



**TripCom**

*Triple Space Communication*

FP6 – 027324

**D1.4**  
**Storage Performance Evaluation**

Vassil Momtchev (ONTO)  
Brahmanda Sapkota (NUIG)  
Andreas Harth (NUIG)  
Omair Shafiq (LFUI)  
Atanas Kiryakov (ONTO)

**EXECUTIVE SUMMARY**

Deliverable D1.4 “Storage Performance Evaluation” presents extensive performance and scalability evaluation of the TripCom storage infrastructure. The storage component plays critical role in the overall system productivity. It ensures durability of Triple Space information model and the internal components data of a single kernel. This document summarizes common RDF tasks to be automated by a triple store, investigates the associated performance factors and evaluates the results of WP1 efforts to implement a high-performance and scalable data layer.

## Document Information

<b>IST Project Number</b>	FP6 – 02734	<b>Acronym</b>	TripCom
<b>Full Title</b>	Triple Space Communication		
<b>Project URL</b>	<a href="http://www.tripcom.org">http://www.tripcom.org</a>		
<b>Document URL</b>	<a href="http://www.tripcom.org/docs/del/D1.4.pdf">http://www.tripcom.org/docs/del/D1.4.pdf</a>		
<b>EU Project Officer</b>	Werner Janusch		

<b>Deliverable</b>	<b>Number</b>	1.4	<b>Title</b>	Storage Performance Evaluation.
<b>Work Package</b>	<b>Number</b>	1	<b>Title</b>	Storage










<b>Date of Delivery</b>	<b>Contractual</b>	M30	<b>Actual</b>	M32
<b>Status</b>	version 1.0		final <input type="checkbox"/>	
<b>Nature</b>	prototype <input type="checkbox"/> report <input checked="" type="checkbox"/> dissemination <input type="checkbox"/>			
<b>Dissemination level</b>	public <input checked="" type="checkbox"/> consortium <input type="checkbox"/>			

<b>Authors (Partner)</b>	Vassil Momtchev (ONTO), Brahmananda Sapkota (NUIG), Andreas Harth (NUIG), Omair Safiq (LFUI), Atanas Kiryakov (ONTO)			
<b>Resp. Author</b>	Vassil Momtchev		<b>E-mail</b>	vassil.momtchev@ontotext.com
	<b>Partner</b>	Ontotext	<b>Phone</b>	+359-2-8091553

<b>Abstract (for dissemination)</b>	Deliverable D1.4 “Storage Performance Evaluation” presents extensive performance and scalability evaluation of the TripCom storage infrastructure. The storage component plays critical role in the overall system productivity. It ensures durability of Triple Space information model and the internal components data of a single kernel. This document summarizes common RDF tasks to be automated by a triple store, investigates the associated performance factors and evaluates the results of WP1 efforts to implement a high-performance and scalable data layer.
<b>Keywords</b>	RDF, triple store, scalability, performance, evaluation

<b>Version Log</b>			
<b>Issue Date</b>	<b>Rev. No.</b>	<b>Author</b>	<b>Change</b>
September 6, 2008	1	Vassil Momtchev	Initial version and ToC
October, 10, 2009	2	Vassil Momtchev	Performance factor specification
December 15, 2009	3	Vassil Momtchev	Pre-final version of the document
December 20, 2009	4	Vassil Momtchev	Final version

## PROJECT CONSORTIUM INFORMATION

Semantic Technology Institute <a href="http://www.sti-innsbruck.at">http://www.sti-innsbruck.at</a>	<p>STI</p>  <p>STI · INNSBRUCK</p>	Prof. Dr. Dieter Fensel Semantic Technology Institute (STI) Innsbruck, Austria E-mail: dieter.fensel@sti-innsbruck.at
National University of Ireland, Galway <a href="http://www.deri.ie">http://www.deri.ie</a>	<p>NUIG</p>  <p>National University of Ireland, Galway <i>Ollscoil na hÉireann, Gaillimh</i></p>	NUIG Dr. Laurentiu Vasiliu Digital Enterprise Research Institute (DERI) Galway, Ireland Email: laurentiu.vasiliu@deri.org
University of Stuttgart <a href="http://www.iaas.uni-stuttgart.de/">http://www.iaas.uni-stuttgart.de/</a>	<p>USTUTT</p>  <p>Universität Stuttgart</p>	USTUTT Prof.Dr. Frank Leymann Inst. f'ur Architektur von Anwendungssystemen (IAAS) Stuttgart, Germany E-mail: frank.leymann@informatik.uni-stuttgart.de
Vienna university of Technology <a href="http://www.complang.tuwien.ac.at/">http://www.complang.tuwien.ac.at/</a>	<p>TUW</p>  <p>TECHNISCHE UNIVERSITÄT WIEN VIENNA UNIVERSITY OF TECHNOLOGY</p>	TUW Prof.Dr. eva K'uhn Institut f'ur Computersprachen Vienna, Austria E-mail: eva@complang.tuwien.ac.at
Free University Berlin <a href="http://www.ag-nbi.de/">http://www.ag-nbi.de/</a>	<p>FUB</p>  <p>Freie Universität Berlin</p>	FUB Prof. Dr.-Ing. Robert Tolksdorf AG Netzbasierte Informationssysteme Berlin, Germany E-mail : tolk@inf.fu-berlin.de
Ontotext Lab, Sirma Group Corp. <a href="http://www.ontotext.com/">http://www.ontotext.com/</a>	<p>ONTO</p>  <p>Ontotext Knowledge and Language Engineering Lab of Sirma</p>	ONTO Atanas Kiryakov, Vassil Momtchev, Ontotext Lab, Sirma Group Corp. Sofia, Bulgaria E-mail: vassil.momtchev@ontotext.com
Profium OY <a href="http://www.profium.com/">http://www.profium.com/</a>	<p>Profium</p>  <p>profium</p>	Profium Dr. Janne Saarela Profium OY Espoo, Finland E-mail: janne.saarela@profium.com
CEFRIEL SCRL. <a href="http://www.cefriel.it/">http://www.cefriel.it/</a>	<p>Cefriel</p>  <p>FORGING INNOVATION BY TECHNOLOGY</p>	CEFRIEL Davide Cerri CEFRIEL SCRL. Milano, Italy E-mail: cerri@cefriel.it
Telefonica I+D <a href="http://www.tid.es/">http://www.tid.es/</a>	<p>TID</p>  <p>TELÉFÓNICA INVESTIGACIÓN Y DESARROLLO</p>	Noelia Pérez Crespo Telefonica I+D Madrid, España E-mail: npc@tid.es

**TABLE OF CONTENTS**

Executive Summary .....	II
Project Consortium Information .....	IV
Table of Contents .....	V
List Of Abbreviations .....	1
1 Introduction.....	2
2 Storage Tasks and Performance Factors .....	3
2.1 Loading and pre-processing the ontology.....	3
2.2 Query processing .....	4
3 Evaluated Systems .....	5
3.1 Triple Reasoning and Rule Entailment Engine.....	5
3.2 Yet Another RDF Store .....	5
4 Scalable Storage Infrastructure .....	6
4.1.1 Triple Reasoning and Rule Entailment Engine.....	6
4.1.2 Yet Another RDF Store .....	8
4.1.3 Comparison with public scalability results of other systems.....	9
5 Query Execution Performance.....	12
5.1 Berlin SPARQL Benchmark (BSBM) .....	12
5.1.1 Comparison with public results of other systems .....	14
6 Discussions .....	17
7 Conclusion .....	18
References.....	19
Appendix.....	21

**LIST OF ABBREVIATIONS**

<b>ANSI</b>	American National Standards Institute
<b>BSD</b>	Berkeley Software Distribution
<b>DBMS</b>	Database Management Systems
<b>ER</b>	Entity Relationship
<b>FOAF</b>	Friend Of a Friend
<b>GNU</b>	GNU's Not Unix
<b>GPL</b>	GNU General Public Licence
<b>HTTP</b>	Hyper Text Transfer Protocol
<b>iTQL</b>	Interactive Tucana Query Language
<b>JRDF</b>	Java RDF
<b>LAN</b>	Local Area Network
<b>LFUI</b>	Leopold-Franzen-Universität Innsbruck
<b>LGPL</b>	GNU Lesser General Public Licence
<b>LUBM</b>	Lehigh University Benchmark
<b>N3</b>	Notation 3
<b>N3QL</b>	N3 Query Language
<b>NDM</b>	Oracle Spatial Network Data Model
<b>OASIS</b>	Organization for the Advancement of Structured Information Standards
<b>ORDI</b>	Ontology Representation and Data Integration
<b>OWL</b>	Web Ontology Language
<b>OWLIM</b>	OWL In Memory
<b>RDBMS</b>	Relational DBMS
<b>RDF</b>	Resource Description Framework
<b>RDFS</b>	RDF Schema
<b>RDQL</b>	RDF Data Query Language
<b>ROI</b>	RDF Input/Output
<b>SAIL</b>	Storage And Inference Layer
<b>SOFA</b>	Simple Ontology Framework API
<b>SOAP</b>	Simple Object Access Protocol
<b>SeRQL</b>	Sesame RDF Query Language
<b>SEQUEL</b>	Structured English Query Language
<b>SPARQL</b>	SPARQL Protocol and RDF Query Language
<b>SQL</b>	Structured Query Language
<b>TAP</b>	The Alpiri Project
<b>TCP</b>	Transmission Control Protocol
<b>URI</b>	Uniform Resource Identifier
<b>URL</b>	Uniform Resource Locator
<b>W3C</b>	World Wide Web Consortium
<b>WSDL</b>	Web Service Description Language
<b>WSMO</b>	Web Service Modeling Language
<b>XML</b>	Extensible Markup Language
<b>YARS</b>	Yet Another RDF Store
<b>YARSQL</b>	YARS Query Language

## 1 INTRODUCTION

TripCom storage infrastructure is responsible to ensure the durability of the Triple Space information model and the internal components data. The storage infrastructure is designed to work in a highly heterogeneous environment and to support different use case scenarios. The storage component plays a critical role in the overall system productivity, so a special investigation is conducted to test its scalability and performance limits. An extended RDF data model specification is presented in [4] to enable the representation of different schemata and information models associated with meta-data. The extended data model uses named graphs and triplesets – a novel primitive, which allows the efficient association of meta-data to portions of a dataset (RDF mutli-graph, as in SPARQL). Hence, multiple independent types of meta-data could be expressed for every tripleset like author, creation date and space.

The document investigates the scalability and performance limits of the persistence storage for a single kernel instance in TripCom. The performance and scalability are important system properties that indicate the amounts of data and requests that a single kernel can handle. Chapter 2 investigates metrics to measure common RDF tasks automated by the triple stores and analyzes the associated performance factors. Thus, the evaluation of the storage infrastructure performance is divided in two categories (i) data loading and (ii) query answering. In chapter 3 a short summary of the evaluated systems is presented together with a brief discussion on the evolution of the software during the evaluation process. Next, using a standard benchmark test, we measure the amount of RDF data that could be loaded in a storage instance and how it is affected by the different performance factors. A comparison with the best known public results is also presented. Chapter 5 surveys the data retrieval task by using a suite of benchmarks to measure the query answering performance over standard datasets. Lastly, a discussion is placed to comment the results and suggest new direction of refinements of the implementations.

The investigated storage infrastructure and data model is specified in ORDI framework. The framework defines an abstract persistence interface and adds a minimal layer of logic to automate configuration related tasks. The storage is integrated in overall TripCom storage infrastructure by TSAdapter, which is responsible to process all coordination interfaces and convert the ORDI data types to serializable types supported by the other components. The performance of TSAdapter is heavily affected by the deployment scenarios and whether in process or inter-process communication is used. Thus, to obtain results which are directly comparable with the state-of-the-art systems the interfaces of the underlying storage engines are used.

## 2 STORAGE TASKS AND PERFORMANCE FACTORS

Performance is a quantity used to characterize the amount of useful work accomplished by a specific computer system. This chapter compares the tasks automated by semantic storage systems and defines measurable performance factors to be evaluated and considered.

The storage infrastructure is responsible to automate the tasks of persistent RDF storage and query evaluation. Their efficiency and performance depends on several key factors, which has to be considered during the evaluation process:

- Scale of information – the number of asserted RDF statements is a key factor for the performance of a system. In chapter 4 a detailed discussion for the scalability of Triple Space storage is presented.
- Hardware resources – the performance has a direct correlation to the available computer resources. More random access memory means that the number of disk operations could be reduced (in some extreme cases even to be eliminated).
- Reasoning complexity and dataset specificity – even the tractable fragments of logic could cause degradation of the performance for large knowledge bases. The highly interconnected datasets with long transitive and symmetric chains will increase the computational complexity many times.
- Number of clients – the number of concurrent clients affects the scalability of the system and increases the sum execution time (the time required to complete a specific task) because the need of internal synchronization and isolation mechanisms.

In [15], Weithoner identifies two different stages in the ontology processing discussed in subsections below.

### 2.1 Loading and pre-processing the ontology

The persisted RDF data in a repository must remain there until it is removed by an authorized user. To guarantee the consistency of the data this task includes the processing by the engine to write the data to a persistence storage device and to generate all indexes required for immediate data retrieval. A measurable metric for the persistence of RDF data is the time needed to process a specific number of statements for a period of time. For all our tests we select the quantity:

$$\text{Loading ontology performance} = \text{thousands statements} / \text{second}$$

The high-performance triple stores utilize comprehensive data structures in order to minimize the computational time to persist the input data. The RDF data persistence metric depends on several factors:

- Materialization – whether the implicit statements inferred by the reasoner are indexed and stored in the model. Systems which implement reasoning with *total*



- materialization*, a term used to denote that all implicit data is stored, tend to have a slower loading time but a faster response time.
- Complexity of the data model – The processing of a specific number of triple statements typically consumes less bandwidth than a quadruple data model. The Triple Space storage data model defined in [4] is an extension of RDF named graph specification with rich meta-data support. Such model would require fewer statements to express a given knowledge, but the processing of each one would take more time.
  - Guarantees for atomicity, consistency, isolation and durability – the support of different levels of transactional isolation requires the implementation of sophisticated locking mechanisms, which causes an overhead.

## 2.2 Query processing

The task of query processing (in the context of SPARQL query language, which does not support data modifications) includes the following activities:

1. Parsing: validate the syntax of a given query,
2. Preparation: generate an execution plan,
3. Evaluation: bind all free variables and
4. Fetching: return the result to the client.

The list contains important factors to affect the query evaluation performance, besides the query complexity (e.g. the number of triple patterns):

- Deduction – the opposite approach of *materialization*, where the inference is performed upon a query request. This process is also referred as backward-chaining.
- Size of the returned result – a test SPARQL query with LIMIT may take significantly less time for overall execution, because less data has to be transferred between the semantic storage and client's processes. This factor is important when comparing the query response time with no LIMIT executed over different scales of information.
- Query expressivity – a query with support of features like full-text search or negation clauses may reduce significantly the complexity (e.g. to replace multiple LIKE clauses with a single full text expression)

Unlike the loading ontology performance metric to test the query evaluation performance a benchmark specific value is needed. For example it could be the time to execute a set of benchmark queries or the average execution time. The quantity to measure the query evaluation process gives less direct indication of the overall performance because it is affected by many independent parameters like query and data set complexity.

### 3 EVALUATED SYSTEMS

The chapter introduces the evaluated systems in the current document. The two evaluated implementations are optimized to handle different usage scenarios. TRREE engine is a reasoning and storage system to demonstrate high scalability and performance. YARS focuses over the distribution of the RDF data model to a cluster of machines and increases the system performance by parallelization. In YARS2 a limited tractable reasoning is introduced, but it still not sufficient to answer in a sound way the most of standard OWL benchmarks.

#### 3.1 Triple Reasoning and Rule Entailment Engine

Triple Reasoning and Rule Entailment Engine is a scalable reasoning and persistence engine integrated in ORDI framework and TripCom storage infrastructure. The engine is a complete rewrite compared the previous versions and realizes an efficient forward-chaining reasoning algorithm. It also provides native support for the triplesets. TRREE is distributed into two variants – SwiftTRREE (in-memory version) and BigTRREE (file-based version). For the purpose of the current evaluation process two different versions of BigTRREE are tested:

- BigTRREE 3.0.x – the first release of the BigTRREE engine with native tripleset support
- BigTRREE 3.1.x – version to demonstrates major performance improvements and new internal file structure

In the evaluation process our team accomplished multiple optimizations and refinements. The sequence of versions 3.1.x is regarded as significant improvements of 3.0.x version family and demonstrates superior performance results.

