KODIAK DATA

TECHNICAL WHITE PAPER

ClickHouse on MemCloud
Performance Benchmarks by Altinity

March 2018

Version 1.0

Table of Contents 2

1 Abstract

Companies all over the world are constantly looking to improve the performance of systems and reduce the cost of those systems (OPEX and CAPEX). One of the ways they are accomplishing this is to move to cloud-based systems. In this whitepaper we will discuss a modern and cost-effective solution for data analytics deployed on Kodiak Data MemCloud.

2 Introduction

Performance in the cloud has always been a challenge when you are talking about data analytics. Altinity has partnered with Kodiak Data to provide the first managed cloud offering of ClickHouse DBMS for analytics. Altinity needed a high-performance, cost-effective solution to host ClickHouse for their customers, which was provided by Kodiak Data.

In enterprise data analytics there has always been a need to get the data faster and, at the same time, reduce the OPEX costs and the total cost of ownership to the companies' consuming products. This challenge extends beyond the enterprise to SMB as well as small startups. Being able to access data faster allows companies to gain insights into sales, marketing and customer service more quickly.

Challenges that are faced by all companies in this area are deeper than one might think. Solving the issue of setting up a high-speed DBMS and then having to hire personnel to maintain the database can be very costly. Each phase of deploying an MPP DBMS system presents significant challenges, including planning, procurement (Software and Hardware), deployment and testing, and putting into production. These have been simplified by MemCloud. The operational aspects of having the personnel to maintain the DBMS and keep the system protected from data loss is outside the scope of this paper, and is equivalent in both solutions.

Kodiak Data understands these challenges in the enterprise and in small business, and delivers a robust, high-speed cloud offering for Altinity's ClickHouse software stack with low-cost overhead from deployment to production turnover.

The challenge to deploying performance-based software in the cloud is that most clouds were not designed for high performance bandwidth and low latency. Altinity tested ClickHouse with AWS EC2 instances vs Kodiak Data instances and found that Kodiak Data was overall faster in performance and better than half the cost of other AWS solutions. Altinity tested ClickHouse vs. Amazon Redshift software and the cost was one fifth and over five times faster in completing the test runs.

Eight tests were performed on both software platforms and done in AWS EC2, as well as Kodiak Data's MemCloud offering.

Specific information on the testing can be found in the details below but the end result was a performance enhancement over AWS EC2 and Redshift systems by over 2 times in worst case scenario and over 6.5 times in the best case . . . all at better then a 70% cost reduction over other solutions based on performance cost.

To match the performance tested it would cost significantly more to deploy on AWS vs. MemCloud month over month. Kodiak Data and Altinity have partnered to provide a platform that not only can be easily deployed for Altinity's customers, but also provide the Clickhouse DBMS capabilities for Kodiak Data's customers.

Data below covers the results of the testing committed by Altinity, and you can see the blog post for the data here.

https://www.altinity.com/blog/2018/1/4/clickhouse-on-the-cloud-benchmarks

3 MPP analytic DBMS overview

A DBMS structure has a source that collects the data from specified locations. Once it has that data it will start to process it, as well as store the raw data form as it moves through the DBMS, and then -- based on queries from the company -- will present the data output in a structured format. Companies run reports on everything from production lifecycle to consumption of the product in the field, from customer responses for support to sales by volume and cost metrics. The source of the tests used in this evaluation was a series of New York City taxi rides. The queries were based on ride metrics: passenger counts, times, destinations and pickups. environmental variables.

4 Test overview

Baseline tests were designed to use a data package of NYC taxi rides using several platforms to run the same base queries. After loading the data into the DBMS, eight queries will be run against the model data and results captured, with the time for the query to complete being the evaluation criteria. The test was conducted with three platforms.

