



WHITE PAPER

Accelerating Apache ActiveMQ with SanDisk® Flash Storage and Torusware Speedus

April 2015



SanDisk®
a Western Digital brand

Western Digital Technologies, Inc.
951 SanDisk Drive, Milpitas, CA 95035

www.SanDisk.com

Table of Contents

Executive Overview	3
Needs Assessment	3
ActiveMQ Applications	4
Enterprise Service Bus	4
Transactional Messaging	4
High-Performance Market Data Distribution	4
Mission-Critical Deployments	4
Web Streaming of Data	4
Infrastructure for the Internet of Things	4
ActiveMQ Performance and Scalability	5
The SanDisk and Torusware Solution	5
About SanDisk and Torusware	6
References	7
Sonic Test Harness Configuration Example	7

Executive Overview

The performance and scalability of messaging middleware, a central part of most software architectures, is key for a wide range of organizations. In these scenarios, the main performance bottlenecks are generally I/O operations, namely storage I/O and communications I/O. This limitation can be significantly alleviated by the combination of SanDisk flash storage and Torusware Speedus, which are able to increase significantly the performance of Apache ActiveMQ™, the leading open-source messaging library. This whitepaper reports on our claim that by significantly reducing I/O overhead, it is possible to increase performance and scalability of messaging middleware, while reducing risks and costs of enterprise integration.

Needs Assessment

Messaging middleware is the primary choice to reliably deliver messages among applications. Popular messaging middleware are Apache ActiveMQ, IBM WebSphere MQ, Tibco Enterprise Message Service™, Progress SonicMQ, and Pivotal RabbitMQ™. Middleware performance and scalability is key for the overall system as all parties depend upon it, and generally the main performance bottleneck is I/O, namely storage I/O and communications I/O. This whitepaper reports our findings on optimizing ActiveMQ by means of SanDisk flash storage and Torusware Speedus, with performance improvements of up to an order of magnitude.

Apache ActiveMQ is one of the most popular and powerful open-source messaging and Integration Patterns server currently available, thanks to its speed, full support of JMS 1.1 and J2EE 1.4, and cross languages and protocols support. Typical Apache ActiveMQ deployments are found in enterprise service bus (ESB) deployments, transactional messaging, high-performance market data distribution, messaging middleware across Web/HTTP clusters (via a RESTful API), and web streaming of data, such as live IM conversations, or dynamically updated live web content such as news or stock prices.

SanDisk heritage of innovation in flash storage technologies has led to the development of SanDisk flash-optimized software technology that enables applications to capitalize on the benefit of flash scalability and performance. Multi-core processors allow parallel execution of multiple threads. But most software was not written to exploit multi-core processors or flash memory. Cores go unused. Flash IOPS go unexploited. Precious IT budget is inefficiently used.

Torusware develops and commercializes Torusware Speedus, a low-latency high-performance TCP socket library. Torusware Speedus allows C/C++/Java applications to exchange messages in less than 250 nanoseconds—critical for latency-sensitive applications, such as in capital markets and real-time analytics.

The combination of SanDisk and Torusware technologies offers cost-effective application scaling by enabling DRAM-like speed for the storage while removing most of the communications overhead, thus being able to fully exploit the capability of multi-core processor's parallel execution of multiple threads. Superior overall performance can be achieved while minimizing infrastructure investment, being able to provide persistent storage at minimum performance cost. This SanDisk and Torusware joint solution delivers higher availability for mission-critical workloads with persistence at very high data throughput.

ActiveMQ Applications

Apache ActiveMQ has a wide range of applications with diverse persistence and performance requirements .

Enterprise Service Bus

ActiveMQ is a key component of ESB solutions such as Apache ServiceMix, Apache Camel and Mule. ESBs are lightweight frameworks that provide flexible, loosely coupled integration across a broad range of enterprise applications. The ESB enterprise messaging backbone, generally ActiveMQ, must provide guaranteed delivery of messages among services and endpoints, performing protocol conversions on the fly, and even fault tolerance with near real-time failover in case of system faults . The messaging system acts as the nervous system of the ESB, dictating its overall performance and the rate of the business requests delivered by the system .

Transactional Messaging

ActiveMQ has applicability in transactional messaging, a use case focused on transactionality, persistence and guaranteed delivery . Here clusters of servers are required to increase availability . Furthermore, if a node is offline, then all the messages will be persisted for when the node comes back online .

High-Performance Market Data Distribution

This use case demands high throughput and performance for the distribution of massive volumes of rapidly changing data as fast as possible . Furthermore, persistence is generally not required, as old data can be ignored; the latest prices are what matters, so real-time low-latency solutions are a design principle.