#### 3.2 Yet Another RDF Store

Yet Another RDF Store is distributed RDF storage engine that is integrated in ORDI framework and TripCom storage infrastructure. In its latest version it is part of Semantic Web Search Engine (SWSE). SWSE is a distributed architecture for reasoning, storing and querying of RDF data with context, where the core is YARS2: a SPARQL compliant engine operating over multiple machines. A new feature in YARS2 is the reasoning support using the external Scalable Authoritative OWL Reasoner (SOAR) engine. SOAR is decoupled from YARS2 implementation and aims using carefully crafted subset of OWL

Reasoning is supported via the SAOR engine, which is decoupled from YARS2. The SAOR engine is designed to accept as input a set of ground statements and produce, through forward-chaining materialisation, a set of inferred statements according to a set of given rules. There are given sets of rules within SAOR with different levels of complexity and efficiency.

## 4 SCALABLE STORAGE INFRASTRUCTURE

This chapter investigates the scalability limits of the Triple Space storage. A standard benchmark is used to test the storage infrastructure and report a comparison with other storage systems.

In terms of software engineering, Bondi [6], specifies scalability as a property of a process, which indicates its ability to either handle growing amounts of work in a graceful manner, or to be readily enlarged. Strictly speaking in the database community, the term database scalability is one that can be upgraded to process more transactions by adding new processors, devices and storage. The advance of modern hardware and the usage of SAN solutions combined with the low cost of the disks have virtually eliminated practical limits to raw storage capacity, [8]. Despite that fact, the most of existing semantic storage systems are still unable to process virtually unlimited number of triples, [7]. Thus, the scope of the current benchmarking is mainly focused evaluating: 1) dataset size / loading time 2) dataset size / query answering time. A more comprehensive investigation of the second metric is placed also in chapter 4 (see a detailed explanation in the next paragraphs).

To test the ability of scaling in loading very large data set a list of tasks is specified to be completed:

- Processing the input RDFXML files
- Permanently persisting the data and building indexes
- Reasoning (if appropriate)
- Query answering and completeness/soundness

### 4.1.1 Triple Reasoning and Rule Entailment Engine

At the present time a very limited set of public data sources may supply the required size of information and list of queries to verify the correctness of the import. The Lehigh University Benchmark (LUBM) is an outstanding benchmark for OWL repository scalability, [5] and it is widely adopted by the community. It employs synthetically generated datasets using a fixed OWL ontology of university organizations. The complexity of the language constructs used is between the ones of OWL Horst. 14 queries are defined by the LUBM benchmark. These queries will be used to check the query evaluation correctness and if the system is able to answer in reasonable time period. The biggest standard dataset is LUBM(50, 0), which contains synthetic data for 50 universities with specific seed (a pseudorandom number).

For the scalability tests the biggest standard version of LUBM with 50 universities is skipped, because its size of 6.89 million explicit statements is still too insignificant for the Triple Space storage scalability targets. A more challenging datasets like LUBM(1000, 0) and LUBM(8000, 0) are used, which contain respectively 133 and 1086 million explicit statements.

The main objective of the evaluation process is to determine the way the different performance factors described in chapter 2 affects the RDF data persistence. Several different system setups are tested:


Hardware and JVM memory	Semantics	LUBM parameters	Millions of statements	Load time sec	Speed st/sec
E5420 @ 2.50GHz(4HT) 2GHz <sup>1</sup> , 13,32GB <sup>2</sup>	none	(1000, 0)	138	2,609	52,894
E5420 @ 2.50GHz(4HT) 2GHz, 9,65GB	OWL-Horst	(1000, 0)	138	11,056	12,482
x86_64 2 x Dual Core Opteron 2GHz, 3GB	OWL-Horst	(1000, 0)	138	14,771	9,343
E5420 @ 2.50GHz(4HT) 2GHz, 8GB	none	(8000, 0)	1068	20,867	51,181
E5420 @ 2.50GHz(4HT) 2GHz, 32GB	OWL-Horst	(8000, 0)	1068	89,672	11,910
x86_64 2 x Dual Core Opteron 2GHz, 12GB	OWL-Horst	(8000, 0)	1068	129,866	8,224
E5420 @ 2.50GHz(4HT) 2GHz, 32GB	none	(20000, 0)	2760	62,003	44,516
E5420 @ 2.50GHz(4HT) 2GHz, 56GB	none	(50000, 0)	6675	151,997	43,914
E5420 @ 2.50GHz(4HT) 2GHz, 56GB	OWL-Horst	(25000, 0)	3336	302,124	11,042

**Table 1 LUBM loading with BigTRREE 3.0.x**

Table 1 indicates that if no reasoning is used BigTRREE 3.0.x is able to sustain a continuous speed of few thousands RDF statements per seconds even for multi-billion data sets. The degrade in the average loading speed between LUBM(1000, 0) and LUBM(8000, 0) (a 8 times data set increase) is less than 4%. Beyond 1 billion statements the performance decreases faster and for LUBM(50000, 0) it drops to nearly 20% compared to LUBM(1000, 0). As conclusion BigTRREE 3.0.x demonstrates excellent scalability up to 1 billion RDF statements and very good scalability beyond if no reasoning is used.


<sup>1</sup> 2GHz indicates the clock speed of the RAM

<sup>2</sup> For LUBM(1000, 0) the minimal required JVM memory is used



**Figure 1 BigTRREE loading speed for LUBM 1000, 8000, 20000 and 50000**

The runs to use OWL-Horst inference indicate a faster drop in the performance with the increase of the data set and a much slower performance. As discussed in the performance factors section the phenomenon is due the need to materialize additional implicit statements and to execute forward-chaining inference rules.



**Figure 2 BigTRREE loading speed for LUBM 1000, 8000 and 25000**

#### 4.1.2 Yet Another RDF Store

YARS2 can, according to standard settings, index approx. 150M statements per 1GB of main memory. However, this can be adjusted in either direction with a trade off on speed of lookups vs. main memory used. In [1], we evaluated YARS2 using a dataset with similar size to LUBM(50000). The most efficient set of rules, as identified in [2] was used in [3] to infer 1.925b new statements from 1.1b input statements collected from 6.5m sources on the Web in a total time of ~16.3 hours. Such is completely linear and

would be in fact faster for execution for more sanitary datasets such as the controlled LUBM test-bed. However, *intersectionOf*, *transitivity* and *someValuesFrom* reasoning required for full LUBM support are not included in the optimised ruleset. Such can be supported at a greater cost. As such, SAOR is designed for operation on Web data and has only been evaluated in this regard. LUBM benchmarks have not been performed.

### 4.1.3 Comparison with public scalability results of other systems

In the recent years there is increased interest towards the evaluation of scalable triple stores. This section contains references to deployed system capable of loading LUBM test.

**AllegroGraph:** “AllegroGraph RDFStore11 is a RDF database with support for “SPARQL, RDFS++, and Prolog reasoning from Java applications”, [9]. The store provides specific support of geospatial data and social network analysis. The system is capable of returning complete results on LUBM(50) benchmark with their RDFS++ reasoning algorithm, [10].

**Virtuoso** “is a high-performance object-relational SQL database offering diverse data and metadata management facilities” like XML management, RDBMS integration, full-text indexing and RDF support. The experiments of benchmarking Virtuoso 5.5 under LUBM(8000,0) are presented in [11].

**Oracle** is major vendor of RDBMS systems, where the RDF support is added since version 10g R2. In 11g the engine is significantly improved in terms of performance and reasoning. OWL-Prime is fragment similar to OWL-Horst. ORACLE 11g LUBM results are published in [12] and [13].

Semantic Repository	LUBM parameters	Scale in millions	Speed St/sec	Time (hours)	Semantics	Named graphs
BigTRREE 3.0b1	1000, 0	138	12,482	3.22	OWL-Horst	+
BigTRREE 3.1	8000, 0	1,068	11,910	23.15	OWL-Horst	+
BigTREEE 3.0b1	20000, 0	2,760	10,534	17.22	none	+
BigTREEE 3.0b1	25000, 0	3,336	10,534	83.92	OWL-Horst	+
BigTREEE 3.0b1	50000, 0	6,675	10,534	42.22	none	+
AllegroGraph 3.0 Fed.	8000, 0	1,107	37,273	8.25	none	+
Openlink Virtuoso, 5.0	8000, 0	1,068	36,900	8.04	none	+
ORACLE 11g	1000, 0	138	7,152	10.72	OWL-Prime	+
ORACLE 11g	1000, 0	1,068	6,667	89.00	OWL-Prime	+
DAML DB	8000, 0	850	10,266	23.00	OWL-Horst	-

<b>DAML DB</b>	8000, 0	220	6,111	10.00	OWL-Horst	-
<b>YARS2 Federated</b>	50000,0 <sup>3</sup>	~ 7,000	-	? <sup>4</sup>	none	+

**Table 2 State-of-the-art in large scale LUBM loading**

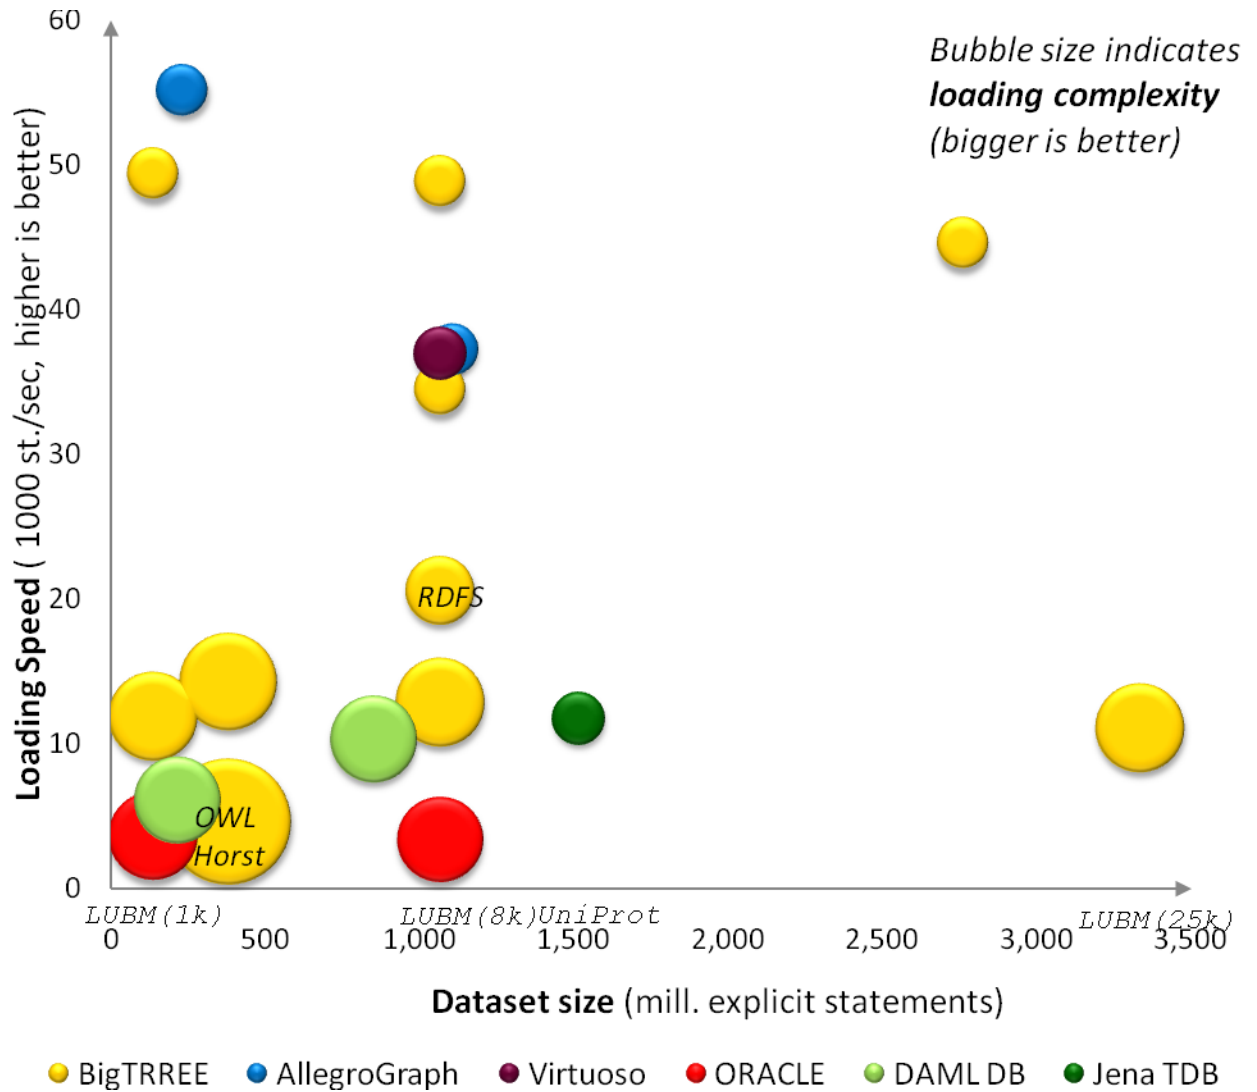
The performance is measured for loading and for query evaluation – two tasks which provide indication about the performance regarding all the activities and thus provide a consistent picture about the efficiency of a specific tool or approach. Without such complete picture, there is no empirical proof for the performance regarding some of the activities. For instance, inference can take place both at load or at query time or can be split into different phases between those two modalities of usage. There is no perfect inference schema or configuration for semantic repository – different approaches provide optimal results in different situations; forward-chaining on top of very dynamic data sets with large “inferred closure” is suboptimal. The same applies to indexing – loading can be very efficient if little or no indices are created and stored; this however will have impact on the query performance.

When evaluating the performance of a triple stores or a specific configuration, one should pay most attention to the results from full-cycle-benchmarks, which allow consistent interpretation of its advantages and the disadvantages. If for instance, only the results of loading a LUBM dataset are presented, without query performance evaluation, there is no trustworthy indication about the complexity of the inference and indexing involved in this task.

Figure 1 presents state-of-the-art in the LUBM loading and the best performing and scaling systems. The bubble size indicates the loading complexity or the amount of pre-processing during loading (e.g. materialization of inferred statements). For better readability the LUBM(50000, 0) results are removed from the picture.

<sup>3</sup> The dataset has similar volume to LUBM(50000) but it's different

<sup>4</sup> The loading process took few days and required some manual configurations



**Figure 3 State-of-the-art in LUBM loading**



## 5 QUERY EXECUTION PERFORMANCE

LUBM is widely used benchmark to measure the loading performance, the soundness and the completeness of RDF data persistence. While it provides a good basis for evaluation of query performance with respect to inference, it is not designed as a comprehensive test with respect to many of the more classical aspects of query evaluation in DBMS. For instance, LUBM fails to address phenomena such as “warm databases” (data and query result caching) and multiple simultaneous clients of the database engine. For this reason, query evaluation performance here is mostly reported on the basis of the BSBM benchmark. Results for BigTREE’s query performance under LUBM can be found in [16].

### 5.1 Berlin SPARQL Benchmark (BSBM)

According its authors, BSBM benchmark is designed to compare triple store implementation that exposes a SPARQL endpoint with realistic workloads of e-commerce use case motivated queries, [18]. The queries simulate the navigation and exploring of products by one or multiple end-users. The benchmark supports the generation different scales and data sets:

<b>Number of Producers</b>	<b>14</b>	<b>60</b>	<b>1422</b>	<b>5,618</b>
<b>Number of Product Features</b>	2,860	4,745	23,833	47,884
<b>Number of Product Types</b>	55	151	731	2011
<b>Number of Vendors</b>	8	34	722	2,854
<b>Number of Offers</b>	13,320	55,700	1,416,240	5,696,520
<b>Number of Reviewers</b>	339	1432	36,249	146,054
<b>Number of Reviews</b>	6,660	27,850	708,120	2,848,260
<b>Total Number of Instances</b>	23,922	92,757	2,258,129	9,034,027
<b>Number of RDF Statements</b>	250K	1M	25M	100M

Table 3 Berlin SPARQL Benchmark metrics, [14]

The query evaluation is measured by the number of the executed query mixes per hour (QMpH). The tested snapshot version of YARS2 is able to process only 6 of the 12 reference queries.

Query	1	4	5	6	9	11
Times Executed	128	128	128	128	128	128
AQET <sup>5</sup> [sec]	3.99963	3.964045	5.32727	1.453865	0.00924	0.010277
AQET(geom.) [sec]	3.99894	3.963239	5.312318	1.453844	0.00826	0.009947
QPS <sup>6</sup>	0.25	0.25	0.19	0.69	108.23	97.31
minQET [sec]	3.86558	3.838074	4.23802	1.411688	0.00668	0.007461
maxQET [sec]	4.24874	4.178616	5.891759	1.482375	0.03291	0.030276
Average result [count/bytes]	0.66	0	3.02	1.02	0	10
Min result [count/bytes]	0	0	0	1	0	10
Max result count [count/bytes]	6	0	5	2	0	10
Timeouts	0	0	0	0	0	0

Table 4 YARS2 detailed BSBM results (250K statements dataset)

For the evaluation of BigTRREE 3.1.x the data sets of 250K and 1M RDF statements are skipped. In Table 6 are presented the summary results of BSBM of 25M with 1, 4, 8 and 16 clients. All tests are performed on E5420 @ 2.50GHz (4HT) 2GHz with JVM 1.6.