1. Clickhouse on AWS instances (three different EC2 instances), which will provide data from small, medium and large Clickhouse deployment
    a. AWS r4.xlarge x1 (4vCPU (13.5 ECU), 30.5 GB RAM, EBS (Amazon Elastic Block Storage)
    b. AWS i3.2xlarge x1 (8vCPU (27 ECU), 61 GB RAM, 1 x 1990 GB NVMe SSD
    c. AWS i3.4xlarge x1 (16vCPU (53 ECU), 122 GB RAM, 2 x 1990 GB NVMe SSD
2. Clickhouse on Kodiak Data MemCloud
    a. Kodiak A.1 (8vCPU, 16GB RAM, 2TB NVMe SSD vDisk storage)
    b. Kodiak A.3 (3 x 8vCPU, 16GB RAM, 2TB NVMe SSD vDisk storage)

3. Redshift DC2.8xlarge (36vCPU, 244GB RAM, 2TB NVMe SSD)

All of the above units except the AWS Redshift unit (largest available at the time of test) were loaded with the same version of Linux and the same version of ClickHouse software. The tests were run based on the queries below and timed (with modifications for Redshift).

```
Q1:
SELECT passenger_count, avg(total_amount)
FROM yellow_tripdata_staging
 GROUP BY passenger_count;
Q2:
SELECT passenger_count, toYear(tpep_pickup_datetime) AS year, count(*)   FROM
yellow_tripdata_staging
 GROUP BY passenger_count, year
Q3:
SELECT passenger_count, toYear(tpep_pickup_datetime) AS year, round(trip_distance) AS
distance, count(*)
  FROM yellow_tripdata_staging
 GROUP BY passenger_count, year, distance
 ORDER BY year, count(*) DESC;
Q4:
-- ClickHouse
SELECT dictGetString('taxi_zones', 'zone', toUInt64(pickup_location_id)) AS zone,
count(*)
  FROM yellow_tripdata
 GROUP BY pickup_location_id
 ORDER BY count(*) DESC
 LIMIT 10;
```

```
-- RedShift
SELECT min(zone), count(*)
  FROM yellow_tripdata_staging td
  left join taxi_zones tz on tz.locationid = td.pickup_location_id
 GROUP BY pickup_location_id
 ORDER BY count(*) DESC
 LIMIT 10;

Q5:
-- ClickHouse
SELECT count(*)
```

```
  FROM yellow_tripdata
 WHERE dictGetString('taxi_zones', 'zone', toUInt64(pickup_location_id)) = 'Midtown
East';

-- RedShift
SELECT count(*)
  FROM yellow_tripdata_staging td
  left join taxi_zones tz on tz.locationid = td.pickup_location_id
 WHERE zone = 'Midtown East';
Q6:

-- ClickHouse
SELECT dictGetFloat32('weather', 'min_temperature', ('GHCND:USW00094728',
pickup_date)) AS t,
       count() AS c
  FROM yellow_tripdata
  GROUP BY t
  ORDER BY c DESC
  LIMIT 10;

-- RedShift:
SELECT min_temperature t, count(*) c
  FROM yellow_tripdata_staging td
  left join central_park_weather_observations w on w.date =
trunc(td.tpep_pickup_datetime)
 GROUP BY t
 ORDER BY c DESC
 LIMIT 10;
Q7:
-- ClickHouse
SELECT dictGetFloat32('weather', 'snowfall', ('GHCND:USW00094728', pickup_date)) AS
t,
       avg((tpep_dropoff_datetime - tpep_pickup_datetime) / 60) AS d
  FROM yellow_tripdata_staging
  GROUP BY t
  ORDER BY t DESC
  LIMIT 10;


-- RedShift
SELECT snowfall t,
       avg(datediff(minute,tpep_pickup_datetime,tpep_dropoff_datetime)) AS d
  FROM yellow_tripdata_staging td
  left join central_park_weather_observations w on w.date =
trunc(td.tpep_pickup_datetime)
 GROUP BY t
 ORDER BY t DESC
 LIMIT 10;

Q8:
-- ClickHouse
```

```
SELECT dictGetUInt32('weather', 'snowfall', ('GHCND:USW00094728', pickup_date)) AS t,
       sum(passenger_count) AS c
  FROM yellow_tripdata
  GROUP BY t
  ORDER BY t DESC
  LIMIT 10;

-- RedShift
SELECT snowfall t,
       sum(passenger_count) AS d
  FROM yellow_tripdata_staging td
  left join central_park_weather_observations w on w.date =
trunc(td.tpep_pickup_datetime)
 GROUP BY t
 ORDER BY t DESC
 LIMIT 10;
```

