Improving Altibase Performance with Solarflare 10GbE Server Adapters and OpenOnload

SolarflareSolutionBrief

Summary

As today's corporations process more and more data, the business ramifications of faster and more resilient database management systems can be staggering. Tens of processing hours can result in the loss of hundreds of millions of dollars and system outages can cause severe revenue loss, security threats, and eroded customer confidence leading to additional lost business. The Altibase HDB™ database addresses these needs with higher performance and availability.

To demonstrate significantly improved performance using Solarflare 10GbE server adapters and OpenOnload® application acceleration middleware, Altibase conducted benchmark and replication performance tests. The results showed that database benchmark performance increased 38% using Solarflare SFN7122F 10GbE server adapters and the Linux kernel network driver compared with the performance using 1GbE server adapters. When using OpenOnload with the SFN7122F adapters, database benchmark performance increased an additional 20%.

High availability performance testing with asynchronous replication and the Solarflare SFN7122F 10GbE server adapters and the Linux kernel network driver resulted in a 21% increase compared with the performance using 1GbE server adapters. When using OpenOnload with the SFN7122F adapters, asynchronous replication performance increased an additional 27%. Additional asynchronous replication testing verified that OpenOnload performance also scales with parallel replication.

The following Solution Brief documents the testing configurations and methodologies. The benchmarks conclusively demonstrated that the performance of Altibase HDB's asynchronous replication improved drastically when using Solarflare SFN7122F 10GbE server adapters with OpenOnload application acceleration middleware.

Background

Altibase

Since 1999, Altibase® has pioneered in-memory databases and in-memory computing with a rich installation base consisting of over 550 global enterprise clients and thousands of mission critical deployments. As the go-to source for mission critical in-memory for the fastest databases (IMDBMS or MMDBMS), Altibase is the only company to offer a feature rich, ACID compliant database that combines the unparalleled speed of in-memory computing along with the economical scalability of traditional disks in a single unified engine.

sales@solarflare.com

US 1.949.581.6830 x2930

UK +44 (0)1223 477171

HK +852 2624-8868

www.solarflare.com

SolarflareSolutionBrief

Altibase HDB

The Altibase HDB is a hybrid database that combines an in-memory database and on-disk database in a single uniform engine and interface, enabling real-time access to large volumes of data while simultaneously simplifying and revolutionizing high-intensity data processing. While Altibase HDB is SQL standards compliant and adheres to all fundamental principles of traditional relational databases, its unique hybrid architecture allows users to enhance performance significantly when needed. HDB is several times faster than the leading in-memory only databases and 10 times faster than conventional on-disk databases.

HDB combines several key services such as high-availability, fault tolerance and load balancing with its replication feature that maintains an up-to-date backup of the database on an active server. In the event the primary server becomes unavailable, services immediately resume from an identical database on an alternate server for a non-stop operating environment. The replication is a TCP/IP based service, supports up to 32-way replication, and is available with a choice of two modes of replication: asynchronous and synchronous.

Asynchronous Mode Replication

Asynchronous mode focuses on high performance and is the default mode of Altibase HDB replication. The term asynchronous signifies that the local server does not wait until a remote server is done applying a transaction. **Figure 1** below illustrates the normal flow of tasks performed in asynchronous mode.