Scale factor	70812	70812	70812	70812
Number of warmup runs	32	32	32	32
Number of clients	1	4	8	16
Seed	808080	808080	808080	808080
Number of query mix runs (without warmups)	128	128	128	128
min Querymix runtime	2.794	8.2124	8.9149	15.453
max Querymix runtime	4.2249	18.0242	24.4157	56.5178
Total runtime (sum)	419.814	1547.436	2042.548	4209.545
Total runtime (actual)	419.814	389.65	258.67	271.285
QMpH	1097.63	1182.6	1781.42	1698.58
CQET	3.2798	12.08934	15.9574	32.88707
CQET (geom.)	3.26776	11.91046	15.71952	32.1926

Table 5 BigTRREE Berlin SPARQL Benchmark summary results (25M statements data set)

The results in Table 7 confirm the tendencies for the smaller data set.

<sup>5</sup> AQET stands for Average Query Execution Time

<sup>6</sup> QPS stands for Queries per Second

Scale factor	284826	284826	284826	284826
Number of warmup runs	32	32	32	32
Number of clients	1	4	8	16
Seed	808080	808080	808080	808080
Number of query mix runs (without warmups)	128	128	128	128
min Querymix runtime	4.7606	19.2003	22.687	30.8348
max Querymix runtime	15.073	56.1505	54.1392	126.2152
Total runtime (sum)	715.257	3854.427	4385.675	8365.895
Total runtime (actual)	715.257	970.792	554.324	537.44
QMpH	644.24	474.66	831.28	857.4
CQET	5.58794	30.11271	34.26309	65.35855
CQET (geom.)	5.53512	29.55032	33.42904	63.47983

Table 6 BigTRREE Berlin SPARQL Benchmark summary results (100M statements data set)

The BSBM test demonstrates a more complex dependency of the scalability by the number of the clients and data set size.




Figure 4 BigTRREE 3.1 executed query mixes per hour for multiple clients

### 5.1.1 Comparison with public results of other systems

The public BSBM results indicate similar scalability dependencies for Sesame Native beta2 and Virtuoso 5.0.9-snapshot (Triple Store). The reports do not include the tests 100M with 4, 8 and 16 clients.

Scale factor	70812	70812	70812	70812	284826
Number of warmup runs	32	32	32	32	32
Number of clients	1	4	8	16	1
Seed	808080	863528	888326	975932	808080
Number of query mix runs (without warmups)	128	128	128	128	128
min Querymix runtime	2.8671	6.3815	18.0862	23.3588	11.4085
max Querymix runtime	5.3755	16.7253	31.3221	62.0002	17.32
Total runtime (sum)	511.013	1654.524	3058.967	6068.448	1830.22
Total runtime (actual)	512.32	416.057	387.238	386.606	1831.57
QMpH	899.44	1107.54	1189.97	1191.91	251.59
CQET	3.99229	12.92597	23.89818	47.40975	14.2986
CQET (geom.)	3.96589	12.84624	23.75684	46.9392	14.2411

Table 7 Sesame Native Berlin SPARQL Benchmark summary results, [14]

Virtuoso 5.0.9-snapshot indicates a better performance on the bigger data sets compared to Sesame:

Scale factor	70812	70812	70812	70812	284826
Number of warmup runs	32	32	32	32	32
Number of clients	1	4	8	16	1
Seed	808080	863528	888326	975932	808080
Number of query mix runs (without warmups)	128	128	128	128	128
min Querymix runtime	0.9894	1.7055	2.1738	5.2109	3.4652
max Querymix runtime	2.2318	3.9565	11.367	23.1092	136.174
Total runtime (sum)	178.373	325.291	599.38	1500.384	774.591
Total runtime (actual)	179.775	82.357	76.745	96.808	775.89
QMpH	2563.2	5595.17	6004.32	4759.92	593.9
CQET	1.39354	2.54134	4.68266	11.72175	6.05149
CQET (geom.)	1.37896	2.49754	4.4735	11.39683	5.11118

Table 8 Virtuoso public BSBM summary results, [14]

On Figure 5 and Figure 6 a comparison between BigTRREE 3.1 and other triple stores is presented. The results are not indicative because they are performed on similar class hardware, but not identical. They support the WP1 objectives to implement a storage system focussed in the scalability of large data sets. Another reason for the lower results in the multi-client scenarios is the lack of a query plan caching in BigTRREE. With the increase of the data set and decrease the number of queries (e.g. clients) the disadvantage disappears.




Figure 5 Total runtime (sum) to complete BSBM for 25M data set (1, 4, 8 and 16 clients) and 100M (1 client)




Figure 6 Executed query mixes per hour for 25M data set (1, 4, 8 and 16 clients) and 100M (1 client)

## 6 DISCUSSIONS

TripCom storage infrastructure supports the configuration of different implementation. In this chapter we present a general discussion for the tradeoffs between parallelisation and distributions. While most of the semantic store providers count on improved performance and scalability through distribution of the data across multiple machines, the economic appropriateness of such approach is still not obvious. The major problem when a DBMS should manage data spread across several machines is known as “remote join”. In essence, this is a situation when the engine running on one machine should perform large number of requests to an engine running on different machines. Given that remote calls, across processes and machines, are considerably slower than inter-process calls the worst case performance penalties for distributed-data solution are enormous. These penalties were empirically observed, measured and analysed in [17]. On the other hand, it seems that volumes of data up to 10 billion statements can be handled in comfort by a store like BigTRREE working on machine with 64GB of RAM. Such configuration with 8-16 CPU cores can be assembled at a cost below \$10,000 and has the potential to deliver considerably better price-performance ratio, compared to a distributed solution. We find such approach more promising as a mid-term strategy for scaling up the storage layer of TripCom, where the web-scale data distribution is handled in a much more comprehensive manner at higher levels of the architecture. The key for efficient usage of multi-core solution is development of parallelisation algorithms for query evaluation and reasoning, which impose minimal cost for synchronisation across the threads.

## 7 CONCLUSION

A high-performance and scalable storage system is a key prerequisite to achieve the goals of Triple Space computing. The document identifies two common tasks automated by the RDF stores 1) ontology loading/pre-processing and 2) query evaluation. LUBM benchmark is used to measure the ontology loading and pre-processing performance for data sets up to 6,675 billion of statements for BigTRREE system with native support of ORDI data model and 7 billion statements for YARS2 using distributed environment. Another LUBM system deployment with 3,336 billion explicit statements and OWL-Horst reasoning is realized and able to return a complete query result. BSBM benchmark indicates state-of-the-art performance for SPARQL query execution over large data set. As conclusion WP1 implemented a configurable storage infrastructure capable of processing billions of RDF statement in a scalable and high-performance way.

**REFERENCES**

- [1] Andreas Harth, Jürgen Umbrich, Aidan Hogan, Stefan Decker. "YARS2: A Federated Repository for Searching and Querying Graph Structured Data". In Proceedings of 6th International Semantic Web Conference (ISWC2007), Busan, Korea, 2007.
- [2] Aidan Hogan, Andreas Harth, Axel Polleres. "SAOR: Authoritative Reasoning for the Web". In Proceedings of the 3rd Asian Semantic Web Conference (ASWC2008), Bangkok, Thailand, 2008 (to appear).
- [3] Aidan Hogan, Andreas Harth, Axel Polleres. "Scalable Authoritative OWL Reasoning on a Billion Triples". In Proceedings of Billion Triple Semantic Web Challenge 2008, at the 7th International Semantic Web Conference (ISWC2008), Karlsruhe, Germany, 2008.
- [4] Vassil Momtchev and Atanas Kiryakov. D1.2 Specification of the Store Architecture and Interfaces. TripCom Deliverable, October 2006. <http://tripcom.org/docs/del/Tripcom-D12.pdf>
- [5] Guo, Y; Pan, Z; and Heflin, J. (2004). An Evaluation of Knowledge Base Systems for Large OWL Datasets. Journal of Web Semantics, 3(2), 2005, pp. 158-182. <http://www.websemanticsjournal.org/ps/pub/2005-16>
- [6] André B. Bondi, 'Characteristics of scalability and their impact on performance', Proceedings of the 2nd international workshop on Software and performance, Ottawa, Ontario, Canada, 2000, ISBN 1-58113-195-X, pages 195 – 203
- [7] LargeTripleStores, <http://esw.w3.org/topic/LargeTripleStores>
- [8] Base One. "Database Scalability - Dispelling myths about the limits of database-centric architecture", 2007. Retrieved on May 23, 2007.
- [9] <http://agraph.franz.com/allegrograph/>
- [10] Franz Inc. (2008). AllegroGraph 3.0 Benchmark for LUBM-50-0. [http://agraph.franz.com/allegrograph/agraph\\_bench\\_lubm50.lhtml](http://agraph.franz.com/allegrograph/agraph_bench_lubm50.lhtml) (as of 16 Jun, 2008).
- [11] Erling, O. (OpenLink Software). (2008). LUBM and Virtuoso. <http://virtuoso.openlinksw.com/wiki/main/Main/VOSArticleLUBMBenchmark>. 1 Feb, 2008.
- [12] Oracle Corp. (2008). Oracle Semantic Technologies Inference Best Practices with RDFS/OWL. An Oracle White Paper. February 2008.
- [13] Oracle Corp. (2008). Semantic Technologies Product Performance. [http://www.oracle.com/technology/tech/semantic\\_technologies/htdocs/performance.html](http://www.oracle.com/technology/tech/semantic_technologies/htdocs/performance.html), as of 18th of June, 2008.
- [14] <http://www4.wiiss.fu-berlin.de/bizer/BerlinSPARQLBenchmark/results/index.html>



- [15] Weithoner, M.L.T., Liebig, T., Luther, M., Bohm, S.: What's Wrong with OWL Benchmarks? Proc. of the Second Int. Workshop on Scalable Semantic Web Knowledge Base Systems (SSWS 2006) 101-114
- [16] Ontotext AD. (2008). OWLIM – Pragmatic OWL Semantic Repository. Presentation, Nov 2008. <http://www.ontotext.com/owlim/OWLIMPres.pdf>
- [17] Erling, Orri. (2008). Towards Web Scale RDF. In Proc. 4th International Workshop on Scalable Semantic Web Knowledge Base Systems (SSWS2008). <http://virtuoso.openlinksw.com/dataspace/dav/wiki/Main/VOSArticles/VOSArticleWebScaleRDF.pdf>
- [18] <http://www4.wiwiss.fu-berlin.de/bizer/BerlinSPARQLBenchmark/spec/index.html>

**APPENDIX A**

## BigTREE 3.0b5 LUBM (25000, 0) query time

<b>Total</b>	<b>349,235</b>	<b>33,274,986</b>	<b>36.54</b>
<b>Query 1</b>	47	4	11.75
<b>Query 2</b>	1,107,961	2,528	438.28
<b>Query 3</b>	41	6	6.83
<b>Query 4</b>	975	34	28.68
<b>Query 5</b>	68	719	0.09
<b>Query 6</b>	300,515	260,959,126	0.00
<b>Query 7</b>	165	67	2.46
<b>Query 8</b>	890	7,790	0.11
<b>Query 9</b>	2,689,860	6,801,440	0.40
<b>Query 10</b>	44	4	11.00
<b>Query 11</b>	37	224	0.17
<b>Query 12</b>	140	15	9.33
<b>Query 13</b>	283,087	116,707	2.43
<b>Query 14</b>	505,453	197,961,139	0.00

**APPENDIX B**

## YARS2 250K BSBM results (single client)

Reading Test Driver data...done

Reading query mix file: querymix.txt

Reading query ignore file: ignoreQueries.txt

Starting test...

```
-32: 0.00ms, total: 15381ms
-31: 0.00ms, total: 14770ms
-30: 0.00ms, total: 14916ms
-29: 0.00ms, total: 14346ms
-28: 0.00ms, total: 15018ms
-27: 0.00ms, total: 14452ms
-26: 0.00ms, total: 14452ms
-25: 0.00ms, total: 14542ms
-24: 0.00ms, total: 14439ms
-23: 0.00ms, total: 15147ms
-22: 0.00ms, total: 14993ms
-21: 0.00ms, total: 14795ms
-20: 0.00ms, total: 13623ms
-19: 0.00ms, total: 14768ms
-18: 0.00ms, total: 14875ms
-17: 0.00ms, total: 14162ms
-16: 0.00ms, total: 15488ms
-15: 0.00ms, total: 14744ms
-14: 0.00ms, total: 14786ms
-13: 0.00ms, total: 15358ms
-12: 0.00ms, total: 14665ms
-11: 0.00ms, total: 14159ms
-10: 0.00ms, total: 14975ms
-9: 0.00ms, total: 14835ms
-8: 0.00ms, total: 14847ms
-7: 0.00ms, total: 15097ms
-6: 0.00ms, total: 15069ms
-5: 0.00ms, total: 14951ms
-4: 0.00ms, total: 14544ms
-3: 0.00ms, total: 14826ms
-2: 0.00ms, total: 13482ms
-1: 0.00ms, total: 14582ms
0: 15297.28ms, total: 15306ms
1: 14845.59ms, total: 14858ms
2: 14814.02ms, total: 14820ms
3: 14948.99ms, total: 14957ms
4: 14189.51ms, total: 14197ms
5: 15255.32ms, total: 15272ms
6: 14473.37ms, total: 14482ms
7: 15069.98ms, total: 15080ms
8: 14460.57ms, total: 14468ms
9: 15518.00ms, total: 15526ms
10: 14694.10ms, total: 14700ms
11: 14610.85ms, total: 14618ms
12: 14064.48ms, total: 14070ms
13: 14601.37ms, total: 14609ms
14: 15026.44ms, total: 15035ms
15: 13969.26ms, total: 13977ms
16: 14486.06ms, total: 14496ms
17: 14951.26ms, total: 14958ms
18: 15042.16ms, total: 15054ms
```

19: 14730.48ms, total: 14738ms  
20: 15346.98ms, total: 15358ms  
21: 13605.05ms, total: 13612ms  
22: 14675.98ms, total: 14683ms  
23: 14879.47ms, total: 14890ms  
24: 15116.34ms, total: 15123ms  
25: 14345.25ms, total: 14353ms  
26: 14334.15ms, total: 14340ms  
27: 14751.34ms, total: 14758ms  
28: 14640.75ms, total: 14649ms  
29: 14741.81ms, total: 14747ms  
30: 14902.25ms, total: 14910ms  
31: 15099.24ms, total: 15108ms  
32: 14608.37ms, total: 14616ms  
33: 14630.91ms, total: 14637ms  
34: 14553.21ms, total: 14573ms  
35: 14888.88ms, total: 14900ms  
36: 14985.27ms, total: 14995ms  
37: 15188.38ms, total: 15195ms  
38: 14492.27ms, total: 14500ms  
39: 15243.45ms, total: 15251ms  
40: 13646.10ms, total: 13654ms  
41: 15050.64ms, total: 15057ms  
42: 15162.85ms, total: 15169ms  
43: 14776.98ms, total: 14784ms  
44: 14162.27ms, total: 14169ms  
45: 15093.45ms, total: 15101ms  
46: 14998.09ms, total: 15005ms  
47: 15406.10ms, total: 15411ms  
48: 15462.39ms, total: 15469ms  
49: 14905.50ms, total: 14912ms  
50: 14189.60ms, total: 14198ms  
51: 14436.63ms, total: 14446ms  
52: 14675.78ms, total: 14683ms  
53: 15132.63ms, total: 15140ms  
54: 14581.73ms, total: 14589ms  
55: 15193.07ms, total: 15201ms  
56: 14461.14ms, total: 14468ms  
57: 14461.22ms, total: 14468ms  
58: 14128.55ms, total: 14139ms  
59: 14409.76ms, total: 14417ms  
60: 14961.56ms, total: 14968ms  
61: 15170.17ms, total: 15178ms  
62: 14536.81ms, total: 14546ms  
63: 15112.66ms, total: 15120ms  
64: 14930.63ms, total: 14938ms  
65: 14263.62ms, total: 14271ms  
66: 15007.38ms, total: 15014ms  
67: 15251.06ms, total: 15259ms  
68: 15403.07ms, total: 15413ms  
69: 14551.06ms, total: 14559ms  
70: 14959.18ms, total: 15056ms  
71: 14692.08ms, total: 14699ms  
72: 14080.72ms, total: 14087ms  
73: 14651.23ms, total: 14819ms  
74: 14133.34ms, total: 14140ms  
75: 15120.68ms, total: 15160ms  
76: 14928.60ms, total: 14935ms  
77: 15347.78ms, total: 15354ms  
78: 13679.89ms, total: 13687ms  
79: 13902.89ms, total: 13911ms  
80: 14808.87ms, total: 14816ms  
81: 15309.40ms, total: 15317ms