5 Test Information and Results

The base control is the same data and queries were used in both ClickHouse and Redshift software. The benchmarks were evaluated on performance (in this case, time to complete task) and cost of the systems that the software is running on in the cloud.

A. MemCloud vs. AWS ClickHouse in small deployment (Cloud)
    a. Test setup is described in section 4
    b. All the tests were run and timed as well as cost input below
    c. Test results
        i. AWS r4.xlarge x1 took 162 seconds to complete
        ii. Kodiak Data MemCloud A.1 took 50 Seconds to complete

    Result = ClickHouse on MemCloud was 3+ times faster at slightly less cost

B. MemCloud vs AWS ClickHouse in medium deployment (Cloud)
    a. Test Setup is described in section 4
    b. All the tests were run and timed as well as cost input below
    c. Test results
        i. AWS i3.2xlarge x1 took 86 seconds to complete
        ii. Kodiak Data MemCloud A.1 took 50 seconds to complete

    Result = ClickHouse on MemCloud was 1.7 times faster at 2.6 times less cost

C. MemCloud vs. AWS ClickHouse in large deployment (Cloud)
    a. Test setup is described in section 4
    b. All the tests were run and timed as well as cost input below

c. Test results
    i. AWS i3.4xlarge x1 took 44 seconds to complete
    ii. Kodiak Data MemCloud A.3 took 18 seconds to complete

Result = ClickHouse on MemCloud was 2.4 times faster at less than 50% the cost

D. MemCloud ClickHouse vs. Redshift in largest AWS deployment (Cloud)
    a. In this test ClickHouse software was deployed one for one (same OS and configuration on Kodiak MemCloud A.3 x1 and AWS dc2.8xlarge [RedShift] x1)
    b. Test setup is described in section 4
    c. All the tests were run and timed as well as cost input below
    d. Test results
        i. AWS dc2.8xlarge x1 took 122 seconds to complete
        ii. Kodiak Data MemCloud A.3 took 18 seconds to complete

Result = ClickHouse on MemCloud was 6.7 times faster at less than 14% of the cost