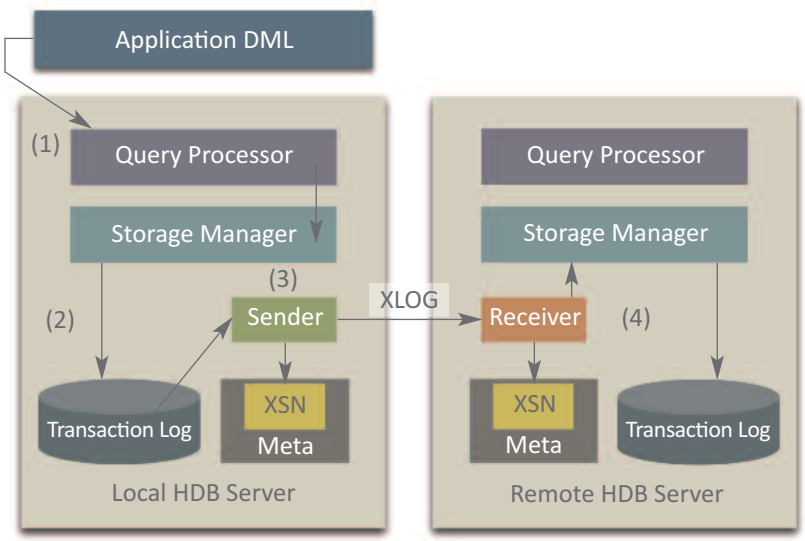

Figure 1. Asynchronous Mode Replication Task Flow.

SOLARFLARE®

ALTIBASE

sales@solarflare.com

US 1.949.581.6830 x2930

UK +44 (0)1223 477171

HK +852 2624-8868

www.solarflare.com

Asynchronous Mode Replication Task Flow (previous page)

1. An application transaction occurs on a local server (master transaction), and a DML statement is executed on a table that is a target for replication.
2. Altibase HDB applies the changes to the data from the master transaction and creates the transaction log records.
3. The sender thread collects transaction logs recorded by the master transaction, converts them into XLOGs, records the XSN in the replication metadata table, and sends XLOGs out to the receiver thread on the remote server. On the local server, Altibase HDB commits the master transaction, and continues processing other transactions without waiting for any information related to the replication.
4. The receiver thread receives the XLOGs and commits the replicated transaction to its database instance on the remote server.

Synchronous Mode Replication

While asynchronous mode focuses on high performance, synchronous mode puts the focus on data consistency. Synchronous mode keeps all replication targets synchronized by updating all nodes as part of a single atomic transaction. In synchronous mode, when a master transaction occurs on a local server, the local server commits the transaction only after it has received confirmation that each of the change logs has been properly applied on the remote server(s). Synchronous mode guarantees 100% data integrity and is recommended if high availability of a database is the main focus, and an Active-Active architecture implementation is planned. **Figure 2** illustrates the normal flow tasks performed in synchronous mode replication.

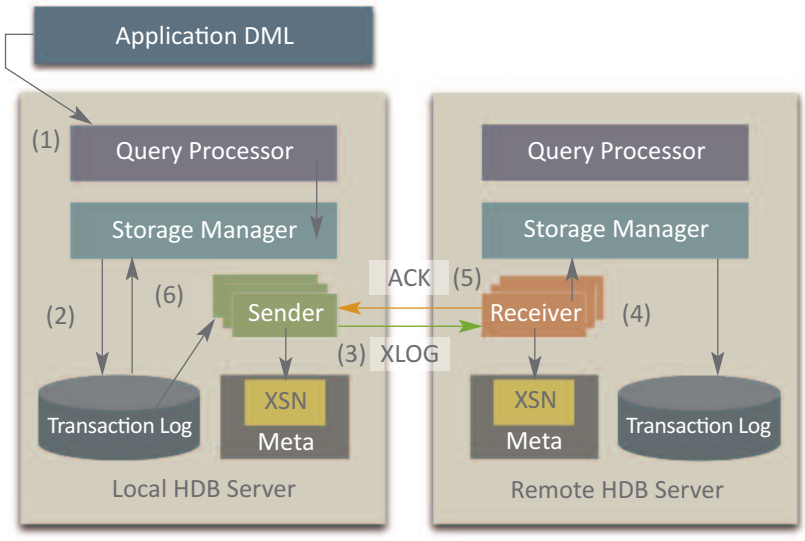

Figure 2. Synchronous Mode Replication Task Flow.

sales@solarflare.com

US 1.949.581.6830 x2930

UK +44 (0)1223 477171

HK +852 2624-8868

www.solarflare.com

Solarflare

Solarflare is the leading provider of application-intelligent networking I/O software and hardware that accelerate, monitor and secure network data, and is the pioneer in high-performance, low-latency 10/40GbE server networking solutions. The company's products are widely used in scale-out server environments such as database, grid, electronic trading, high performance computing, cloud, virtualization and big data.