Mission-Critical Deployments

ActiveMQ is used in really mission critical deployments, such as in the Next Generation Air Transportation System of the U .S. Federal Aviation Administration (FAA), helping increase security and reduce delays in the most congested airports . Here real-time low-latency solutions are a must .

Web Streaming of Data

This use case focuses on AJAX support in ActiveMQ, supporting real-time data streaming into web browsers . Typical examples are streaming stock prices, live IM conversations or dynamically updated live content and news. Here, ActiveMQ is integrated into web containers, communicating web browsers via HTTP GET/POST operations . A combination of persistence with low latency is required, to be able to cope with disconnected IM clients and real-time updates of a high-volume of web browsers, respectively .

Infrastructure for the Internet of Things

ActiveMQ delivers information reliably between multiple devices, systems and applications, both on-premises or through the (hybrid) cloud . All assets will be connected, creating efficient and smart solutions in the Internet of Things . Here the infrastructure must cope with high volumes of messages, both incoming and outgoing .

ActiveMQ Performance and Scalability

Apache ActiveMQ performance and scalability is limited by storage and network I/O, both for horizontal and especially for vertical scaling . Messaging applications that are concerned about guaranteed message delivery must specify messages as persistent, although it comes at the cost of storage or database overhead. Furthermore, even non-persistence configurations, when running out of memory (a common scenario when storing messages for disconnected clients), swap messages out of active memory to storage.

The more reliable the message delivery, the more overhead is required to achieve it . This trade-off between reliability and performance is a key design consideration. Performance is maximized avoiding persistence, while reliability is maximized by activating persistency and transacted sessions.

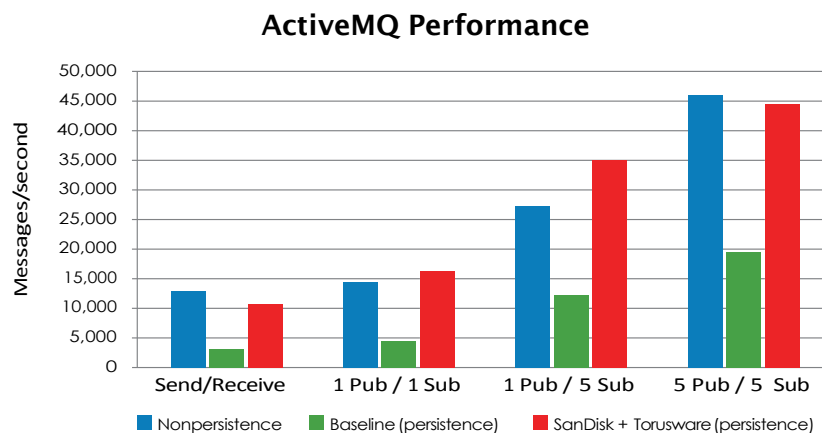
Apache ActiveMQ is highly flexible; it can be used in many different use cases, with multiple requirements and deployment options. Therefore, to adapt to all these scenarios, ActiveMQ is highly configurable . However, default parameters might not be optimal for a number of scenarios, so fine tuning is always a resort to be considered. However, it might be difficult to achieve performance gains across different infrastructure and varying message traffic conditions.

The SanDisk and Torusware Solution

Apache ActiveMQ performance and scalability can be increased by removing the main sources of overhead, storage and network I/O, by means of SanDisk flash storage and Torusware Speedus, respectively. To demonstrate the benefits of this joint solution, we conducted an internal test using Sonic Harness ActiveMQ benchmark .

Here are the experimental configuration details:

- Hardware configuration: HP ProLiant server SL250s Gen8, with two Intel® 64GB E5-2660 processors, one 1TB 7.2k 2.5in 6G SATA HDD, one 500GB Optimus Eco® SSD
- Apache ActiveMQ 5.10.0 with OpenJDK JVM 1.7.0.55 on CentOS 6.4
- Torusware Speedus v1.0
- Message sizes are 1KB in length, text, generated using ConstantMsgGenerator
- Persistence using the default configuration, KahaDB with sync writes
- SonicMQ Progress Test Harness v8.0.2



The figure above shows the performance of Apache ActiveMQ on four configurations, a send/receive with a single queue, and three publisher/subscriber (Pub/Sub) scenarios: first, a single publisher with a single subscriber, second, a single publisher with five subscribers (one topic), and finally, five publishers with five subscribers (five topics) . All four configuration are using transacted delivery .