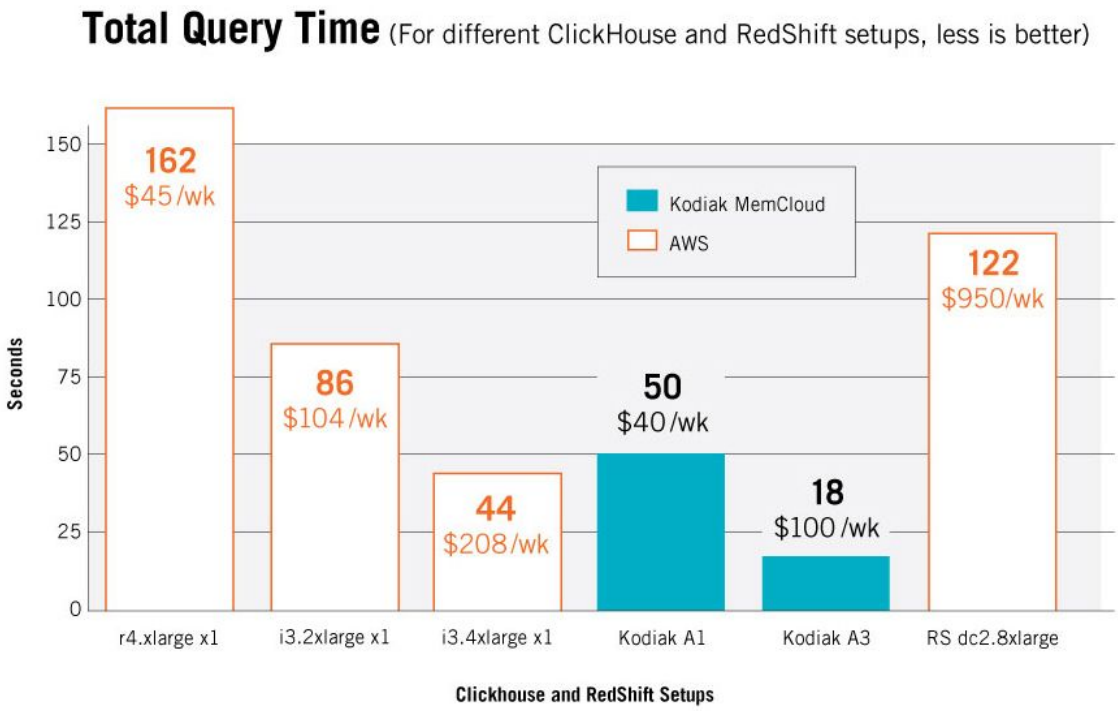

6 Base Testing Information

The raw systems used for this testing were:

ClickHouse:

- AWS r4.xlarge x1: 4 vCPU, 13.5 ECU, 30.5GB RAM, EBS
- AWS i3.2xlarge x1: 8 vCPU, 27 ECU, 61GB RAM, 1 x 1990 NVMe SSD
- AWS i3.4xlarge x1: 16 vCPU, 53 ECU, 122GB RAM, 2 x 1990 NVMe SSD
- Kodiak A.1: 8 vCPU, 16GB vRAM, 2TB NVMe SSD vDisk
- Kodiak A.3: 3 * (8 vCPU, 16GB vRAM, 2TB NVMe SSD vDisk)

RedShift:

- RedShift dc2.8xlarge: 36 vCPU, 244GB vRAM, 2.5TB NVMe SSD

It should be noted that all the prices are taken from public sources. Prices are approximate, taken from:

1. For Amazon, prices refer to https://aws.amazon.com/ec2/pricing/on-demand/
2. RedShift prices can be found at https://aws.amazon.com/redshift/pricing/
3. http://calculator.s3.amazonaws.com/index.html can be used to estimate discounted long term commitments with Amazon
4. Kodiak A.1 and A.3 prices are discounted quotes for Altinity clients

7 Summary

ClickHouse on Kodiak Data MemCloud ran faster and more cost effectively than either ClickHouse or RedShift on standard AWS EC2 instances. Customers interested in high performance and lower costs for their ClickHouse operations should evaluate the Kodiak Data cloud platform.

About Kodiak Data

Kodiak Data is a leading cluster virtualization technology company that allows customers to easily deploy and scale Big Data infrastructure in both public and private clouds. The Kodiak Data Virtual Cluster Infrastructure (VCI) platform is the only solution than can create, within minutes, code-ready virtual clusters that run at memory-speed and scale to the needs of big data applications. For more information about Kodiak Data, visit www.kodiakdata.com.

About Altinity

Altinity is the leading service provider for ClickHouse – an open-source column-store analytic database. Altinity helps customers around the world take advantage of the cutting-edge technology

for managing and analyzing extremely large volumes of data without expensive license costs. Altinity provides 24/7 ClickHouse Support, Professional Services, and Training. Visit altinity.com or follow us @AltinityDB. For more information, please contact info@altinity.com.

KODIAK DATA, INC.

2570 W El Camino Real #500, Mountain View, CA 94040 | 650 383 8374
info@kodiakdata.com | www.kodiakdata.com

Altinity

Office 2, Derby House, 123 Watling Street
Gillingham, Kent, ME7 2YY
info@altinity.com | www.altinity.com

KDWP01802022-01