Solarflare Flareon® Server I/O Adapters

Solarflare Flareon 10/40GbE PCIe 3.0 server I/O adapters deliver industry-leading message rates with the lowest latency and jitter over standard Ethernet along with low CPU-utilization, enabling the industry's best performance and scalability for enterprise data center environments including database applications.

Figure 3. Solarflare Flareon Ultra SFN7122F 10GbE PCIe 3.0 Server I/O Adapter.

OpenOnload®

OpenOnload application acceleration middleware is an accelerated network stack that provides acceleration of TCP- and UDP-based applications. OpenOnload comprises a user-level shared library that implements the protocol stack and a supporting kernel module. By operating in user space, OpenOnload dramatically reduces CPU interrupts, data copies and context switching – resulting in reduced latency and higher message rates. OpenOnload is binary compatible with the industry standard BSD Sockets API, thereby providing acceleration with no need to run a new protocol such as RDMA, RoCE or iWARP, on the wire. OpenOnload is ideal for applications that benefit from lower latency (with decreasing jitter) and higher throughput, including financial services, geosciences, biosciences, weather and other HPC applications, and high performance databases.

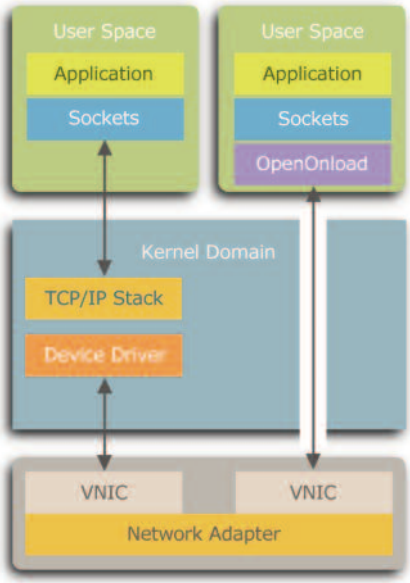

Figure 4. OpenOnload.

SOLARFLARE®

ALTIBASE

sales@solarflare.com

US 1.949.581.6830 x2930

UK +44 (0)1223 477171

HK +852 2624-8868

www.solarflare.com

SolarflareSolutionBrief

Test Objective

Altibase conducted a series of performance tests with two separate test beds to validate that:

- Solarflare Flareon 10GbE server adapters and OpenOnload can improve benchmark and replication performance with Altibase HDB, and
- Solarflare OpenOnload performance improvement can scale with the parallel replication option

Test Bed A

Table 1. System Under Test

Component	Description
Server	Intel® Xeon® E5 processor@2.4GhZ and 6 cores
OS	CentOS release 6.3 (final) with Linux kernel 2.6.32-279.el6.x86_64
Database	Altibase HDB version 6.3.1.2.2
Network Adapters	1 Gigabit adapter
	Solarflare Flareon Ultra SFN7122F 10GbE SFP+ server adapter with OpenOnload

Configuration

- The SUT was configured as both the active database and the clients
- Logging is to disk

Test Codes

- Standard BMT
- AT-P26-Replication
- AT-P55-ReplicationSync

Standard BMT Benchmark

The BMT benchmark tests the performance characteristics of Altibase HDB on the testing environment. The BMT test consists of a simple C program that performs a user provided number of SELECT, INSERT, UPDATE and DELETE operations against the respective database.

sales@solarflare.com
US 1.949.581.6830 x2930
UK +44 (0)1223 477171
HK +852 2624-8868
www.solarflare.com