82: 15142.94ms, total: 15150ms  
 83: 14187.91ms, total: 14194ms  
 84: 14683.86ms, total: 14693ms  
 85: 14716.63ms, total: 14723ms  
 86: 14870.05ms, total: 14891ms  
 87: 14784.14ms, total: 14791ms  
 88: 15046.43ms, total: 15053ms  
 89: 13635.18ms, total: 13676ms  
 90: 14770.49ms, total: 14776ms  
 91: 15110.26ms, total: 15117ms  
 92: 15341.24ms, total: 15348ms  
 93: 15183.02ms, total: 15193ms  
 94: 14865.13ms, total: 14872ms  
 95: 15152.78ms, total: 15159ms  
 96: 14663.08ms, total: 14670ms  
 97: 14108.80ms, total: 14136ms  
 98: 15270.03ms, total: 15277ms  
 99: 14858.92ms, total: 14866ms  
 100: 15264.04ms, total: 15314ms  
 101: 15028.58ms, total: 15038ms  
 102: 14378.20ms, total: 14389ms  
 103: 14797.61ms, total: 14804ms  
 104: 15105.98ms, total: 15125ms  
 105: 14733.89ms, total: 14743ms  
 106: 15292.78ms, total: 15298ms  
 107: 14670.24ms, total: 14677ms  
 108: 14830.03ms, total: 14873ms  
 109: 14659.75ms, total: 14669ms  
 110: 14752.33ms, total: 14779ms  
 111: 14309.64ms, total: 14316ms  
 112: 14445.99ms, total: 14454ms  
 113: 14483.92ms, total: 14568ms  
 114: 15040.85ms, total: 15048ms  
 115: 15030.95ms, total: 15037ms  
 116: 14214.61ms, total: 14223ms  
 117: 14936.69ms, total: 14942ms  
 118: 14838.87ms, total: 14854ms  
 119: 14648.61ms, total: 14655ms  
 120: 14243.94ms, total: 14252ms  
 121: 15386.15ms, total: 15394ms  
 122: 14106.25ms, total: 14112ms  
 123: 15053.53ms, total: 15060ms  
 124: 14939.42ms, total: 14945ms  
 125: 14914.77ms, total: 14922ms  
 126: 15021.96ms, total: 15029ms  
 127: 15097.57ms, total: 15104ms

Scale factor: 666  
 Number of warmup runs: 32  
 Seed: 808080  
 Number of query mix runs (without warmups): 128 times  
 min/max Querymix runtime: 13.6051s / 15.5180s  
 Total runtime: 1889.834 seconds  
 QMPH: 243.83 query mixes per hour  
 CQET: 14.76433 seconds average runtime of query mix  
 CQET (geom.): 14.75838 seconds geometric mean runtime of query mix

Metrics for Query: 1  
 Count: 128 times executed in whole run  
 AQET: 3.999629 seconds (arithmetic mean)  
 AQET(geom.): 3.998940 seconds (geometric mean)  
 QPS: 0.25 Queries per second  
 minQET/maxQET: 3.86558151s / 4.24873695s

```

Average result count: 0.66
min/max result count: 0 / 6
Number of timeouts: 0

Metrics for Query: 4
Count: 128 times executed in whole run
AQET: 3.964045 seconds (arithmetic mean)
AQET(geom.): 3.963239 seconds (geometric mean)
QPS: 0.25 Queries per second
minQET/maxQET: 3.83807443s / 4.17861580s
Average result count: 0.00
min/max result count: 0 / 0
Number of timeouts: 0

Metrics for Query: 5
Count: 128 times executed in whole run
AQET: 5.327270 seconds (arithmetic mean)
AQET(geom.): 5.312318 seconds (geometric mean)
QPS: 0.19 Queries per second
minQET/maxQET: 4.23802048s / 5.89175916s
Average result count: 3.02
min/max result count: 0 / 5
Number of timeouts: 0

Metrics for Query: 6
Count: 128 times executed in whole run
AQET: 1.453865 seconds (arithmetic mean)
AQET(geom.): 1.453844 seconds (geometric mean)
QPS: 0.69 Queries per second
minQET/maxQET: 1.41168790s / 1.48237489s
Average result count: 1.02
min/max result count: 1 / 2
Number of timeouts: 0

Metrics for Query: 9
Count: 128 times executed in whole run
AQET: 0.009239 seconds (arithmetic mean)
AQET(geom.): 0.008263 seconds (geometric mean)
QPS: 108.23 Queries per second
minQET/maxQET: 0.00667593s / 0.03290882s
Average result (Bytes): 0.00
min/max result (Bytes): 0 / 0
Number of timeouts: 0

Metrics for Query: 11
Count: 128 times executed in whole run
AQET: 0.010277 seconds (arithmetic mean)
AQET(geom.): 0.009947 seconds (geometric mean)
QPS: 97.31 Queries per second
minQET/maxQET: 0.00746126s / 0.03027585s
Average result count: 10.00
min/max result count: 10 / 10
Number of timeouts: 0

```

**APPENDIX C****BigTRREE 3.1.x 100M BSBM results (1, 4, 8 and 16 clients)**

Reading Test Driver data...done

Starting test...

```
-32: 0.00ms, total: 11206ms
-31: 0.00ms, total: 6825ms
-30: 0.00ms, total: 5911ms
-29: 0.00ms, total: 5982ms
-28: 0.00ms, total: 5771ms
-27: 0.00ms, total: 6061ms
-26: 0.00ms, total: 5921ms
-25: 0.00ms, total: 6056ms
-24: 0.00ms, total: 6442ms
-23: 0.00ms, total: 5766ms
-22: 0.00ms, total: 6454ms
-21: 0.00ms, total: 5636ms
-20: 0.00ms, total: 5687ms
-19: 0.00ms, total: 5069ms
-18: 0.00ms, total: 5363ms
-17: 0.00ms, total: 5724ms
-16: 0.00ms, total: 6625ms
-15: 0.00ms, total: 5823ms
-14: 0.00ms, total: 5747ms
-13: 0.00ms, total: 6021ms
-12: 0.00ms, total: 5046ms
-11: 0.00ms, total: 5100ms
-10: 0.00ms, total: 5961ms
-9: 0.00ms, total: 5792ms
-8: 0.00ms, total: 5981ms
-7: 0.00ms, total: 4983ms
-6: 0.00ms, total: 5657ms
-5: 0.00ms, total: 5889ms
-4: 0.00ms, total: 5460ms
-3: 0.00ms, total: 5515ms
-2: 0.00ms, total: 5241ms
-1: 0.00ms, total: 5500ms
0: 6593.20ms, total: 6612ms
1: 5395.68ms, total: 5409ms
2: 5254.64ms, total: 5270ms
3: 6144.29ms, total: 6155ms
4: 5350.58ms, total: 5362ms
5: 5340.49ms, total: 5353ms
6: 6129.68ms, total: 6142ms
7: 4962.93ms, total: 4974ms
8: 5069.18ms, total: 5081ms
9: 5011.88ms, total: 5023ms
10: 6522.13ms, total: 6534ms
11: 5376.84ms, total: 5387ms
12: 5941.63ms, total: 5999ms
13: 15072.97ms, total: 15084ms
14: 5320.71ms, total: 5331ms
15: 5196.59ms, total: 5207ms
16: 5110.67ms, total: 5121ms
17: 4918.80ms, total: 4929ms
18: 5520.87ms, total: 5533ms
19: 5572.65ms, total: 5586ms
20: 5094.91ms, total: 5106ms
```

21: 6006.84ms, total: 6017ms  
22: 6174.38ms, total: 6186ms  
23: 5799.24ms, total: 5810ms  
24: 4997.85ms, total: 5008ms  
25: 5388.70ms, total: 5400ms  
26: 6181.84ms, total: 6193ms  
27: 5532.08ms, total: 5543ms  
28: 4760.61ms, total: 4772ms  
29: 5357.57ms, total: 5369ms  
30: 5496.37ms, total: 5507ms  
31: 6465.89ms, total: 6476ms  
32: 5928.03ms, total: 5939ms  
33: 5094.49ms, total: 5105ms  
34: 5311.72ms, total: 5323ms  
35: 5887.26ms, total: 5898ms  
36: 5754.83ms, total: 5767ms  
37: 5523.20ms, total: 5534ms  
38: 5435.45ms, total: 5446ms  
39: 5192.60ms, total: 5204ms  
40: 6260.11ms, total: 6271ms  
41: 5188.42ms, total: 5199ms  
42: 5189.39ms, total: 5199ms  
43: 4833.96ms, total: 4845ms  
44: 5022.93ms, total: 5033ms  
45: 6700.52ms, total: 6711ms  
46: 5498.00ms, total: 5509ms  
47: 5886.09ms, total: 5896ms  
48: 5107.31ms, total: 5117ms  
49: 5012.28ms, total: 5024ms  
50: 5721.36ms, total: 5731ms  
51: 5475.67ms, total: 5486ms  
52: 5154.07ms, total: 5169ms  
53: 4936.84ms, total: 4947ms  
54: 5315.56ms, total: 5332ms  
55: 5138.02ms, total: 5150ms  
56: 6393.26ms, total: 6404ms  
57: 5121.44ms, total: 5132ms  
58: 5910.45ms, total: 5926ms  
59: 5028.65ms, total: 5039ms  
60: 5237.17ms, total: 5251ms  
61: 6299.98ms, total: 6310ms  
62: 6584.50ms, total: 6595ms  
63: 5339.57ms, total: 5349ms  
64: 5184.92ms, total: 5195ms  
65: 4868.55ms, total: 4879ms  
66: 5355.34ms, total: 5370ms  
67: 5507.98ms, total: 5519ms  
68: 6999.74ms, total: 7015ms  
69: 5210.56ms, total: 5221ms  
70: 5115.11ms, total: 5130ms  
71: 5022.41ms, total: 5033ms  
72: 5322.19ms, total: 5337ms  
73: 5713.86ms, total: 5726ms  
74: 5447.18ms, total: 5462ms  
75: 6440.32ms, total: 6452ms  
76: 5836.13ms, total: 5851ms  
77: 5013.71ms, total: 5025ms  
78: 5943.58ms, total: 5954ms  
79: 5302.99ms, total: 5314ms  
80: 5217.32ms, total: 5228ms  
81: 6221.03ms, total: 6237ms  
82: 5200.44ms, total: 5210ms  
83: 5068.02ms, total: 5078ms



84: 5906.75ms, total: 5917ms  
 85: 5197.03ms, total: 5207ms  
 86: 4878.95ms, total: 4890ms  
 87: 6147.02ms, total: 6158ms  
 88: 6191.82ms, total: 6206ms  
 89: 5182.89ms, total: 5194ms  
 90: 6345.19ms, total: 6359ms  
 91: 4952.91ms, total: 5006ms  
 92: 5084.72ms, total: 5099ms  
 93: 6105.19ms, total: 6116ms  
 94: 4801.59ms, total: 4816ms  
 95: 6128.51ms, total: 6139ms  
 96: 5082.98ms, total: 5098ms  
 97: 5172.31ms, total: 5183ms  
 98: 5310.32ms, total: 5325ms  
 99: 5834.84ms, total: 5845ms  
 100: 5248.39ms, total: 5258ms  
 101: 4910.45ms, total: 4924ms  
 102: 5644.55ms, total: 5655ms  
 103: 5908.59ms, total: 5923ms  
 104: 5010.41ms, total: 5022ms  
 105: 5192.73ms, total: 5204ms  
 106: 6029.17ms, total: 6039ms  
 107: 5833.79ms, total: 5849ms  
 108: 4791.57ms, total: 4802ms  
 109: 5236.49ms, total: 5251ms  
 110: 5097.81ms, total: 5109ms  
 111: 6576.71ms, total: 6591ms  
 112: 5084.25ms, total: 5095ms  
 113: 5317.48ms, total: 5332ms  
 114: 5394.89ms, total: 5406ms  
 115: 5555.21ms, total: 5570ms  
 116: 5201.20ms, total: 5212ms  
 117: 5141.73ms, total: 5156ms  
 118: 5743.96ms, total: 5756ms  
 119: 6171.65ms, total: 6187ms  
 120: 5181.19ms, total: 5191ms  
 121: 5361.79ms, total: 5376ms  
 122: 5660.84ms, total: 5672ms  
 123: 5557.92ms, total: 5568ms  
 124: 5525.27ms, total: 5536ms  
 125: 5286.46ms, total: 5301ms  
 126: 5214.20ms, total: 5225ms  
 127: 6517.65ms, total: 6533ms

Scale factor: 284826  
 Number of warmup runs: 32  
 Seed: 808080  
 Number of query mix runs (without warmups): 128 times  
 min/max Querymix runtime: 4.7606s / 15.0730s  
 Total runtime: 715.257 seconds  
 QMPH: 644.24 query mixes per hour  
 CQET: 5.58794 seconds average runtime of query mix  
 CQET (geom.): 5.53512 seconds geometric mean runtime of query mix

Metrics for Query: 1  
 Count: 128 times executed in whole run  
 AQET: 0.060056 seconds (arithmetic mean)  
 AQET(geom.): 0.045302 seconds (geometric mean)  
 QPS: 16.65 Queries per second  
 minQET/maxQET: 0.03528900s / 0.96855900s  
 Average result count: 8.00  
 min/max result count: 1 / 10

```

Number of timeouts:      0

Metrics for Query:      2
Count:                  768 times executed in whole run
AQET:                   0.189629 seconds (arithmetic mean)
AQET(geom.):            0.178799 seconds (geometric mean)
QPS:                    5.27 Queries per second
minQET/maxQET:         0.15603700s / 1.27559300s
Average result count:  19.23
min/max result count:  7 / 32
Number of timeouts:    0

Metrics for Query:      3
Count:                  128 times executed in whole run
AQET:                   0.041619 seconds (arithmetic mean)
AQET(geom.):            0.041427 seconds (geometric mean)
QPS:                    24.03 Queries per second
minQET/maxQET:         0.03615200s / 0.08649000s
Average result count:  6.88
min/max result count:  0 / 10
Number of timeouts:    0

Metrics for Query:      4
Count:                  128 times executed in whole run
AQET:                   0.118816 seconds (arithmetic mean)
AQET(geom.):            0.103081 seconds (geometric mean)
QPS:                    8.42 Queries per second
minQET/maxQET:         0.07676000s / 0.96856200s
Average result count:  7.16
min/max result count:  0 / 10
Number of timeouts:    0

Metrics for Query:      5
Count:                  128 times executed in whole run
AQET:                   0.547721 seconds (arithmetic mean)
AQET(geom.):            0.456304 seconds (geometric mean)
QPS:                    1.83 Queries per second
minQET/maxQET:         0.09615500s / 1.51848200s
Average result count:  5.00
min/max result count:  5 / 5
Number of timeouts:    0

Metrics for Query:      6
Count:                  128 times executed in whole run
AQET:                   2.715204 seconds (arithmetic mean)
AQET(geom.):            2.668257 seconds (geometric mean)
QPS:                    0.37 Queries per second
minQET/maxQET:         2.48479800s / 12.13626800s
Average result count:  247.69
min/max result count:  1 / 30401
Number of timeouts:    0

Metrics for Query:      7
Count:                  512 times executed in whole run
AQET:                   0.098076 seconds (arithmetic mean)
AQET(geom.):            0.091449 seconds (geometric mean)
QPS:                    10.20 Queries per second
minQET/maxQET:         0.05855900s / 0.97618800s
Average result count:  14.93
min/max result count:  1 / 132
Number of timeouts:    0

Metrics for Query:      8

```

```

Count:                256 times executed in whole run
AQET:                 0.055225 seconds (arithmetic mean)
AQET(geom.):         0.054557 seconds (geometric mean)
QPS:                  18.11 Queries per second
minQET/maxQET:       0.01360500s / 0.08434200s
Average result count: 4.85
min/max result count: 0 / 15
Number of timeouts:  0

```

```

Metrics for Query:    9
Count:                512 times executed in whole run
AQET:                 0.024138 seconds (arithmetic mean)
AQET(geom.):         0.024043 seconds (geometric mean)
QPS:                  41.43 Queries per second
minQET/maxQET:       0.01865800s / 0.03577100s
Average result (Bytes): 6771.43
min/max result (Bytes): 2001 / 12744
Number of timeouts:  0

```

```

Metrics for Query:    10
Count:                256 times executed in whole run
AQET:                 0.072561 seconds (arithmetic mean)
AQET(geom.):         0.067460 seconds (geometric mean)
QPS:                  13.78 Queries per second
minQET/maxQET:       0.05701100s / 0.23070300s
Average result count: 2.00
min/max result count: 0 / 10
Number of timeouts:  0

```

```

Metrics for Query:    11
Count:                128 times executed in whole run
AQET:                 0.022681 seconds (arithmetic mean)
AQET(geom.):         0.018681 seconds (geometric mean)
QPS:                  44.09 Queries per second
minQET/maxQET:       0.01657200s / 0.18872100s
Average result count: 10.00
min/max result count: 10 / 10
Number of timeouts:  0

```

```

Metrics for Query:    12
Count:                128 times executed in whole run
AQET:                 0.199641 seconds (arithmetic mean)
AQET(geom.):         0.163875 seconds (geometric mean)
QPS:                  5.01 Queries per second
minQET/maxQET:       0.09248000s / 1.44813600s
Average result (Bytes): 1425.98
min/max result (Bytes): 1397 / 1460
Number of timeouts:  0

```

Reading Test Driver data...done

Starting test...

