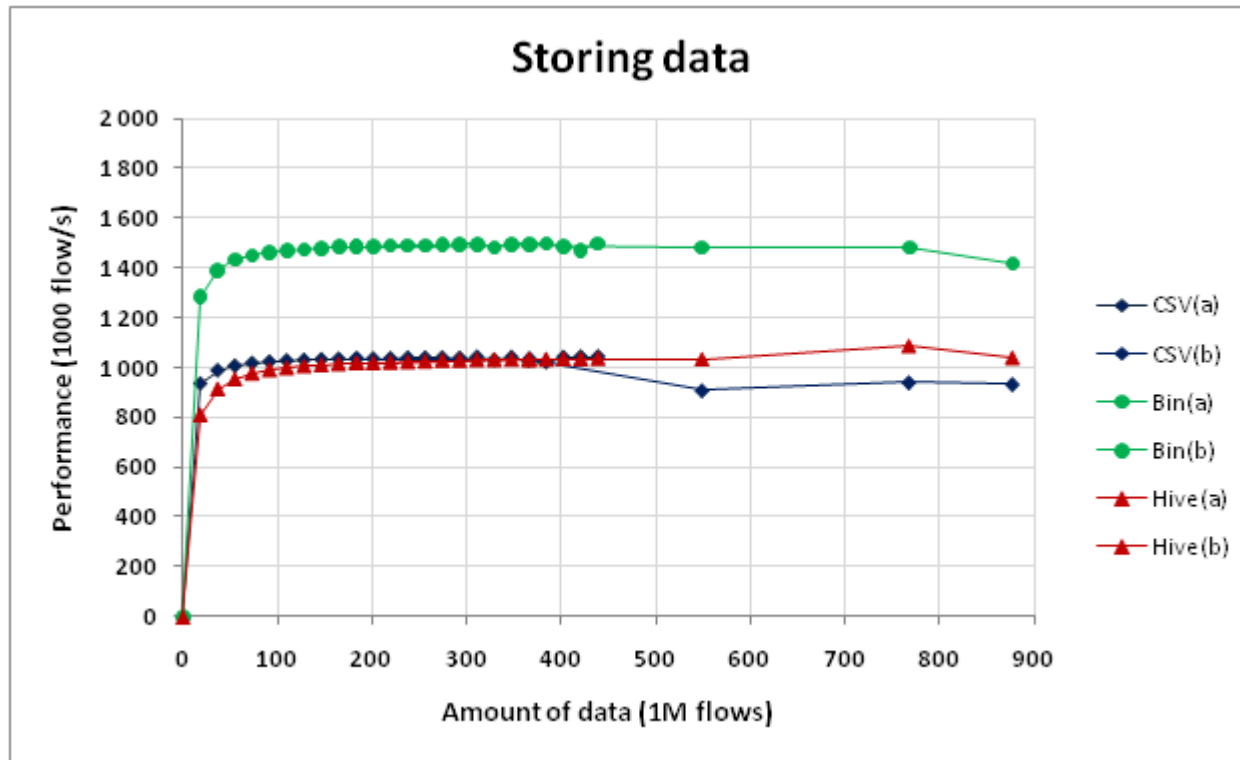
# Hadoop summary

---

- Java and long latency means low performance/effectivity per single node
  - Query over 877. mil records – performance lower than 1 mil. records/s per node
  - Query over 8 billion records – performance lower than 2 mil. records/s per node
  - nfdump on single node reaches 4+ mil. records/s per node
- Parallel queries improve single node perf.

# Hadoop results

- Data upload (877 mil. toků)



# nfdist

---

- Tool utilizing hdfs as a storage
- NfDump files are upload to hdfs
- Distributed nfdumps retrives data from hdfs and results are merged by nfcats tool

# NfDist summary

---

- Outdated with limitations
- Old nfdump format
- Limited by HDFS block size
- Performance per single node similar to Hadoop

# ElasticSearch

---

- 9+1 nodes
- Intel(R) Xeon(R) CPU E5-2650 v2 @ 2.60GHz
- Each node 8GB RAM



# Results Elasticsearch

---

- Queries over whole data set
- Query 2 – 1x 30s, 3x10s
- Query 3 – 1-2s
- Query 4 – 5200s

# ElasticSearch summary

---

- Extremely slow upload due to indexing
  - 877 mil. toků in 9hours without replication
  - 46 hours with replication
- Large index
  - Index is 4 times larger than data
- Fast response to filtration queries
  - Around 1s
- Limited by RAM