Schema

```
create table tb_test1 (
     c1 integer primary key,
     c2 char(500),
     c3 integer,
     c4 char(500),
     c5 char(500),
     c6 char(500),
     c7 char(500),
     c8 char(500),
     c9 integer,
     c10 date
);
Create index TB_TEST1_IDX on TB_TEST1 (c1, c3);
```

Test Bed A Test Results

Standard BMT Benchmark

For this benchmark:

- 8 clients were used
- INSERT and DELETE of tables were performed
- Transaction: Commit/10 statements
- Results are in TPS (Transactions per Second)
- 1G: 1Gigabit Ethernet adapter
- 10G: Solarflare SFN7122F 10GbE adapter
- Onload: SFN7122F with OpenOnload

	1G	10G	Onload	Ratio 10G vs. 1G	Ratio Onload vs. 1G
Insert	45,119	60,127	69,105	133%	153%
Delete	52,180	72,113	86,493	138%	166%

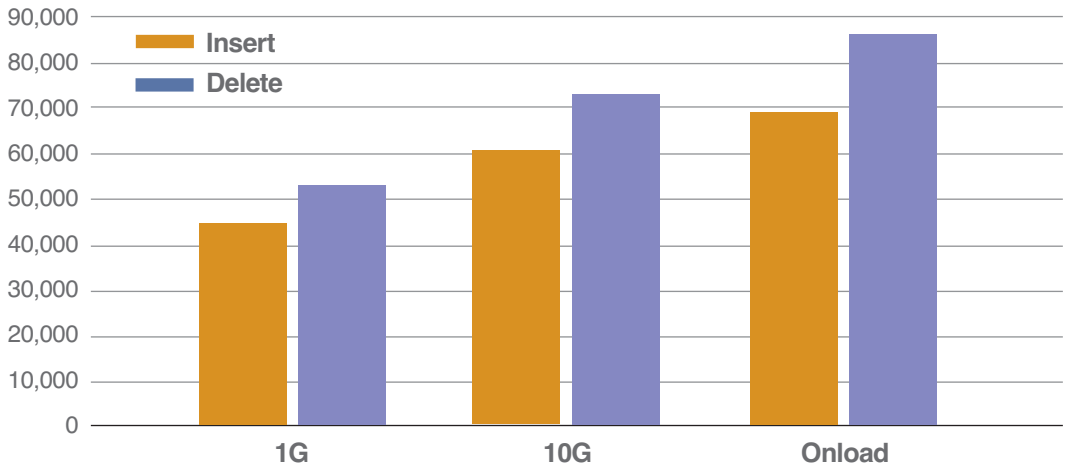

Figure 5. Test bed "A" Standard BMT Benchmark Results.

sales@solarflare.com

US 1.949.581.6830 x2930

UK +44 (0)1223 477171

HK +852 2624-8868

www.solarflare.com

SolarflareSolutionBrief

As shown in the results, the Flareon SFN7122F network adapter showed a performance increase of 33% for INSERT and 38% for DELETE compared with using a 1 Gigabit Ethernet adapter. When OpenOnload kernel bypass middleware was used, the Flareon SFN7122F performance improvement increased to 53% for INSERT and 66% for DELETE compared with the performance using a 1 Gigabit Ethernet adapter.

Replication Test Codes

One of Altibase HDB's key features is its built in replication engine. This feature can be utilized both for load balancing and high availability purposes.