Generating queries...done

Clients created.

```

Thread 4: query mix -30: 0.00ms, total: 69995.56ms
Thread 1: query mix -32: 0.00ms, total: 71863.18ms
Thread 2: query mix -31: 0.00ms, total: 72455.29ms
Thread 3: query mix -29: 0.00ms, total: 73535.01ms
Thread 2: query mix -26: 0.00ms, total: 25924.16ms
Thread 1: query mix -27: 0.00ms, total: 33169.12ms
Thread 4: query mix -28: 0.00ms, total: 40742.30ms
Thread 3: query mix -25: 0.00ms, total: 44655.72ms

```

Thread 1: query mix -23: 0.00ms, total: 34640.36ms  
Thread 3: query mix -21: 0.00ms, total: 28571.27ms  
Thread 2: query mix -24: 0.00ms, total: 52427.42ms  
Thread 4: query mix -22: 0.00ms, total: 50673.31ms  
Thread 2: query mix -18: 0.00ms, total: 26316.12ms  
Thread 3: query mix -19: 0.00ms, total: 34501.72ms  
Thread 1: query mix -20: 0.00ms, total: 46400.34ms  
Thread 4: query mix -17: 0.00ms, total: 44314.13ms  
Thread 3: query mix -15: 0.00ms, total: 35202.78ms  
Thread 2: query mix -16: 0.00ms, total: 44200.46ms  
Thread 1: query mix -14: 0.00ms, total: 43496.13ms  
Thread 3: query mix -12: 0.00ms, total: 30735.66ms  
Thread 4: query mix -13: 0.00ms, total: 41901.20ms  
Thread 2: query mix -11: 0.00ms, total: 33906.94ms  
Thread 1: query mix -10: 0.00ms, total: 36703.36ms  
Thread 2: query mix -7: 0.00ms, total: 23095.95ms  
Thread 3: query mix -9: 0.00ms, total: 37243.94ms  
Thread 1: query mix -6: 0.00ms, total: 28107.74ms  
Thread 4: query mix -8: 0.00ms, total: 53587.00ms  
Thread 3: query mix -4: 0.00ms, total: 32356.33ms  
Thread 2: query mix -5: 0.00ms, total: 45894.21ms  
Thread 1: query mix -3: 0.00ms, total: 30285.06ms  
Thread 4: query mix -2: 0.00ms, total: 30178.61ms  
Thread 3: query mix -1: 0.00ms, total: 22307.92ms  
Warmup phase ended...

Starting actual run...

Thread 2: query mix 2: 26158.54ms, total: 26207.74ms  
Thread 1: query mix 1: 29668.28ms, total: 29699.83ms  
Thread 4: query mix 0: 36773.96ms, total: 36799.30ms  
Thread 3: query mix 3: 41178.78ms, total: 41209.51ms  
Thread 2: query mix 4: 35360.20ms, total: 35386.20ms  
Thread 1: query mix 5: 31928.85ms, total: 31947.95ms  
Thread 3: query mix 7: 28346.52ms, total: 28362.93ms  
Thread 4: query mix 6: 44638.59ms, total: 44654.25ms  
Thread 1: query mix 9: 29450.71ms, total: 29468.63ms  
Thread 2: query mix 8: 34256.47ms, total: 34271.88ms  
Thread 3: query mix 10: 43890.46ms, total: 43906.35ms  
Thread 4: query mix 11: 44139.45ms, total: 44155.32ms  
Thread 2: query mix 13: 43981.43ms, total: 43997.98ms  
Thread 1: query mix 12: 56150.48ms, total: 56167.07ms  
Thread 4: query mix 15: 30678.63ms, total: 30695.74ms  
Thread 3: query mix 14: 46481.21ms, total: 46497.30ms  
Thread 2: query mix 16: 29966.44ms, total: 29983.79ms  
Thread 1: query mix 17: 28990.33ms, total: 29006.84ms  
Thread 4: query mix 18: 26761.85ms, total: 26778.37ms  
Thread 3: query mix 19: 35067.53ms, total: 35091.29ms  
Thread 2: query mix 20: 28707.12ms, total: 28726.22ms  
Thread 4: query mix 22: 29642.59ms, total: 29659.51ms  
Thread 1: query mix 21: 37183.77ms, total: 37201.00ms  
Thread 2: query mix 24: 24503.38ms, total: 24520.89ms  
Thread 3: query mix 23: 39617.00ms, total: 39637.27ms  
Thread 4: query mix 25: 31080.89ms, total: 31098.49ms  
Thread 1: query mix 26: 34401.87ms, total: 34422.83ms  
Thread 3: query mix 28: 26830.14ms, total: 26847.94ms  
Thread 2: query mix 27: 39314.99ms, total: 39335.24ms  
Thread 4: query mix 29: 27453.00ms, total: 27472.51ms  
Thread 1: query mix 30: 28526.38ms, total: 28544.03ms  
Thread 2: query mix 32: 31189.01ms, total: 31210.32ms  
Thread 3: query mix 31: 33810.74ms, total: 33830.11ms  
Thread 4: query mix 33: 28118.36ms, total: 28136.66ms  
Thread 1: query mix 34: 31042.28ms, total: 31060.64ms  
Thread 2: query mix 35: 29836.03ms, total: 29855.91ms

Thread 3: query mix 36: 31542.33ms, total: 31564.53ms  
Thread 4: query mix 37: 32527.06ms, total: 32543.53ms  
Thread 1: query mix 38: 35490.17ms, total: 35554.14ms  
Thread 2: query mix 39: 27509.40ms, total: 27527.29ms  
Thread 4: query mix 41: 30298.58ms, total: 30314.86ms  
Thread 3: query mix 40: 39703.49ms, total: 39719.04ms  
Thread 1: query mix 42: 31257.41ms, total: 31275.01ms  
Thread 2: query mix 43: 27903.04ms, total: 27919.26ms  
Thread 4: query mix 44: 27297.72ms, total: 27314.06ms  
Thread 3: query mix 45: 33225.02ms, total: 33242.11ms  
Thread 1: query mix 46: 29867.88ms, total: 29884.06ms  
Thread 2: query mix 47: 32088.32ms, total: 32105.54ms  
Thread 4: query mix 48: 29012.15ms, total: 29028.18ms  
Thread 3: query mix 49: 27713.56ms, total: 27730.03ms  
Thread 1: query mix 50: 29336.11ms, total: 29353.49ms  
Thread 2: query mix 51: 33427.44ms, total: 33442.59ms  
Thread 4: query mix 52: 28099.10ms, total: 28117.56ms  
Thread 3: query mix 53: 26292.79ms, total: 26308.79ms  
Thread 1: query mix 54: 30246.66ms, total: 30264.20ms  
Thread 2: query mix 55: 26672.84ms, total: 26689.21ms  
Thread 3: query mix 57: 24124.14ms, total: 24139.50ms  
Thread 4: query mix 56: 39246.04ms, total: 39264.92ms  
Thread 2: query mix 59: 23426.00ms, total: 23442.45ms  
Thread 1: query mix 58: 36852.09ms, total: 36868.60ms  
Thread 3: query mix 60: 26861.20ms, total: 26876.05ms  
Thread 4: query mix 61: 37177.71ms, total: 37193.43ms  
Thread 1: query mix 63: 26237.39ms, total: 26253.39ms  
Thread 2: query mix 62: 37126.74ms, total: 37142.37ms  
Thread 3: query mix 64: 29943.32ms, total: 29960.58ms  
Thread 4: query mix 65: 23974.95ms, total: 23990.97ms  
Thread 1: query mix 66: 26056.53ms, total: 26071.63ms  
Thread 2: query mix 67: 31986.19ms, total: 32003.64ms  
Thread 4: query mix 69: 25497.80ms, total: 25514.64ms  
Thread 3: query mix 68: 41864.53ms, total: 41886.40ms  
Thread 1: query mix 70: 28054.98ms, total: 28069.90ms  
Thread 2: query mix 71: 27164.75ms, total: 27182.05ms  
Thread 4: query mix 72: 27571.10ms, total: 27587.02ms  
Thread 3: query mix 73: 29059.57ms, total: 29075.68ms  
Thread 1: query mix 74: 30713.07ms, total: 30728.55ms  
Thread 2: query mix 75: 30184.50ms, total: 30201.66ms  
Thread 3: query mix 77: 21055.97ms, total: 21071.17ms  
Thread 4: query mix 76: 36842.03ms, total: 36858.23ms  
Thread 2: query mix 79: 27776.89ms, total: 27793.09ms  
Thread 1: query mix 78: 37301.20ms, total: 37319.81ms  
Thread 3: query mix 80: 30126.27ms, total: 30143.10ms  
Thread 4: query mix 81: 32378.26ms, total: 32400.02ms  
Thread 1: query mix 83: 25293.71ms, total: 25310.00ms  
Thread 2: query mix 82: 28799.08ms, total: 28815.64ms  
Thread 3: query mix 84: 34768.78ms, total: 34785.35ms  
Thread 4: query mix 85: 25482.65ms, total: 25499.63ms  
Thread 1: query mix 86: 25561.81ms, total: 25577.99ms  
Thread 2: query mix 87: 35054.00ms, total: 35069.49ms  
Thread 4: query mix 89: 24030.15ms, total: 24045.68ms  
Thread 3: query mix 88: 28080.29ms, total: 28100.48ms  
Thread 1: query mix 90: 33366.53ms, total: 33383.53ms  
Thread 2: query mix 91: 23882.67ms, total: 23899.72ms  
Thread 4: query mix 92: 25324.84ms, total: 25343.19ms  
Thread 1: query mix 94: 20377.59ms, total: 20394.26ms  
Thread 3: query mix 93: 38362.28ms, total: 38381.91ms  
Thread 2: query mix 95: 24757.82ms, total: 24778.07ms  
Thread 4: query mix 96: 23813.03ms, total: 23830.17ms  
Thread 1: query mix 97: 24646.64ms, total: 24662.28ms  
Thread 3: query mix 98: 27444.94ms, total: 27461.76ms

Thread 4: query mix 100: 22783.27ms, total: 22845.59ms  
 Thread 2: query mix 99: 33280.03ms, total: 33297.42ms  
 Thread 1: query mix 101: 23494.65ms, total: 23511.48ms  
 Thread 4: query mix 103: 21795.72ms, total: 21812.25ms  
 Thread 2: query mix 104: 22960.23ms, total: 22977.58ms  
 Thread 3: query mix 102: 35818.42ms, total: 35839.38ms  
 Thread 1: query mix 105: 26667.58ms, total: 26688.18ms  
 Thread 3: query mix 108: 19404.16ms, total: 19424.26ms  
 Thread 4: query mix 106: 33951.26ms, total: 33968.88ms  
 Thread 2: query mix 107: 32965.38ms, total: 32987.96ms  
 Thread 1: query mix 109: 26114.09ms, total: 26135.31ms  
 Thread 3: query mix 110: 26568.56ms, total: 26585.96ms  
 Thread 2: query mix 112: 23193.07ms, total: 23208.97ms  
 Thread 4: query mix 111: 32311.45ms, total: 32328.36ms  
 Thread 1: query mix 113: 28969.40ms, total: 28986.57ms  
 Thread 3: query mix 114: 24580.40ms, total: 24597.06ms  
 Thread 2: query mix 115: 26261.06ms, total: 26277.15ms  
 Thread 4: query mix 116: 24831.61ms, total: 24847.45ms  
 Thread 1: query mix 117: 25724.27ms, total: 25742.11ms  
 Thread 3: query mix 118: 28615.84ms, total: 28631.41ms  
 Thread 2: query mix 119: 25655.61ms, total: 25674.05ms  
 Thread 4: query mix 120: 24803.93ms, total: 24820.43ms  
 Thread 1: query mix 121: 24991.98ms, total: 25012.30ms  
 Thread 3: query mix 122: 27270.56ms, total: 27285.76ms  
 Thread 2: query mix 123: 25505.01ms, total: 25521.69ms  
 Thread 4: query mix 124: 27858.46ms, total: 27874.64ms  
 Thread 1: query mix 125: 27274.80ms, total: 27292.78ms  
 Thread 3: query mix 126: 19200.26ms, total: 19213.88ms  
 Thread 2: query mix 127: 19288.55ms, total: 19302.34ms  
 Benchmark run completed in 970.791758s

Scale factor: 284826  
 Number of warmup runs: 32  
 Number of clients: 4  
 Seed: 808080  
 Number of query mix runs (without warmups): 128 times  
 min/max Querymix runtime: 19.2003s / 56.1505s  
 Total runtime (sum): 3854.427 seconds  
 Total actual runtime: 970.792 seconds  
 QMPH: 474.66 query mixes per hour  
 CQET: 30.11271 seconds average runtime of query mix  
 CQET (geom.): 29.55032 seconds geometric mean runtime of query mix

Metrics for Query: 1  
 Count: 128 times executed in whole run  
 AQET: 0.092901 seconds (arithmetic mean)  
 AQET(geom.): 0.076240 seconds (geometric mean)  
 QPS: 42.74 Queries per second  
 minQET/maxQET: 0.03997800s / 0.46494700s  
 Average result count: 8.00  
 min/max result count: 1 / 10  
 Number of timeouts: 0

Metrics for Query: 2  
 Count: 768 times executed in whole run  
 AQET: 0.294216 seconds (arithmetic mean)  
 AQET(geom.): 0.270738 seconds (geometric mean)  
 QPS: 13.49 Queries per second  
 minQET/maxQET: 0.16462800s / 1.61929600s  
 Average result count: 19.23  
 min/max result count: 7 / 32  
 Number of timeouts: 0

Metrics for Query: 3  
 Count: 128 times executed in whole run  
 AQET: 0.115764 seconds (arithmetic mean)  
 AQET(geom.): 0.094400 seconds (geometric mean)  
 QPS: 34.30 Queries per second  
 minQET/maxQET: 0.04467000s / 0.83524400s  
 Average result count: 6.88  
 min/max result count: 0 / 10  
 Number of timeouts: 0

Metrics for Query: 4  
 Count: 128 times executed in whole run  
 AQET: 0.178933 seconds (arithmetic mean)  
 AQET(geom.): 0.153360 seconds (geometric mean)  
 QPS: 22.19 Queries per second  
 minQET/maxQET: 0.08579100s / 1.39441900s  
 Average result count: 7.16  
 min/max result count: 0 / 10  
 Number of timeouts: 0

Metrics for Query: 5  
 Count: 128 times executed in whole run  
 AQET: 7.091967 seconds (arithmetic mean)  
 AQET(geom.): 5.366304 seconds (geometric mean)  
 QPS: 0.56 Queries per second  
 minQET/maxQET: 0.50699100s / 21.72135900s  
 Average result count: 5.00  
 min/max result count: 5 / 5  
 Number of timeouts: 0

Metrics for Query: 6  
 Count: 128 times executed in whole run  
 AQET: 15.966659 seconds (arithmetic mean)  
 AQET(geom.): 15.618299 seconds (geometric mean)  
 QPS: 0.25 Queries per second  
 minQET/maxQET: 6.74919600s / 34.08094600s  
 Average result count: 247.69  
 min/max result count: 1 / 30401  
 Number of timeouts: 0

Metrics for Query: 7  
 Count: 512 times executed in whole run  
 AQET: 0.662243 seconds (arithmetic mean)  
 AQET(geom.): 0.561403 seconds (geometric mean)  
 QPS: 6.00 Queries per second  
 minQET/maxQET: 0.07060800s / 2.37355700s  
 Average result count: 14.93  
 min/max result count: 1 / 132  
 Number of timeouts: 0

Metrics for Query: 8  
 Count: 256 times executed in whole run  
 AQET: 0.431152 seconds (arithmetic mean)  
 AQET(geom.): 0.359520 seconds (geometric mean)  
 QPS: 9.21 Queries per second  
 minQET/maxQET: 0.01889500s / 2.18005100s  
 Average result count: 4.85  
 min/max result count: 0 / 15  
 Number of timeouts: 0

Metrics for Query: 9  
 Count: 512 times executed in whole run  
 AQET: 0.127354 seconds (arithmetic mean)

```

AQET(geom.):          0.116165 seconds (geometric mean)
QPS:                  31.18 Queries per second
minQET/maxQET:       0.02915800s / 1.19600200s
Average result (Bytes): 6771.43
min/max result (Bytes): 2001 / 12744
Number of timeouts:   0

```

```

Metrics for Query:    10
Count:                256 times executed in whole run
AQET:                 0.286919 seconds (arithmetic mean)
AQET(geom.):          0.243162 seconds (geometric mean)
QPS:                  13.84 Queries per second
minQET/maxQET:       0.07806000s / 1.57588500s
Average result count: 2.00
min/max result count: 0 / 10
Number of timeouts:   0

```

```

Metrics for Query:    11
Count:                128 times executed in whole run
AQET:                 0.113234 seconds (arithmetic mean)
AQET(geom.):          0.103512 seconds (geometric mean)
QPS:                  35.06 Queries per second
minQET/maxQET:       0.05786300s / 0.98969600s
Average result count: 10.00
min/max result count: 10 / 10
Number of timeouts:   0

```

```

Metrics for Query:    12
Count:                128 times executed in whole run
AQET:                 0.193427 seconds (arithmetic mean)
AQET(geom.):          0.177073 seconds (geometric mean)
QPS:                  20.53 Queries per second
minQET/maxQET:       0.11793300s / 1.29457200s
Average result (Bytes): 1425.98
min/max result (Bytes): 1397 / 1460
Number of timeouts:   0

```

Reading Test Driver data...done

Starting test...