# Vertica

---

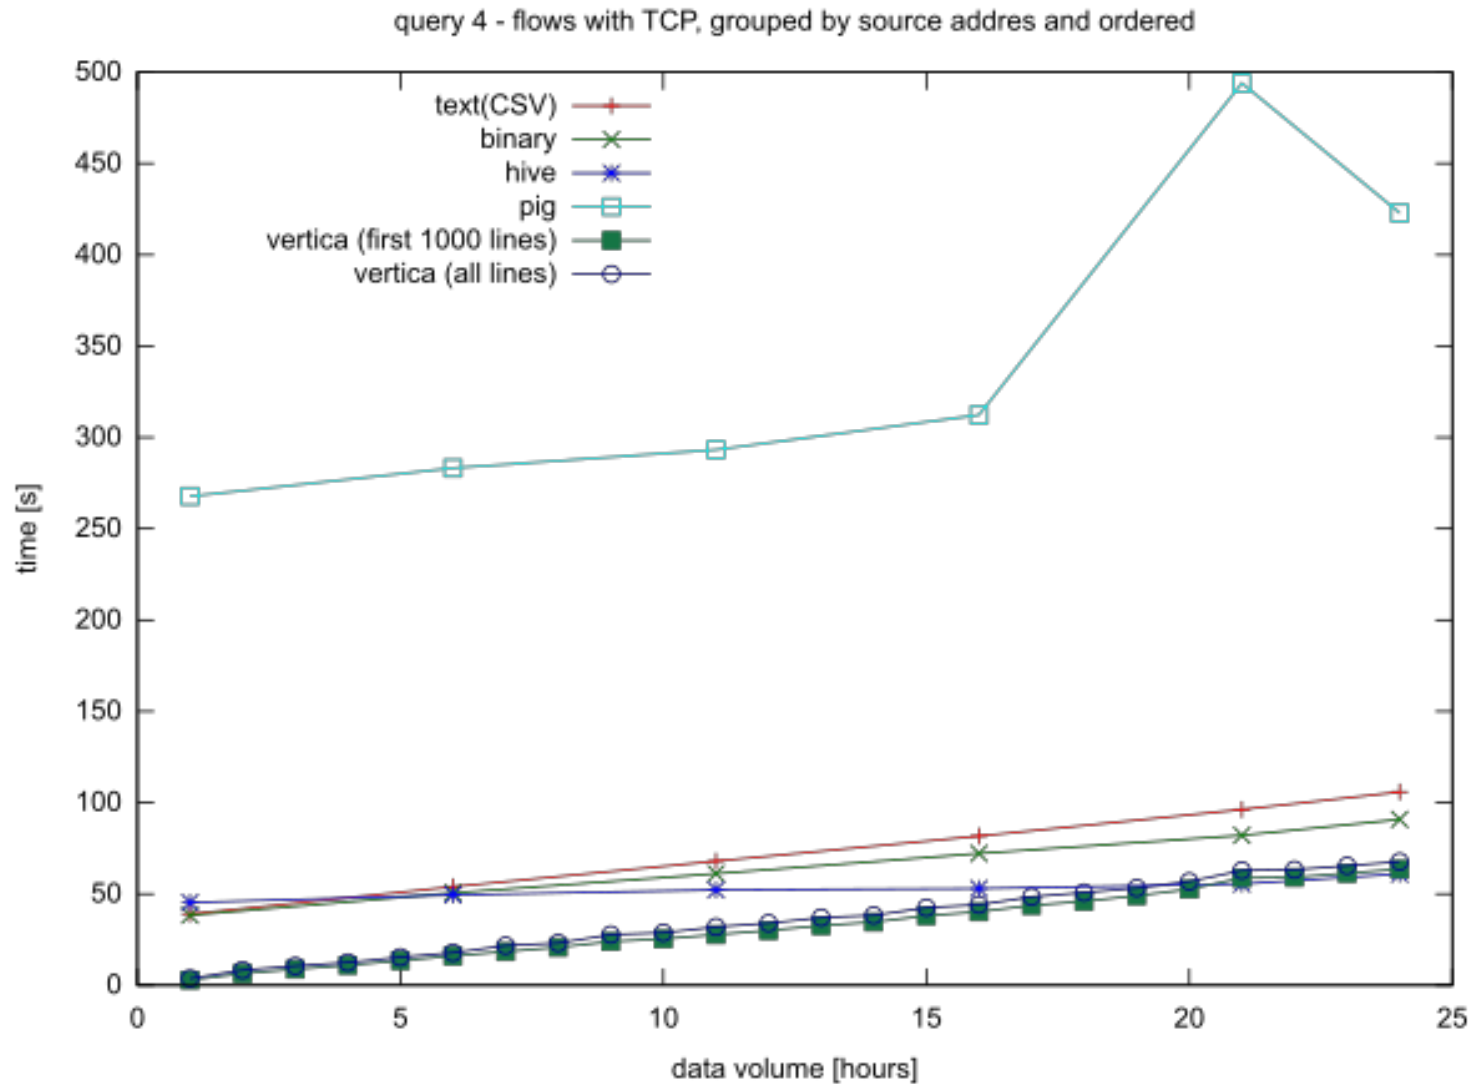
- 3 nodes
- CPU: 2 cores z [Intel E5-2670](#) @ 2600 MHz
- Each node 4 GB RAM