AT-P26-Replication
Schema

```
create table tb_test1 (
     c1 integer primary key,
     c2 char(10),
     c3 integer,
     c4 char(255),
     c5 char(255),
     c6 char(255),
     c7 char(10),
     c8 char(10),
     c9 integer,
     c10 date
);
Primary key(c1, c2);
T1_idx on t1(c10);
```

For this benchmark:

- 8 clients were used
- INSERT and DELETE of tables were performed
- Transaction: Commit/10 statements
- Results are in TPS (Transactions per Second)
- 1G: 1Gigabit Ethernet adapter
- 10G: Solarflare SFN7122F 10GbE adapter
- Onload: SFN7122F with OpenOnload

	1G	10G	Onload
Insert TPS	81,623	83,847	82,974
Replication TPS	36,372	43,841	55,776
Replication Performance Ratio	45%	52%	67%
1G vs. 10G	100%	121%	
1G vs. Onload	100%		153%

SOLARFLARE®

ALTIBASE

sales@solarflare.com

US 1.949.581.6830 x2930

UK +44 (0)1223 477171

HK +852 2624-8868

www.solarflare.com

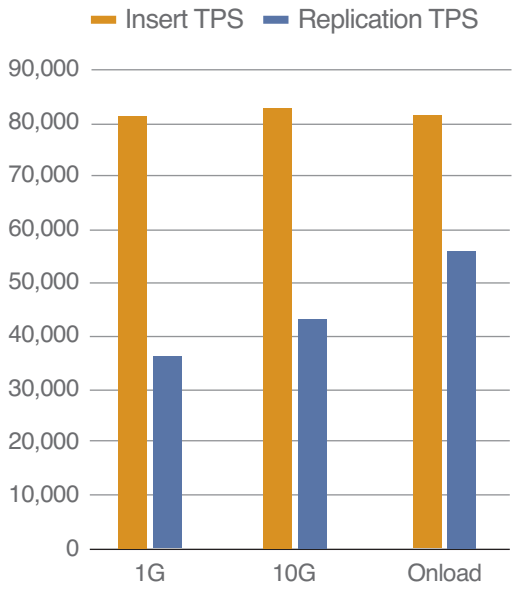

Figure 6. Test bed "A" AT-P26-Replications Standby Server Test.

As shown in the results, the Flareon SFN7122F 10GbE server adapters provided a 21% asynchronous replication performance improvement compared with using 1GbE server adapters. When using OpenOnload with the SFN7122F adapters, performance improvement increased to 53% compared with 1GbE server adapters.

AT-P55-Replication

Schema

```
create table t1(
     c1 integer,
     c2 char(10),
     c3 integer,
     c4 char(255),
     c5 char(255),
     c6 char(255),
     c7 char(10),
     c8 char(10),
     c9 integer,
     c10 date
);
Primary key(c1, c2);
T1_iex on t1(c10);
```

For this benchmark:

- Data: 10 tables with 100,000 records per table
- Transaction: Commit / 50000 record insert
- 1G: 1 Gigabit Ethernet adapter
- 10G: Solarflare SFN7122F 10GbE adapter
- Onload: SFN7122F with OpenOnload

SOLARFLARE®

ALTIBASE

sales@solarflare.com

US 1.949.581.6830 x2930

UK +44 (0)1223 477171

HK +852 2624-8868

www.solarflare.com

SolarflareSolutionBrief

Thread	1G	10G	Onload	Ratio Onload vs 1G
1	82,386	82,339	87,958	107%
8	90,090	163,949	163,618	182%

Table 2. Test bed "A" AT-P55-Replication Test Results.

Unit: RPS (replicated records number per second)

There was a large performance improvement when using the SFN7122F 10GbE server adapter. But there were no observable performance improvements when using the SFN7122F with OpenOnload. This implies that the performance increase was due to increased bandwidth.

Test Bed B

Target Server	Description
Server	Intel® Xeon® E5 -2605 processor @ 2.4GHz and 8 cores with 64GB memory
OS	CentOS 2.6.32
Database	Altibase HDB version 6.5.1.0.0
Client	**Description**
Server	Intel® Xeon® E5603 processor @ 2.53GHz and 16 cores with 64GB memory
OS	Linux 2.6.18
Network	**Description**
Solarflare	Flareon Ultra SFN7122F 10GbE SFP+ server adapter with OpenOnload
Intel	1GbE Adapter

Table 3. Test bed "B" System Under Test.