Generating queries...done

Clients created.

```

Thread 3: query mix -30: 0.00ms, total: 27294.43ms
Thread 7: query mix -26: 0.00ms, total: 30270.61ms
Thread 4: query mix -29: 0.00ms, total: 31061.44ms
Thread 1: query mix -32: 0.00ms, total: 34657.91ms
Thread 5: query mix -27: 0.00ms, total: 35240.41ms
Thread 8: query mix -25: 0.00ms, total: 38161.48ms
Thread 6: query mix -28: 0.00ms, total: 39495.30ms
Thread 2: query mix -31: 0.00ms, total: 44806.37ms
Thread 1: query mix -21: 0.00ms, total: 28509.73ms
Thread 7: query mix -23: 0.00ms, total: 32948.75ms
Thread 6: query mix -18: 0.00ms, total: 25484.76ms
Thread 8: query mix -19: 0.00ms, total: 29798.95ms
Thread 4: query mix -22: 0.00ms, total: 37627.57ms
Thread 5: query mix -20: 0.00ms, total: 36301.97ms
Thread 3: query mix -24: 0.00ms, total: 44284.88ms
Thread 2: query mix -17: 0.00ms, total: 45819.58ms
Thread 4: query mix -12: 0.00ms, total: 24236.35ms
Thread 5: query mix -11: 0.00ms, total: 27288.38ms
Thread 8: query mix -13: 0.00ms, total: 33139.49ms
Thread 6: query mix -14: 0.00ms, total: 36519.29ms

```



Thread 7: query mix -15: 0.00ms, total: 40945.10ms  
Thread 3: query mix -10: 0.00ms, total: 33046.03ms  
Thread 1: query mix -16: 0.00ms, total: 48949.65ms  
Thread 5: query mix -7: 0.00ms, total: 23867.48ms  
Thread 2: query mix -9: 0.00ms, total: 32919.67ms  
Thread 8: query mix -6: 0.00ms, total: 27550.25ms  
Thread 3: query mix -3: 0.00ms, total: 24429.34ms  
Thread 7: query mix -4: 0.00ms, total: 31157.41ms  
Thread 1: query mix -2: 0.00ms, total: 23236.60ms  
Thread 4: query mix -8: 0.00ms, total: 44732.73ms  
Thread 6: query mix -5: 0.00ms, total: 36897.59ms  
Thread 5: query mix -1: 0.00ms, total: 18247.87ms  
Warmup phase ended...

Starting actual run...

Thread 5: query mix 7: 27641.69ms, total: 27681.98ms  
Thread 1: query mix 2: 31663.86ms, total: 31708.68ms  
Thread 4: query mix 5: 33451.30ms, total: 33483.60ms  
Thread 3: query mix 3: 34936.19ms, total: 34983.89ms  
Thread 2: query mix 4: 35591.58ms, total: 35644.52ms  
Thread 8: query mix 1: 37199.54ms, total: 37248.62ms  
Thread 6: query mix 6: 41926.45ms, total: 41969.96ms  
Thread 7: query mix 0: 45005.20ms, total: 45052.11ms  
Thread 1: query mix 9: 27447.30ms, total: 27473.43ms  
Thread 5: query mix 8: 33636.07ms, total: 33658.49ms  
Thread 8: query mix 13: 42110.85ms, total: 42136.07ms  
Thread 3: query mix 11: 48600.25ms, total: 48626.37ms  
Thread 4: query mix 10: 50163.16ms, total: 50245.55ms  
Thread 7: query mix 15: 38791.46ms, total: 38813.21ms  
Thread 6: query mix 14: 46619.86ms, total: 46640.97ms  
Thread 2: query mix 12: 54060.24ms, total: 54086.46ms  
Thread 1: query mix 16: 35576.49ms, total: 35600.65ms  
Thread 5: query mix 17: 33421.53ms, total: 33443.08ms  
Thread 8: query mix 18: 23641.00ms, total: 23659.88ms  
Thread 4: query mix 20: 26522.76ms, total: 26542.92ms  
Thread 1: query mix 24: 22922.02ms, total: 22939.97ms  
Thread 6: query mix 22: 34145.95ms, total: 34164.49ms  
Thread 7: query mix 21: 39227.43ms, total: 39247.16ms  
Thread 3: query mix 19: 42070.61ms, total: 42090.71ms  
Thread 5: query mix 25: 32620.30ms, total: 32639.59ms  
Thread 2: query mix 23: 43619.13ms, total: 43638.19ms  
Thread 8: query mix 26: 37217.80ms, total: 37236.24ms  
Thread 1: query mix 28: 24840.48ms, total: 24859.73ms  
Thread 4: query mix 27: 39771.87ms, total: 39789.61ms  
Thread 7: query mix 30: 33542.97ms, total: 33564.73ms  
Thread 6: query mix 29: 34795.97ms, total: 34816.12ms  
Thread 2: query mix 33: 28915.33ms, total: 28935.70ms  
Thread 5: query mix 32: 38325.90ms, total: 38344.30ms  
Thread 3: query mix 31: 41794.57ms, total: 41815.35ms  
Thread 1: query mix 35: 30732.97ms, total: 30751.67ms  
Thread 8: query mix 34: 35005.74ms, total: 35027.89ms  
Thread 4: query mix 36: 36073.28ms, total: 36091.98ms  
Thread 7: query mix 37: 34154.69ms, total: 34175.15ms  
Thread 2: query mix 39: 30351.08ms, total: 30370.92ms  
Thread 6: query mix 38: 38710.15ms, total: 38726.44ms  
Thread 3: query mix 41: 31263.95ms, total: 31283.28ms  
Thread 8: query mix 43: 28361.40ms, total: 28381.05ms  
Thread 1: query mix 42: 31081.80ms, total: 31100.78ms  
Thread 5: query mix 40: 43622.14ms, total: 43641.10ms  
Thread 4: query mix 44: 31771.96ms, total: 31790.18ms  
Thread 3: query mix 48: 26353.74ms, total: 26374.33ms  
Thread 2: query mix 46: 36347.74ms, total: 36366.10ms  
Thread 8: query mix 49: 26092.84ms, total: 26115.15ms

Thread 6: query mix 47: 35140.87ms, total: 35164.07ms  
Thread 7: query mix 45: 44217.34ms, total: 44236.73ms  
Thread 1: query mix 50: 31751.94ms, total: 31778.34ms  
Thread 4: query mix 52: 30813.39ms, total: 30831.28ms  
Thread 5: query mix 51: 40451.85ms, total: 40469.38ms  
Thread 3: query mix 53: 27209.20ms, total: 27228.06ms  
Thread 8: query mix 55: 27280.16ms, total: 27302.56ms  
Thread 7: query mix 57: 24856.54ms, total: 24879.20ms  
Thread 2: query mix 54: 31517.29ms, total: 31537.78ms  
Thread 4: query mix 59: 23377.41ms, total: 23396.04ms  
Thread 6: query mix 56: 47863.57ms, total: 47888.40ms  
Thread 5: query mix 60: 29699.97ms, total: 29723.29ms  
Thread 1: query mix 58: 43468.86ms, total: 43489.86ms  
Thread 2: query mix 64: 29140.12ms, total: 29158.23ms  
Thread 7: query mix 63: 31320.80ms, total: 31339.45ms  
Thread 4: query mix 65: 26066.38ms, total: 26085.77ms  
Thread 8: query mix 62: 46733.76ms, total: 46750.43ms  
Thread 3: query mix 61: 54139.15ms, total: 54156.56ms  
Thread 6: query mix 66: 29297.80ms, total: 29322.23ms  
Thread 5: query mix 67: 34575.61ms, total: 34596.41ms  
Thread 2: query mix 69: 25583.67ms, total: 25603.41ms  
Thread 7: query mix 70: 30797.49ms, total: 30813.54ms  
Thread 4: query mix 71: 28196.58ms, total: 28216.51ms  
Thread 1: query mix 68: 47108.52ms, total: 47131.71ms  
Thread 8: query mix 72: 32948.86ms, total: 32966.53ms  
Thread 3: query mix 73: 33009.20ms, total: 33025.71ms  
Thread 6: query mix 74: 36180.40ms, total: 36197.92ms  
Thread 7: query mix 77: 25601.87ms, total: 25619.38ms  
Thread 5: query mix 75: 37364.08ms, total: 37382.93ms  
Thread 2: query mix 76: 43128.28ms, total: 43145.02ms  
Thread 1: query mix 79: 34075.03ms, total: 34093.62ms  
Thread 8: query mix 80: 28952.57ms, total: 28970.23ms  
Thread 7: query mix 83: 22686.98ms, total: 22703.84ms  
Thread 6: query mix 82: 29824.65ms, total: 29843.17ms  
Thread 3: query mix 81: 37292.52ms, total: 37310.18ms  
Thread 4: query mix 78: 53722.42ms, total: 53741.33ms  
Thread 1: query mix 86: 27555.95ms, total: 27572.50ms  
Thread 2: query mix 85: 31885.01ms, total: 31902.83ms  
Thread 5: query mix 84: 47017.03ms, total: 47033.46ms  
Thread 6: query mix 89: 28309.41ms, total: 28331.49ms  
Thread 7: query mix 88: 33197.14ms, total: 33216.50ms  
Thread 4: query mix 91: 25746.61ms, total: 25766.07ms  
Thread 1: query mix 92: 25943.92ms, total: 25962.24ms  
Thread 8: query mix 87: 50666.51ms, total: 50725.65ms  
Thread 5: query mix 94: 24439.74ms, total: 24461.32ms  
Thread 3: query mix 90: 50636.42ms, total: 50654.45ms  
Thread 7: query mix 96: 26572.33ms, total: 26593.59ms  
Thread 6: query mix 95: 27181.60ms, total: 27202.39ms  
Thread 4: query mix 97: 30378.93ms, total: 30396.83ms  
Thread 2: query mix 93: 48273.71ms, total: 48291.72ms  
Thread 1: query mix 98: 34975.89ms, total: 34994.09ms  
Thread 5: query mix 100: 29787.98ms, total: 29805.46ms  
Thread 3: query mix 101: 26237.89ms, total: 26258.98ms  
Thread 6: query mix 103: 29576.26ms, total: 29597.68ms  
Thread 8: query mix 99: 45229.39ms, total: 45246.66ms  
Thread 4: query mix 104: 27021.28ms, total: 27039.79ms  
Thread 2: query mix 105: 31400.87ms, total: 31418.88ms  
Thread 7: query mix 102: 43475.66ms, total: 43498.87ms  
Thread 3: query mix 108: 23360.42ms, total: 23381.93ms  
Thread 8: query mix 110: 24819.18ms, total: 24836.30ms  
Thread 6: query mix 109: 27845.85ms, total: 27866.03ms  
Thread 1: query mix 106: 45935.31ms, total: 45955.02ms  
Thread 2: query mix 112: 26116.51ms, total: 26136.54ms

Thread 4: query mix 111: 37275.03ms, total: 37295.86ms  
 Thread 5: query mix 107: 48082.01ms, total: 48102.17ms  
 Thread 3: query mix 114: 30634.18ms, total: 30653.40ms  
 Thread 7: query mix 113: 38798.54ms, total: 38817.05ms  
 Thread 6: query mix 116: 30915.34ms, total: 30936.90ms  
 Thread 1: query mix 117: 30028.76ms, total: 30046.86ms  
 Thread 8: query mix 115: 38957.50ms, total: 38977.22ms  
 Thread 5: query mix 120: 28869.96ms, total: 28886.66ms  
 Thread 4: query mix 119: 30680.75ms, total: 30739.05ms  
 Thread 2: query mix 118: 38992.00ms, total: 39009.71ms  
 Thread 3: query mix 121: 31387.21ms, total: 31404.33ms  
 Thread 6: query mix 123: 31304.10ms, total: 31321.23ms  
 Thread 7: query mix 122: 38812.90ms, total: 38830.51ms  
 Thread 8: query mix 125: 26518.83ms, total: 26544.64ms  
 Thread 1: query mix 124: 26682.27ms, total: 26702.21ms  
 Thread 5: query mix 126: 24179.95ms, total: 24195.56ms  
 Thread 4: query mix 127: 22833.39ms, total: 22849.47ms  
 Benchmark run completed in 554.324475s

Scale factor: 284826  
 Number of warmup runs: 32  
 Number of clients: 8  
 Seed: 808080  
 Number of query mix runs (without warmups): 128 times  
 min/max Querymix runtime: 22.6870s / 54.1392s  
 Total runtime (sum): 4385.675 seconds  
 Total actual runtime: 554.324 seconds  
 QMPH: 831.28 query mixes per hour  
 CQET: 34.26309 seconds average runtime of query mix  
 CQET (geom.): 33.42904 seconds geometric mean runtime of query mix

Metrics for Query: 1  
 Count: 128 times executed in whole run  
 AQET: 0.116431 seconds (arithmetic mean)  
 AQET(geom.): 0.074423 seconds (geometric mean)  
 QPS: 67.95 Queries per second  
 minQET/maxQET: 0.04324200s / 2.09911500s  
 Average result count: 8.00  
 min/max result count: 1 / 10  
 Number of timeouts: 0

Metrics for Query: 2  
 Count: 768 times executed in whole run  
 AQET: 0.346561 seconds (arithmetic mean)  
 AQET(geom.): 0.303735 seconds (geometric mean)  
 QPS: 22.83 Queries per second  
 minQET/maxQET: 0.17544900s / 2.27497100s  
 Average result count: 19.23  
 min/max result count: 7 / 32  
 Number of timeouts: 0

Metrics for Query: 3  
 Count: 128 times executed in whole run  
 AQET: 0.085778 seconds (arithmetic mean)  
 AQET(geom.): 0.076927 seconds (geometric mean)  
 QPS: 92.24 Queries per second  
 minQET/maxQET: 0.04194000s / 0.28456300s  
 Average result count: 6.88  
 min/max result count: 0 / 10  
 Number of timeouts: 0

Metrics for Query: 4  
 Count: 128 times executed in whole run

```

AQET: 0.150389 seconds (arithmetic mean)
AQET(geom.): 0.138792 seconds (geometric mean)
QPS: 52.61 Queries per second
minQET/maxQET: 0.08812600s / 0.43902600s
Average result count: 7.16
min/max result count: 0 / 10
Number of timeouts: 0

Metrics for Query: 5
Count: 128 times executed in whole run
AQET: 10.321151 seconds (arithmetic mean)
AQET(geom.): 7.866798 seconds (geometric mean)
QPS: 0.77 Queries per second
minQET/maxQET: 0.94145900s / 32.24380500s
Average result count: 5.00
min/max result count: 5 / 5
Number of timeouts: 0

Metrics for Query: 6
Count: 128 times executed in whole run
AQET: 19.855620 seconds (arithmetic mean)
AQET(geom.): 19.683890 seconds (geometric mean)
QPS: 0.40 Queries per second
minQET/maxQET: 11.73113100s / 35.25364700s
Average result count: 247.69
min/max result count: 1 / 30401
Number of timeouts: 0

Metrics for Query: 7
Count: 512 times executed in whole run
AQET: 0.174320 seconds (arithmetic mean)
AQET(geom.): 0.158507 seconds (geometric mean)
QPS: 45.39 Queries per second
minQET/maxQET: 0.07234500s / 2.15697500s
Average result count: 14.93
min/max result count: 1 / 132
Number of timeouts: 0

Metrics for Query: 8
Count: 256 times executed in whole run
AQET: 0.099776 seconds (arithmetic mean)
AQET(geom.): 0.091290 seconds (geometric mean)
QPS: 79.30 Queries per second
minQET/maxQET: 0.01550100s / 0.29148300s
Average result count: 4.85
min/max result count: 0 / 15
Number of timeouts: 0

Metrics for Query: 9
Count: 512 times executed in whole run
AQET: 0.050274 seconds (arithmetic mean)
AQET(geom.): 0.041492 seconds (geometric mean)
QPS: 157.37 Queries per second
minQET/maxQET: 0.02355200s / 2.01896100s
Average result (Bytes): 6771.43
min/max result (Bytes): 2001 / 12744
Number of timeouts: 0

Metrics for Query: 10
Count: 256 times executed in whole run
AQET: 0.149327 seconds (arithmetic mean)
AQET(geom.): 0.110015 seconds (geometric mean)
QPS: 52.98 Queries per second

```

```

minQET/maxQET:          0.06234100s / 2.21065600s
Average result count:   2.00
min/max result count:   0 / 10
Number of timeouts:     0

Metrics for Query:      11
Count:                  128 times executed in whole run
AQET:                  0.031363 seconds (arithmetic mean)
AQET(geom.):           0.024837 seconds (geometric mean)
QPS:                   252.26 Queries per second
minQET/maxQET:         0.01491600s / 0.23392100s
Average result count:   10.00
min/max result count:   10 / 10
Number of timeouts:     0

Metrics for Query:      12
Count:                  128 times executed in whole run
AQET:                  0.226412 seconds (arithmetic mean)
AQET(geom.):           0.174637 seconds (geometric mean)
QPS:                   34.94 Queries per second
minQET/maxQET:         0.09737000s / 2.27253200s
Average result (Bytes): 1425.98
min/max result (Bytes): 1397 / 1460
Number of timeouts:     0

```

Reading Test Driver data...done

Starting test...