# Vertica

query 1 - count flows, packets and bytes sum



# Vertica



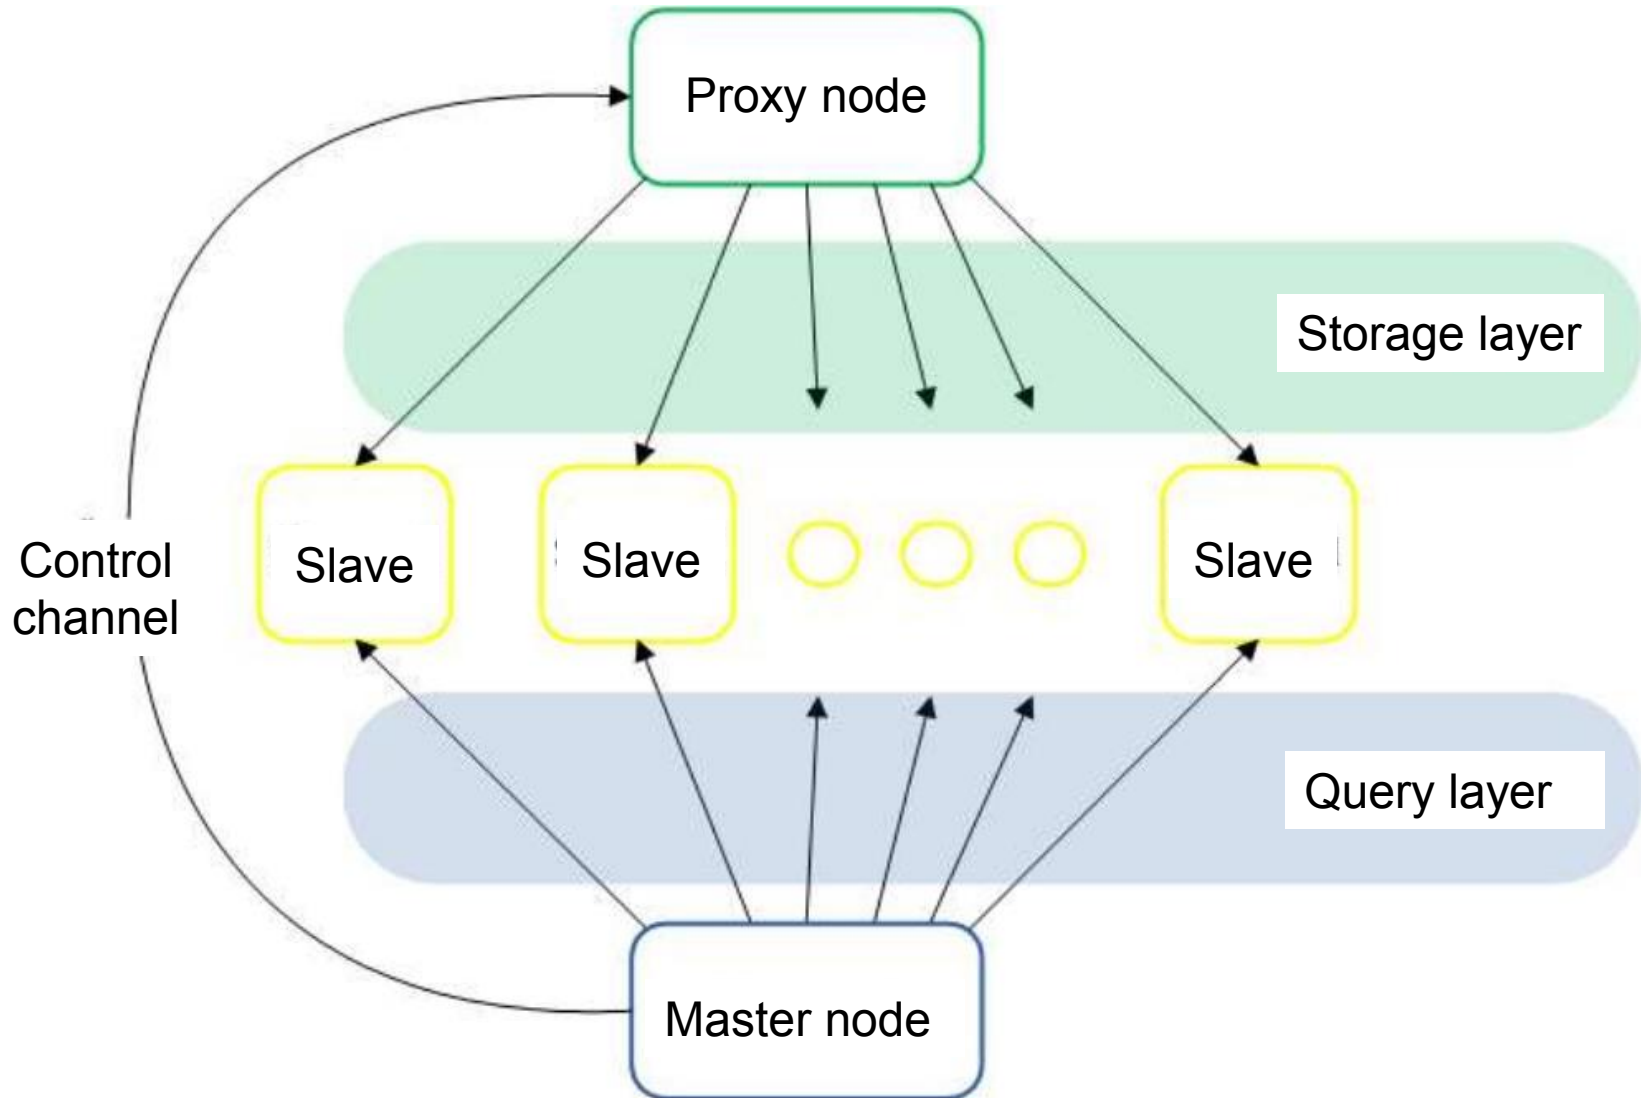
# Vertica summary

---

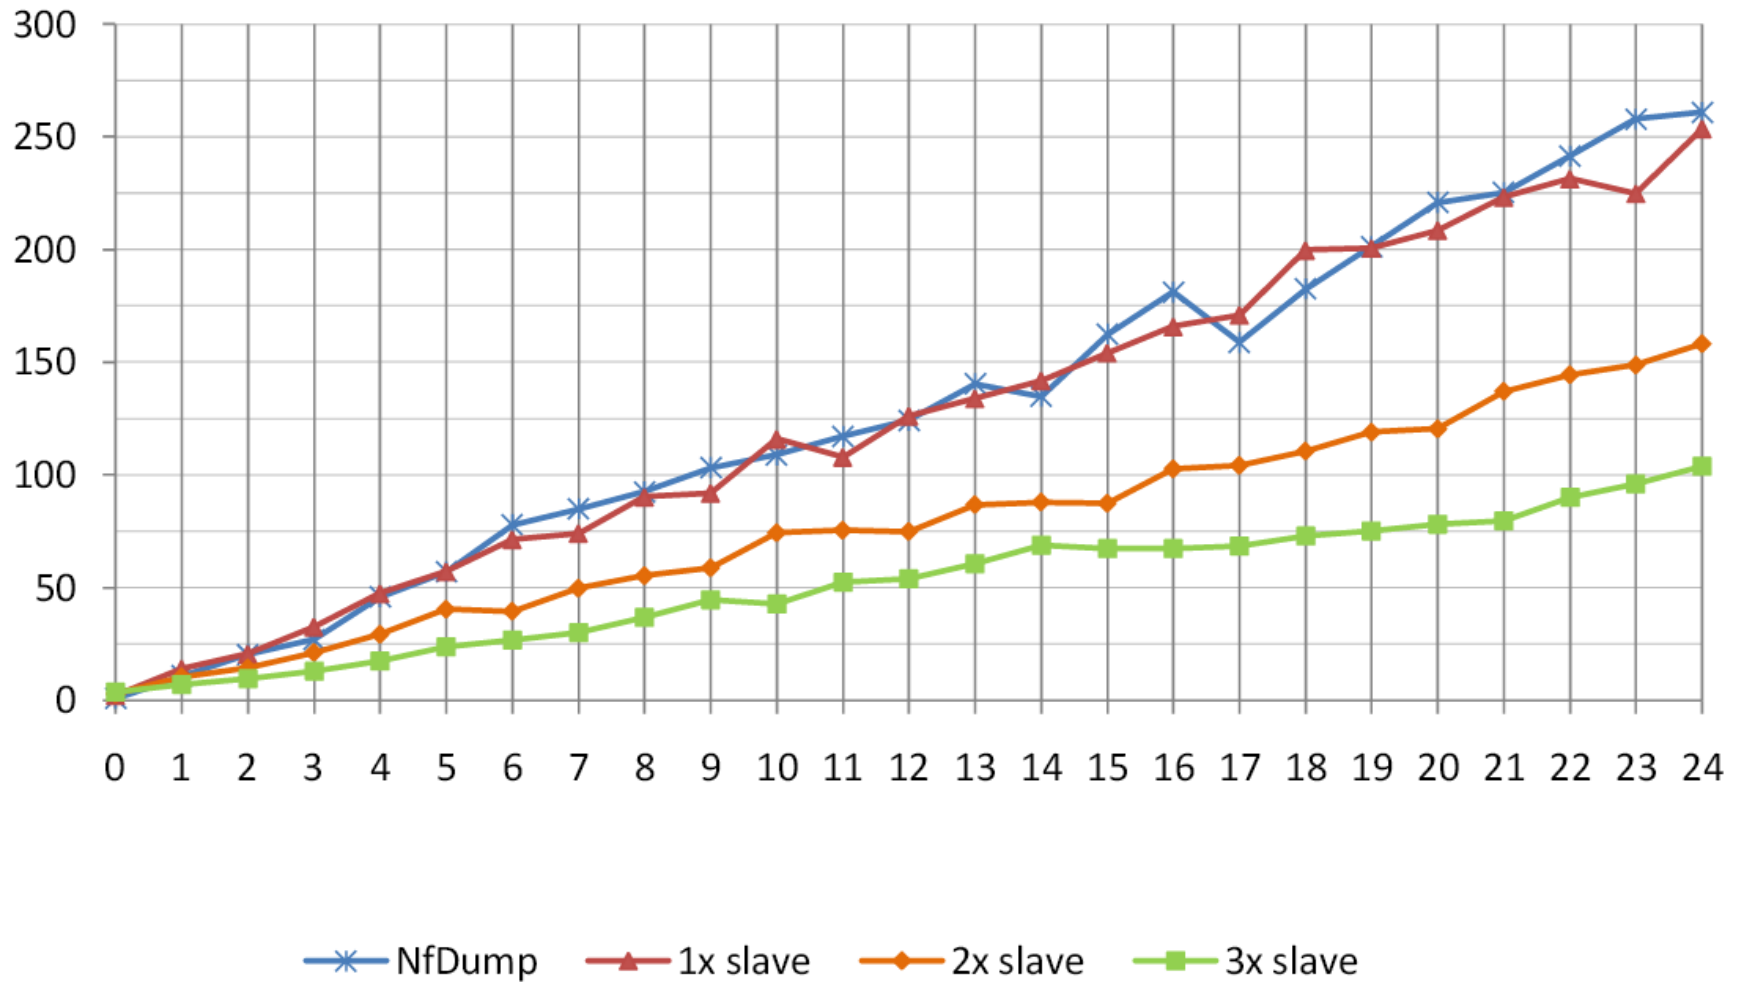
- Vertica is a column based DB
  - Allows to read only necessary fields from the record
  - Exploit thread paralellism
  - Deals with realibility
  - Publicly available up to 3 nodes

# Proprietary implementation

---



# Results





# Summary

---

- Proprietary implementation achieves high performance per single node in both tasks storage and queries.
- Does not support high-availability features and multi-thread support so far

# Conclusion

---

- Publicly available platforms exhibit certain limitations
- Flow collector deals with specific data and queries as such proprietary solution will always offer better parameters
- SecurityCloud project implements open source „big data“ flow collector which will be available 2015

# Acknowledgement

---

- This work is partially supported by  
Technological Agency of Czech Republic

**T A**

**Č R**