The results show that the combination of SanDisk and Torusware rivals and even outperforms non-persistence numbers . Therefore, it is possible to enable persistence without suffering performance losses . This is key to all use cases where guaranteed delivery is essential for the business process, such as ESB, transactional messaging, mission critical deployments, and some scenarios in web streaming of data and the Internet of Things.

Moreover, when comparing baseline (persistence) results, which are backed by a HDD and rely on TCP/IP protocol stack, the introduction of flash storage and Torusware Speedus increases performance between 125% and 265% . This higher throughput supports significantly faster message delivery, particularly when a node requests a significant number of messages (e.g., after a restart).

Publisher/subscriber configurations, which are able to reach 45,000 messages per second, are key to deal with message replication and fault tolerance, as well as multiple communication endpoints, such as in Web streaming of data and the Internet of Things.

About SanDisk and Torusware

At SanDisk we have a singular focus: to expand the possibilities of data storage technology . And we're able to achieve it with a unique approach .

Because storage is our primary concern, SanDisk has the perspective needed to ensure data protection, while balancing reliability and accessibility to suit unique enterprise needs . Our vertical integration—from our controller design to our flash management algorithms and software acceleration solutions—sets SanDisk apart from most of our competitors . From our own flash manufacturing facilities to our extensive portfolio of products, we have complete top-to-bottom capabilities . This unique integration results in a highly reliable enterprise system composed of innovative technologies that work together to benefit your business.

A broad portfolio of data center storage solutions based on end-to-end capabilities is just one reason why enterprise IT and system vendors alike trust SanDisk with their data . Learn more about SanDisk data center storage solutions for the enterprise at www.sandisk.com/enterprise .

Torusware engineers hardware and software to work efficiently in your data center and the cloud . Torusware improves predictability and performance of latency-sensitive applications through low-latency and efficient solutions based on a combination of a high-performance socket implementation, Torusware Speedus, with a deep knowledge on gluing latency-critical software to hardware, thus bridging the performance gap between applications and infrastructure. The integration of Torusware solutions is non-intrusive, no application recoding is required, and it is interoperable with currently existing communication protocols. For more information about Torusware visit www.torusware.com.

References

Sonic Test Harness Benchmark Website: https://community.progress.com/community_groups/technology_partner/w/technologypartner/401.sonic-test-harness.aspx

Progress Software . "Benchmarking Enterprise Messaging Systems":
www.dcs.bbk.ac.uk/~geoff/benchmarking_enterprise_msg_wp.pdf

Sonic Test Harness Configuration Example

These parameters have been modified according each particular test requirement:

```
# Simple point to point example for TestHarness
```

```
mode=PUB
```

```
# **** Connection/Session info ****
```

```
numconnections=1
```

```
numsessions=1
```

```
destname=test1-topic
```

```
numdests=1
```

```
istransacted=true
```

```
msgpertxn=1
```

```
# **** Producer options ****
```

```
deliverymode=PERS
```

```
msgpriority=4
```

```
isreplyto=false
```

```
replytodest=none
```

```
# **** Message generation options ****
```

```
msgtype=TEXT
```

```
msgsize=1k
```

```
msggenclass=ConstantMsgGenerator
```

```
msgproperties=none
```

```
msgcache=1
```

```
# **** Consumer options ****
```

```
durable=true
```

```
keepdurables=false
```

```
ackmode=auto
```

```
isforward=false
```

```
msgselector=none
```

```
recordercount=0
```

```
# **** Test timing and measurement ****
```

```
sleepmsecs=0
```

```
numrptintervals=10
```

```
rptintervalsecs=30
```

```
starter=prompt
```

```
dowarmup=true
```

```
nummessages=0
```

```
checklatency=false
```

```
checkmsgsize=false
```

```
resultsfile=none
```

```
# **** Connection and user specs ****
```

```
jndiurl=none
```

```
user=none
```

```
password=none
```

```
jmsuser=none
```

```
jmspassword=none
```

```
clientid=SimpleSender
```

Specifications are subject to change. ©2015 - 2016 Western Digital Corporation or its affiliates. All rights reserved. SanDisk and the SanDisk logo are trademarks of Western Digital Corporation or its affiliates, registered in the U.S. and other countries. Optimus is a registered trademark of Western Digital Corporation or its affiliates. Other brand names mentioned herein are for identification purposes only and may be the trademarks of their respective holder(s). 20160623

Western Digital Technologies, Inc. is the seller of record and licensee in the Americas of SanDisk® products.