Generating queries...done

Clients created.

```

Thread 16: query mix -18: 0.00ms, total: 52402.04ms
Thread 7: query mix -26: 0.00ms, total: 52470.45ms
Thread 14: query mix -19: 0.00ms, total: 52475.17ms
Thread 13: query mix -21: 0.00ms, total: 57002.73ms
Thread 3: query mix -30: 0.00ms, total: 59799.19ms
Thread 4: query mix -29: 0.00ms, total: 62269.57ms
Thread 10: query mix -23: 0.00ms, total: 67950.73ms
Thread 8: query mix -25: 0.00ms, total: 70974.95ms
Thread 6: query mix -27: 0.00ms, total: 71273.83ms
Thread 1: query mix -32: 0.00ms, total: 71701.57ms
Thread 11: query mix -20: 0.00ms, total: 74127.08ms
Thread 5: query mix -28: 0.00ms, total: 75853.36ms
Thread 12: query mix -22: 0.00ms, total: 80755.52ms
Thread 2: query mix -31: 0.00ms, total: 88718.02ms
Thread 9: query mix -24: 0.00ms, total: 89193.38ms
Thread 15: query mix -17: 0.00ms, total: 95400.28ms
Thread 3: query mix -12: 0.00ms, total: 51857.03ms
Thread 4: query mix -11: 0.00ms, total: 52105.82ms
Thread 1: query mix -7: 0.00ms, total: 45172.32ms
Thread 10: query mix -10: 0.00ms, total: 52075.45ms
Thread 13: query mix -13: 0.00ms, total: 64488.12ms
Thread 11: query mix -6: 0.00ms, total: 53234.26ms
Thread 14: query mix -14: 0.00ms, total: 74934.55ms
Thread 8: query mix -9: 0.00ms, total: 56821.52ms
Thread 7: query mix -15: 0.00ms, total: 76414.00ms
Thread 2: query mix -3: 0.00ms, total: 42408.29ms
Thread 16: query mix -16: 0.00ms, total: 81023.73ms
Thread 12: query mix -4: 0.00ms, total: 53210.91ms
Thread 9: query mix -2: 0.00ms, total: 46054.96ms
Thread 5: query mix -5: 0.00ms, total: 61297.89ms
Thread 6: query mix -8: 0.00ms, total: 67852.87ms
Thread 15: query mix -1: 0.00ms, total: 44074.48ms

```

Warmup phase ended...

Starting actual run...

Thread 14: query mix 7: 53805.96ms, total: 54649.14ms  
Thread 6: query mix 9: 59625.50ms, total: 59856.96ms  
Thread 11: query mix 4: 62111.12ms, total: 62570.48ms  
Thread 9: query mix 2: 63822.12ms, total: 64746.40ms  
Thread 5: query mix 15: 64279.20ms, total: 65130.25ms  
Thread 1: query mix 8: 65450.24ms, total: 66067.86ms  
Thread 8: query mix 5: 71187.44ms, total: 71847.95ms  
Thread 7: query mix 14: 76913.25ms, total: 77329.46ms  
Thread 15: query mix 11: 76498.02ms, total: 77737.61ms  
Thread 12: query mix 1: 77018.48ms, total: 78206.39ms  
Thread 4: query mix 6: 81506.62ms, total: 82644.50ms  
Thread 13: query mix 10: 86356.31ms, total: 87056.85ms  
Thread 3: query mix 3: 91301.87ms, total: 91846.00ms  
Thread 2: query mix 0: 93686.45ms, total: 94167.30ms  
Thread 16: query mix 12: 95456.15ms, total: 95780.66ms  
Thread 14: query mix 16: 57153.21ms, total: 57174.02ms  
Thread 11: query mix 18: 51388.17ms, total: 51410.60ms  
Thread 6: query mix 17: 56280.26ms, total: 56300.35ms  
Thread 5: query mix 20: 61020.54ms, total: 61042.74ms  
Thread 10: query mix 13: 126215.18ms, total: 127044.75ms  
Thread 15: query mix 24: 54171.17ms, total: 54195.09ms  
Thread 3: query mix 28: 47432.72ms, total: 47452.61ms  
Thread 8: query mix 22: 68831.39ms, total: 68852.49ms  
Thread 12: query mix 25: 64575.48ms, total: 64600.56ms  
Thread 4: query mix 26: 63616.73ms, total: 63639.46ms  
Thread 9: query mix 19: 81689.03ms, total: 81712.07ms  
Thread 1: query mix 21: 82796.40ms, total: 82827.99ms  
Thread 13: query mix 27: 71062.07ms, total: 71083.29ms  
Thread 16: query mix 30: 64701.57ms, total: 64723.87ms  
Thread 7: query mix 23: 87244.04ms, total: 87268.75ms  
Thread 2: query mix 29: 70559.50ms, total: 70585.51ms  
Thread 6: query mix 33: 57923.08ms, total: 57941.22ms  
Thread 5: query mix 34: 61454.33ms, total: 61475.78ms  
Thread 10: query mix 35: 64456.85ms, total: 64478.93ms  
Thread 14: query mix 31: 80284.94ms, total: 80314.09ms  
Thread 11: query mix 32: 79275.35ms, total: 79298.99ms  
Thread 12: query mix 39: 56036.56ms, total: 56058.50ms  
Thread 1: query mix 42: 54164.44ms, total: 54184.03ms  
Thread 15: query mix 36: 72674.43ms, total: 72695.99ms  
Thread 3: query mix 37: 69066.45ms, total: 69088.81ms  
Thread 9: query mix 41: 62038.34ms, total: 62066.59ms  
Thread 13: query mix 43: 51015.33ms, total: 51039.43ms  
Thread 16: query mix 44: 56237.12ms, total: 56255.99ms  
Thread 8: query mix 38: 76449.05ms, total: 76470.93ms  
Thread 4: query mix 40: 80062.05ms, total: 80082.45ms  
Thread 2: query mix 46: 73695.79ms, total: 73716.19ms  
Thread 7: query mix 45: 77635.73ms, total: 77657.46ms  
Thread 6: query mix 47: 73223.05ms, total: 73243.38ms  
Thread 10: query mix 49: 55877.43ms, total: 55902.80ms  
Thread 5: query mix 48: 61592.29ms, total: 61618.64ms  
Thread 14: query mix 50: 57268.33ms, total: 57291.48ms  
Thread 1: query mix 53: 50345.97ms, total: 50369.09ms  
Thread 12: query mix 52: 58156.79ms, total: 58176.10ms  
Thread 11: query mix 51: 69377.04ms, total: 69400.15ms  
Thread 13: query mix 57: 53911.77ms, total: 53934.66ms  
Thread 3: query mix 55: 59468.71ms, total: 59496.73ms  
Thread 15: query mix 54: 65874.38ms, total: 65895.30ms  
Thread 8: query mix 59: 57445.11ms, total: 57466.31ms  
Thread 4: query mix 60: 58800.39ms, total: 58817.33ms  
Thread 5: query mix 65: 45725.77ms, total: 45745.49ms

Thread 9: query mix 56: 87301.41ms, total: 87328.54ms  
Thread 10: query mix 64: 55387.79ms, total: 55414.43ms  
Thread 14: query mix 66: 54963.82ms, total: 54991.36ms  
Thread 6: query mix 63: 57608.47ms, total: 57634.55ms  
Thread 16: query mix 58: 89061.83ms, total: 89086.76ms  
Thread 11: query mix 69: 46869.86ms, total: 46891.49ms  
Thread 13: query mix 70: 52296.06ms, total: 52318.50ms  
Thread 3: query mix 71: 52154.63ms, total: 52178.31ms  
Thread 1: query mix 67: 67065.65ms, total: 67087.95ms  
Thread 7: query mix 62: 90627.50ms, total: 90647.10ms  
Thread 15: query mix 72: 63132.43ms, total: 63152.04ms  
Thread 8: query mix 73: 65205.19ms, total: 65226.15ms  
Thread 2: query mix 61: 106673.27ms, total: 106697.00ms  
Thread 12: query mix 68: 89588.43ms, total: 89608.00ms  
Thread 10: query mix 77: 46674.42ms, total: 46695.65ms  
Thread 4: query mix 74: 65670.69ms, total: 65741.64ms  
Thread 16: query mix 80: 54052.14ms, total: 54073.66ms  
Thread 3: query mix 83: 46383.98ms, total: 46403.61ms  
Thread 5: query mix 75: 72477.91ms, total: 72497.05ms  
Thread 6: query mix 79: 63009.54ms, total: 63029.28ms  
Thread 13: query mix 82: 59864.54ms, total: 59881.85ms  
Thread 9: query mix 76: 82748.47ms, total: 82769.89ms  
Thread 11: query mix 81: 73896.69ms, total: 73915.91ms  
Thread 15: query mix 86: 53092.31ms, total: 53110.01ms  
Thread 7: query mix 85: 57201.88ms, total: 57219.88ms  
Thread 4: query mix 91: 46231.42ms, total: 46249.54ms  
Thread 14: query mix 78: 95542.24ms, total: 95567.75ms  
Thread 12: query mix 89: 54812.01ms, total: 54832.40ms  
Thread 16: query mix 92: 48386.25ms, total: 48404.87ms  
Thread 2: query mix 88: 65405.18ms, total: 65425.80ms  
Thread 1: query mix 84: 92776.94ms, total: 92795.02ms  
Thread 5: query mix 94: 50201.43ms, total: 50222.46ms  
Thread 6: query mix 95: 51454.70ms, total: 51475.19ms  
Thread 13: query mix 96: 48028.77ms, total: 48046.40ms  
Thread 9: query mix 97: 52859.39ms, total: 52881.60ms  
Thread 8: query mix 87: 95158.09ms, total: 95177.00ms  
Thread 4: query mix 101: 44626.74ms, total: 44647.69ms  
Thread 7: query mix 100: 57920.33ms, total: 57941.19ms  
Thread 11: query mix 98: 66488.91ms, total: 66506.73ms  
Thread 12: query mix 103: 51886.80ms, total: 51912.21ms  
Thread 10: query mix 90: 104321.13ms, total: 104340.34ms  
Thread 16: query mix 104: 48428.79ms, total: 48454.40ms  
Thread 3: query mix 93: 93551.58ms, total: 93571.41ms  
Thread 6: query mix 108: 46475.53ms, total: 46494.13ms  
Thread 2: query mix 105: 58534.31ms, total: 58555.84ms  
Thread 15: query mix 99: 88358.87ms, total: 88378.99ms  
Thread 14: query mix 102: 82227.71ms, total: 82247.18ms  
Thread 13: query mix 109: 62379.09ms, total: 62398.23ms  
Thread 9: query mix 110: 54345.91ms, total: 54371.00ms  
Thread 4: query mix 112: 51844.60ms, total: 51865.47ms  
Thread 11: query mix 114: 52122.70ms, total: 52142.02ms  
Thread 1: query mix 106: 93159.34ms, total: 93179.47ms  
Thread 16: query mix 117: 50474.23ms, total: 50492.69ms  
Thread 10: query mix 116: 54213.03ms, total: 54238.44ms  
Thread 8: query mix 111: 73160.15ms, total: 73180.96ms  
Thread 5: query mix 107: 92735.80ms, total: 92755.46ms  
Thread 7: query mix 113: 66501.15ms, total: 66520.38ms  
Thread 2: query mix 120: 50710.88ms, total: 50729.77ms  
Thread 12: query mix 115: 68844.88ms, total: 68864.68ms  
Thread 6: query mix 119: 56604.36ms, total: 56626.17ms  
Thread 3: query mix 118: 64847.46ms, total: 64864.60ms  
Thread 15: query mix 121: 53103.77ms, total: 53121.04ms  
Thread 9: query mix 124: 47748.17ms, total: 47768.66ms

Thread 4: query mix 125: 40743.67ms, total: 40806.40ms  
 Thread 13: query mix 123: 49441.26ms, total: 49461.05ms  
 Thread 11: query mix 126: 33791.54ms, total: 33810.88ms  
 Thread 14: query mix 122: 54308.74ms, total: 54326.21ms  
 Thread 1: query mix 127: 30834.84ms, total: 30853.67ms  
 Benchmark run completed in 537.439564s

Scale factor: 284826  
 Number of warmup runs: 32  
 Number of clients: 16  
 Seed: 808080  
 Number of query mix runs (without warmups): 128 times  
 min/max Querymix runtime: 30.8348s / 126.2152s  
 Total runtime (sum): 8365.895 seconds  
 Total actual runtime: 537.440 seconds  
 QMPH: 857.40 query mixes per hour  
 CQET: 65.35855 seconds average runtime of query mix  
 CQET (geom.): 63.47983 seconds geometric mean runtime of query mix

Metrics for Query: 1  
 Count: 128 times executed in whole run  
 AQET: 0.177424 seconds (arithmetic mean)  
 AQET(geom.): 0.124529 seconds (geometric mean)  
 QPS: 87.73 Queries per second  
 minQET/maxQET: 0.04672100s / 2.78305500s  
 Average result count: 8.00  
 min/max result count: 1 / 10  
 Number of timeouts: 0

Metrics for Query: 2  
 Count: 768 times executed in whole run  
 AQET: 0.581124 seconds (arithmetic mean)  
 AQET(geom.): 0.454545 seconds (geometric mean)  
 QPS: 26.79 Queries per second  
 minQET/maxQET: 0.18033500s / 3.97388500s  
 Average result count: 19.23  
 min/max result count: 7 / 32  
 Number of timeouts: 0

Metrics for Query: 3  
 Count: 128 times executed in whole run  
 AQET: 0.254536 seconds (arithmetic mean)  
 AQET(geom.): 0.151139 seconds (geometric mean)  
 QPS: 61.16 Queries per second  
 minQET/maxQET: 0.05524800s / 2.60983900s  
 Average result count: 6.88  
 min/max result count: 0 / 10  
 Number of timeouts: 0

Metrics for Query: 4  
 Count: 128 times executed in whole run  
 AQET: 0.314825 seconds (arithmetic mean)  
 AQET(geom.): 0.230439 seconds (geometric mean)  
 QPS: 49.44 Queries per second  
 minQET/maxQET: 0.09676700s / 3.00246500s  
 Average result count: 7.16  
 min/max result count: 0 / 10  
 Number of timeouts: 0

Metrics for Query: 5  
 Count: 128 times executed in whole run  
 AQET: 19.786027 seconds (arithmetic mean)  
 AQET(geom.): 14.993473 seconds (geometric mean)



```

QPS: 0.79 Queries per second
minQET/maxQET: 1.83400500s / 63.39033800s
Average result count: 5.00
min/max result count: 5 / 5
Number of timeouts: 0

Metrics for Query: 6
Count: 128 times executed in whole run
AQET: 38.863609 seconds (arithmetic mean)
AQET(geom.): 38.111776 seconds (geometric mean)
QPS: 0.40 Queries per second
minQET/maxQET: 14.11860000s / 108.55968000s
Average result count: 247.69
min/max result count: 1 / 30401
Number of timeouts: 0

Metrics for Query: 7
Count: 512 times executed in whole run
AQET: 0.301733 seconds (arithmetic mean)
AQET(geom.): 0.241491 seconds (geometric mean)
QPS: 51.59 Queries per second
minQET/maxQET: 0.08478800s / 2.95802700s
Average result count: 14.93
min/max result count: 1 / 132
Number of timeouts: 0

Metrics for Query: 8
Count: 256 times executed in whole run
AQET: 0.157071 seconds (arithmetic mean)
AQET(geom.): 0.134040 seconds (geometric mean)
QPS: 99.10 Queries per second
minQET/maxQET: 0.02487900s / 2.27214900s
Average result count: 4.85
min/max result count: 0 / 15
Number of timeouts: 0

Metrics for Query: 9
Count: 512 times executed in whole run
AQET: 0.092119 seconds (arithmetic mean)
AQET(geom.): 0.082252 seconds (geometric mean)
QPS: 168.98 Queries per second
minQET/maxQET: 0.02307900s / 0.39083500s
Average result (Bytes): 6771.43
min/max result (Bytes): 2001 / 12744
Number of timeouts: 0

Metrics for Query: 10
Count: 256 times executed in whole run
AQET: 0.136007 seconds (arithmetic mean)
AQET(geom.): 0.121771 seconds (geometric mean)
QPS: 114.45 Queries per second
minQET/maxQET: 0.06444200s / 0.39923800s
Average result count: 2.00
min/max result count: 0 / 10
Number of timeouts: 0

Metrics for Query: 11
Count: 128 times executed in whole run
AQET: 0.042349 seconds (arithmetic mean)
AQET(geom.): 0.036143 seconds (geometric mean)
QPS: 367.57 Queries per second
minQET/maxQET: 0.01499100s / 0.26726700s
Average result count: 10.00