Configuration

- 1 transaction consists of one commit per 10 insert records
- 96 clients

Test Codes

- AT-P26-Replication

sales@solarflare.com

US 1.949.581.6830 x2930

UK +44 (0)1223 477171

HK +852 2624-8868

www.solarflare.com

AT-P26-Replication Schema

```
create table tb_test1
(
     c1 integer primary key,
     c2 char(10),
     c3 integer,
     c4 char(10),
     c5 char(10),
     c6 char(10),
     c7 char(10),
     c8 char(10),
     c9 integer,
     c10 date );
```

Test Bed B Test Results

	No Parallel (10G w/o Onload)	No Parallel (10G with Onload)	6 Parallel (10G w/o Onload)	6 Parallel (10G with Onload)
Replication performance improvement rate	100%(base)	132%	152%	198%

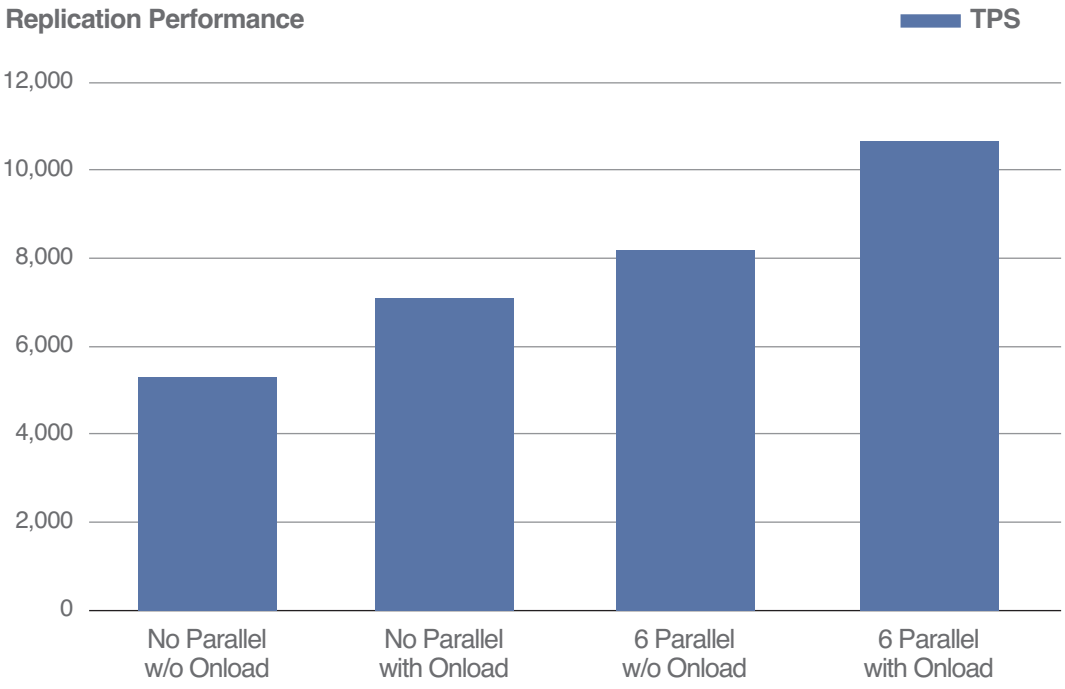

In Test bed “B”, the test results showed that OpenOnload performance improvement for asynchronous replication scales with increased parallel replication option processing. Without the parallel replication option, replication performance with OpenOnload increased by 32% compared with using the kernel driver. With 6 parallel replication threads, performance with OpenOnload increased by 30% compared with using the Linux kernel driver.

SOLARFLARE®

ALTIBASE

sales@solarflare.com

US 1.949.581.6830 x2930

UK +44 (0)1223 477171

HK +852 2624-8868

www.solarflare.com

SF-114775-CD Issue 1
AB_SB_091815