```

```
min/max result count: 10 / 10
Number of timeouts: 0

Metrics for Query: 12
Count: 128 times executed in whole run
AQET: 0.271476 seconds (arithmetic mean)
AQET(geom.): 0.219163 seconds (geometric mean)
QPS: 57.34 Queries per second
minQET/maxQET: 0.11414600s / 3.08111000s
Average result (Bytes): 1425.98
min/max result (Bytes): 1397 / 1460
Number of timeouts: 0
```

**APPENDIX D**

BigTRREE 3.1.x BSBM query results summary

Scale factor	Number of warmup runs	Number of clients	Seed	Number of query mix runs (without warmups)	min Querymix runtime	max Querymix runtime	Total runtime (sum)	Total runtime (actual)	QMpH	CQ ET	CQET (geom.)
<b>70812</b>	32	1	808080	128	2.794	4.2249	419.814	419.814	1097.63	3.2798	3.26776
Query	Times Executed	AQET [sec]	AQET(geom.) [sec]	QPS	minQET [sec]	maxQET [sec]	Average result [count/bytes]	Min result [count/bytes]	Max result [count/bytes]	Timeouts	
<b>1</b>	128	0.049768	0.041712	20.09	0.033315	0.247843	5.27	0	10	0	
<b>2</b>	768	0.197361	0.18567	5.07	0.155792	1.306715	21.33	8	35	0	
<b>3</b>	128	0.047117	0.041969	21.22	0.035789	0.22575	4.37	0	10	0	
<b>4</b>	128	0.090362	0.08544	11.07	0.077569	0.783033	4.76	0	10	0	
<b>5</b>	128	0.322613	0.294702	3.1	0.085224	0.905606	5	5	5	0	
<b>6</b>	128	0.669317	0.666721	1.49	0.632436	1.257132	4.02	1	79	0	
<b>7</b>	512	0.093571	0.087163	10.69	0.027386	0.948893	15.67	1	114	0	
<b>8</b>	256	0.058586	0.055676	17.07	0.013445	0.206057	4.96	0	16	0	

<b>9</b>	512	0.025052	0.023333	39.92	0.017228	0.199358	6661.48	1747	11288	0	
<b>10</b>	256	0.081076	0.068617	12.33	0.057254	0.907642	1.74	0	8	0	
<b>11</b>	128	0.01685	0.016749	59.35	0.015563	0.025921	10	10	10	0	
<b>12</b>	128	0.145786	0.115589	6.86	0.090562	1.015629	1420.2	1397	1453	0	
<b>Scale factor</b>	Number of warmup runs	Number of clients	Seed	Number of query mix runs (without warmups)	min Querymix runtime	max Querymix runtime	Total runtime (sum)	Total runtime (actual)	QMpH	CQ ET	CQET (geom.)
<b>70812</b>	32	4	808080	128	8.2124	18.0242	1547.436	389.65	1182.6	12.08934	11.91046
<b>Query</b>	Times Executed	AQET [sec]	AQET(geom.) [sec]	QPS	minQET [sec]	maxQET [sec]	Average result [count/bytes]	Min result [count/bytes]	Max result count [count/bytes]	Timeouts	
<b>1</b>	128	0.090106	0.068435	44.07	0.036138	1.334832	5.27	0	10	0	
<b>2</b>	768	0.325642	0.294307	12.2	0.162211	1.620651	21.33	8	35	0	
<b>3</b>	128	0.094046	0.07646	42.23	0.040343	0.819323	4.37	0	10	0	
<b>4</b>	128	0.171897	0.152711	23.1	0.084683	1.131607	4.76	0	10	0	
<b>5</b>	128	2.824721	2.395342	1.41	0.237419	8.239437	5	5	5	0	

6	128	3.790936	3.613147	1.05	1.315839	7.676452	4.02	1	79	0	
7	512	0.377526	0.324281	10.52	0.031959	1.781982	15.67	1	114	0	
8	256	0.280276	0.23339	14.17	0.014922	1.606054	4.96	0	16	0	
9	512	0.097844	0.082302	40.59	0.019667	1.23403	6661.48	1747	11288	0	
10	256	0.189614	0.152862	20.94	0.066191	1.297125	1.74	0	8	0	
11	128	0.11691	0.091866	33.97	0.024977	1.475237	10	10	10	0	
12	128	0.205616	0.18013	19.31	0.107503	1.405704	1420.2	1397	1453	0	
<b>Scale factor</b>	Number of warmup runs	Number of clients	Seed	Number of query mix runs (without warmups)	min Querymix runtime	max Querymix runtime	Total runtime (sum)	Total runtime (actual)	QMpH	CQ ET	CQET (geom.)
<b>70812</b>	32	8	808080	128	8.9149	24.4157	2042.548	258.67	1781.42	15.9574	15.71952
<b>Query</b>	Times Executed	AQET [sec]	AQET(geom.) [sec]	QPS	minQET [sec]	maxQET [sec]	Average result [count/bytes]	Min result [count/bytes]	Max result count [count/bytes]	Timeouts	
<b>1</b>	128	0.124584	0.092178	63.38	0.041744	2.506674	5.27	0	10	0	
<b>2</b>	768	0.515222	0.441632	15.33	0.185575	2.982661	21.33	8	35	0	

<b>3</b>	128	0.1219 59	0.0928 84	64.75	0.043822	2.45301	4.37	0	10	0	
<b>4</b>	128	0.2241 49	0.1958 42	35.23	0.089102	1.862246	4.76	0	10	0	
<b>5</b>	128	4.5417 77	3.8403 21	1.74	0.494658	13.666202	5	5	5	0	
<b>6</b>	128	5.5023 52	5.3956 63	1.44	2.714365	8.54667	4.02	1	79	0	
<b>7</b>	512	0.2812 63	0.2282 34	28.07	0.050766	2.672537	15.67	1	114	0	
<b>8</b>	256	0.1535 5	0.1237 66	51.43	0.015369	2.284245	4.96	0	16	0	
<b>9</b>	512	0.0620 08	0.0503 58	127.34	0.021686	0.308398	6661.48	1747	11288	0	
<b>10</b>	256	0.1761 37	0.1509 35	44.83	0.065587	0.486326	1.74	0	8	0	
<b>11</b>	128	0.0339 38	0.0286 86	232.67	0.014315	0.296675	10	10	10	0	
<b>12</b>	128	0.2848 54	0.2414 15	27.72	0.097264	2.000639	1420.2	1397	1453	0	
<b>Scale factor</b>	Number of warmup runs	Number of clients	Seed	Number of query mix runs (without warmups)	min Querymix runtime	max Querymix runtime	Total runtime (sum)	Total runtime (actual)	QMpH	CQ ET	CQET (geom.)
<b>70812</b>	32	16	808080	128	15.453	56.5178	4209.545	271.285	1698.58	32.88707	32.1926
<b>Query</b>	Times Executed	AQET [sec]	AQET(geom.)	QPS	minQET [sec]	maxQET [sec]	Average result	Min result [count/by]	Max result count	Timeou	

			[sec]				[count/bytes]	tes]	[count/bytes]	ts	
<b>1</b>	128	0.204528	0.143414	75.87	0.044167	2.973137	5.27	0	10	0	
<b>2</b>	768	0.894627	0.694753	17.34	0.239051	5.630504	21.33	8	35	0	
<b>3</b>	128	0.195765	0.148002	79.26	0.055463	1.737861	4.37	0	10	0	
<b>4</b>	128	0.400346	0.327318	38.76	0.120457	1.813772	4.76	0	10	0	
<b>5</b>	128	10.075478	8.542204	1.54	1.361994	30.375319	5	5	5	0	
<b>6</b>	128	12.968448	12.618154	1.2	4.075927	18.425354	4.02	1	79	0	
<b>7</b>	512	0.384364	0.31451	40.37	0.062001	3.490542	15.67	1	114	0	
<b>8</b>	256	0.239336	0.184913	64.83	0.015036	3.159334	4.96	0	16	0	
<b>9</b>	512	0.15054	0.104369	103.08	0.022795	2.74055	6661.48	1747	11288	0	
<b>10</b>	256	0.281863	0.190443	55.05	0.071103	2.979177	1.74	0	8	0	
<b>11</b>	128	0.065522	0.04981	236.82	0.014592	0.412031	10	10	10	0	
<b>12</b>	128	0.427208	0.320713	36.32	0.120531	3.074707	1420.2	1397	1453	0	
<b>Scale factor</b>	Number of warmup runs	Number of clients	Seed	Number of query mix runs (without	min Querymix runtime	max Querymix runtime	Total runtime (sum)	Total runtime (actual)	QMpH	CQ ET	CQET (geom.)

				warmups)							
<b>284826</b>	32	1	808080	128	4.7606	15.073	715.257	715.257	644.24	5.58794	5.53512
<b>Query</b>	Times Executed	AQET [sec]	AQET(g eom.) [sec]	QPS	minQET [sec]	maxQET [sec]	Average result [count/bytes]	Min result [count/bytes]	Max result count [count/bytes]	Tim eou ts	
<b>1</b>	128	0.060056	0.045302	16.65	0.035289	0.968559	8	1	10	0	
<b>2</b>	768	0.189629	0.178799	5.27	0.156037	1.275593	19.23	7	32	0	
<b>3</b>	128	0.041619	0.041427	24.03	0.036152	0.08649	6.88	0	10	0	
<b>4</b>	128	0.118816	0.103081	8.42	0.07676	0.968562	7.16	0	10	0	
<b>5</b>	128	0.547721	0.456304	1.83	0.096155	1.518482	5	5	5	0	
<b>6</b>	128	2.715204	2.668257	0.37	2.484798	12.136268	247.69	1	30401	0	
<b>7</b>	512	0.098076	0.091449	10.2	0.058559	0.976188	14.93	1	132	0	
<b>8</b>	256	0.055225	0.054557	18.11	0.013605	0.084342	4.85	0	15	0	
<b>9</b>	512	0.024138	0.024043	41.43	0.018658	0.035771	6771.43	2001	12744	0	
<b>10</b>	256	0.072561	0.06746	13.78	0.057011	0.230703	2	0	10	0	
<b>11</b>	128	0.022681	0.018681	44.09	0.016572	0.188721	10	10	10	0	



<b>12</b>	128	0.1996 41	0.1638 75	5.01	0.09248	1.448136	1425.98	1397	1460	0	
<b>Scale factor</b>	Number of warmup runs	Number of clients	Seed	Number of query mix runs (without warmups)	min Querymix runtime	max Querymix runtime	Total runtime (sum)	Total runtime (actual)	QMpH	CQ ET	CQET (geom.)
<b>284826</b>	32	4	808080	128	19.2003	56.1505	3854.427	970.792	474.66	30.11271	29.55032
<b>Query</b>	Times Executed	AQET [sec]	AQET(geom.) [sec]	QPS	minQET [sec]	maxQET [sec]	Average result [count/bytes]	Min result [count/bytes]	Max result count [count/bytes]	Timeouts	
<b>1</b>	128	0.0929 01	0.0762 4	42.74	0.039978	0.464947	8	1	10	0	
<b>2</b>	768	0.2942 16	0.2707 38	13.49	0.164628	1.619296	19.23	7	32	0	
<b>3</b>	128	0.1157 64	0.0944	34.3	0.04467	0.835244	6.88	0	10	0	
<b>4</b>	128	0.1789 33	0.1533 6	22.19	0.085791	1.394419	7.16	0	10	0	
<b>5</b>	128	7.0919 67	5.3663 04	0.56	0.506991	21.721359	5	5	5	0	
<b>6</b>	128	15.966 659	15.618 299	0.25	6.749196	34.080946	247.69	1	30401	0	
<b>7</b>	512	0.6622 43	0.5614 03	6	0.070608	2.373557	14.93	1	132	0	
<b>8</b>	256	0.4311 52	0.3595 2	9.21	0.018895	2.180051	4.85	0	15	0	

<b>9</b>	512	0.127354	0.116165	31.18	0.029158	1.196002	6771.43	2001	12744	0	
<b>10</b>	256	0.286919	0.243162	13.84	0.07806	1.575885	2	0	10	0	
<b>11</b>	128	0.113234	0.103512	35.06	0.057863	0.989696	10	10	10	0	
<b>12</b>	128	0.193427	0.177073	20.53	0.117933	1.294572	1425.98	1397	1460	0	
<b>Scale factor</b>	Number of warmup runs	Number of clients	Seed	Number of query mix runs (without warmups)	min Querymix runtime	max Querymix runtime	Total runtime (sum)	Total runtime (actual)	QMpH	CQ ET	CQET (geom.)
<b>284826</b>	32	8	808080	128	22.687	54.1392	4385.675	554.324	831.28	34.26309	33.42904
<b>Query</b>	Times Executed	AQET [sec]	AQET(geom.) [sec]	QPS	minQET [sec]	maxQET [sec]	Average result [count/bytes]	Min result [count/bytes]	Max result count [count/bytes]	Timeouts	
<b>1</b>	128	0.116431	0.074423	67.95	0.043242	2.099115	8	1	10	0	
<b>2</b>	768	0.346561	0.303735	22.83	0.175449	2.274971	19.23	7	32	0	
<b>3</b>	128	0.085778	0.076927	92.24	0.04194	0.284563	6.88	0	10	0	
<b>4</b>	128	0.150389	0.138792	52.61	0.088126	0.439026	7.16	0	10	0	
<b>5</b>	128	10.321151	7.866798	0.77	0.941459	32.243805	5	5	5	0	

6	128	19.85562	19.68389	0.4	11.731131	35.253647	247.69	1	30401	0	
7	512	0.17432	0.158507	45.39	0.072345	2.156975	14.93	1	132	0	
8	256	0.099776	0.09129	79.3	0.015501	0.291483	4.85	0	15	0	
9	512	0.050274	0.041492	157.37	0.023552	2.018961	6771.43	2001	12744	0	
10	256	0.149327	0.110015	52.98	0.062341	2.210656	2	0	10	0	
11	128	0.031363	0.024837	252.26	0.014916	0.233921	10	10	10	0	
12	128	0.226412	0.174637	34.94	0.09737	2.272532	1425.98	1397	1460	0	
<b>Scale factor</b>	Number of warmup runs	Number of clients	Seed	Number of query mix runs (without warmups)	min Querymix runtime	max Querymix runtime	Total runtime (sum)	Total runtime (actual)	QMpH	CQ ET	CQET (geom.)
<b>284826</b>	32	16	808080	128	30.8348	126.2152	8365.895	537.44	857.4	65.35855	63.47983
<b>Query</b>	Times Executed	AQET [sec]	AQET(geom.) [sec]	QPS	minQET [sec]	maxQET [sec]	Average result [count/bytes]	Min result [count/bytes]	Max result count [count/bytes]	Timeouts	
<b>1</b>	128	0.177424	0.124529	87.73	0.046721	2.783055	8	1	10	0	
<b>2</b>	768	0.581124	0.454545	26.79	0.180335	3.973885	19.23	7	32	0	

<b>3</b>	128	0.2545 36	0.1511 39	61.16	0.055248	2.609839	6.88	0	10	0	
<b>4</b>	128	0.3148 25	0.2304 39	49.44	0.096767	3.002465	7.16	0	10	0	
<b>5</b>	128	19.786 027	14.993 473	0.79	1.834005	63.390338	5	5	5	0	
<b>6</b>	128	38.863 609	38.111 776	0.4	14.1186	108.55968	247.69	1	30401	0	
<b>7</b>	512	0.3017 33	0.2414 91	51.59	0.084788	2.958027	14.93	1	132	0	
<b>8</b>	256	0.1570 71	0.1340 4	99.1	0.024879	2.272149	4.85	0	15	0	
<b>9</b>	512	0.0921 19	0.0822 52	168.98	0.023079	0.390835	6771.43	2001	12744	0	
<b>10</b>	256	0.1360 07	0.1217 71	114.45	0.064442	0.399238	2	0	10	0	
<b>11</b>	128	0.0423 49	0.0361 43	367.57	0.014991	0.267267	10	10	10	0	
<b>12</b>	128	0.2714 76	0.2191 63	57.34	0.114146	3.08111	1425.98	1397	1460	0	