



# Diamanti Lab Validation

An IDC Lab Validation, *Sponsored by Diamanti* | **June 2020**

# Executive Summary

Containers are a modern computing architecture that is rapidly growing in popularity due to its lightweight, agile, and highly automatable nature that is ideal to support cloud native, microservices applications. The industry has standardized on Kubernetes as the container orchestration and management platform, which has robust capabilities for application deployment patterns, scalability, and integration with networking and storage. However, while Kubernetes is extremely powerful, it is also very complex and requires other elements to make it into a fully operable enterprise platform, such as a management interface, security and access controls, and integrated networking and storage subsystems. While early adopters of Kubernetes focused on stateless applications, enterprises are quickly seeing the benefits of containers and Kubernetes for the rapid development and portability of stateful applications such as databases. These can be run in containers, but present additional challenges in workload management and persistent storage, including considerations regarding data protection and disaster recovery.

## As enterprises adopt containers they are seeking solutions that:

- Provide a full enterprise class platform built around Kubernetes that fully integrate functions such as management, networking, storage, and identity/access management.
- A turnkey platform that makes Kubernetes easier to deploy and manage.
- Support for a wide range of workloads, particularly stateful applications that require persistent storage and advanced data services
- Support agile hybrid cloud deployments that can span multiple datacenters and multiple public clouds.

In this IDC Lab Validation, IDC validated key features and functionality of the Diamanti platform, an integrated hardware and software solution built for Kubernetes. In this evaluation, we verified the key features and performance of the MongoDB Atlas cloud service and MongoDB benchmark workloads running on the Diamanti platform, including:

- Performance of a MongoDB cluster on the Diamanti platform versus the MongoDB Atlas public cloud service.
- Management of MongoDB using MongoDB Cloud Manager, an application for managing MongoDB databases in either cloud-based or on-premises deployments, such as the Diamanti platform, that can operate either as a locally deployed program or as a SaaS.

# Diamanti Platform Overview

The Diamanti platform is a hyperconverged infrastructure designed for modern containerized applications. The I/O optimized hardware is pre-integrated with Kubernetes software for a unified hardware and software solution.

The Diamanti platform uses Kubernetes CSI and CNI drivers to interface to NVMe SSD storage and SR-IOV virtual function networking, both of which are accelerated by dedicated PCI-e processors. The system also integrates a management console which allows users to manage functions such as QoS, multi-zone clustering, data protection, monitoring, and RBAC.

A partnership with MongoDB enables hybrid cloud deployments using SaaS based management with MongoDB Atlas, and the equivalent MongoDB Cloud Manager for on-prem deployment. Workloads that require higher performance, consistency and scalability can be targeted for the Diamanti platform, while maintaining the familiar UI of the MongoDB management and analytics engine.



## Benefits of the Diamanti approach:

- Faster deployment as Kubernetes is pre-installed and integrated with the hardware
- Scale faster by easily adding more nodes
- Full stack support (hardware and software) from a single vendor
- Operationalize Kubernetes with Diamanti's Kubernetes management tools
- Optimized performance with bare metal containers and accelerated networking and storage subsystems

## The specifications for the D20 units used for this lab validation are as follows (per node):

### Compute

- 2x Intel Xeon Silver 4114 2.22 GHz processors with 20 cores total (10 cores per CPU)
- 192GB RAM

### Storage

- Diamanti Ultima Storage Card
- 4 TB data storage (4x 1000GB NVMe SSDs)
- 960GB host OS and container image storage (2x 480GB SATA SSDs)

### Network

- Diamanti Ultima Network Card with 4x10 GbE via QSFP+ module

### Software

- Diamanti Spektra pre-validated and pre-packaged with Kubernetes 1.15, Docker 1.13 and CentOS Linux operating system

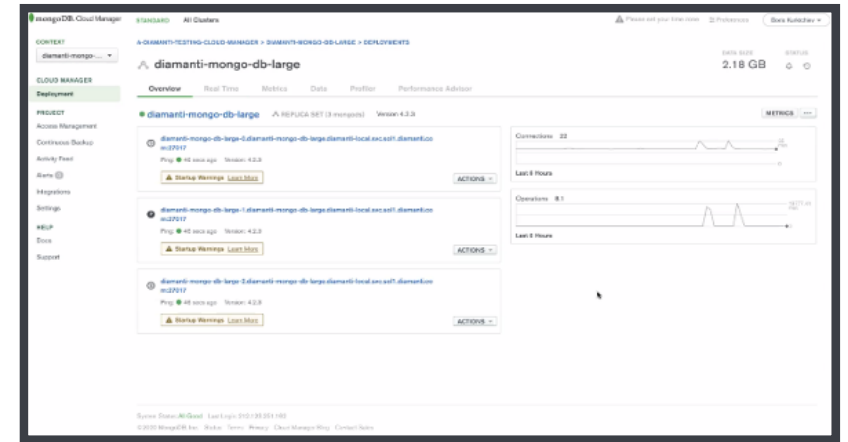
# MongoDB and the Diamanti Platform

## Key Findings

Manage and monitor on-premises instances of MongoDB with a SaaS based management platform.

## Why does it matter?

- MongoDB is a very popular NoSQL DBMS. But as it is scaled up, resource management can be complicated. In addition, many enterprises are moving to hybrid cloud and wish to deploy MongoDB in Kubernetes-governed containers, which can be a tricky proposition due to the skills and expertise required in building and maintaining a Kubernetes container infrastructure. Diamanti makes this option as simple to manage as the MongoDB DBaaS offering, MongoDB Atlas.
- Using this environment, users can gain insight into the performance and health of on-premises MongoDB with a convenient SaaS-like web application.
- MongoDB Cloud Manager and the MongoDB Atlas cloud service are both developed by MongoDB and use the same UI, allowing management of cloud and on-premises MongoDB instances with similar tools.
- Management of both cloud and on-premises MongoDB instances with similar tooling allows customers to deploy in a hybrid scenario and optimize deployment location based on varying requirements, such as security, cost, and performance.



## Validation process

IDC observed the use of the MongoDB Atlas service to provision MongoDB clusters to the public cloud and subsequently manage and monitor them. IDC also observed the use of the MongoDB Cloud Manager to manage and monitor MongoDB instances running locally on the Diamanti platform. In both cases, the UI shared the same design and offered the same management and monitoring functionality in a hosted management platform. In particular, IDC observed the real time monitoring of the performance of MongoDB clusters, with graphical visualizations.



# Performance Tests Summary

## Key findings

A comparison of operations per second between a MongoDB cluster provisioned in the public cloud and on a Diamanti cluster showed that the Diamanti cluster had 27.7x to 30.8x higher throughput with the small instances and 32.6x to 33.3x higher throughput with the large instances using the YCSB open source performance benchmark with varying mixtures of read and write operations. The Diamanti instances also showed more consistent throughput and latency compared with the Atlas cloud instance, as can be seen in comparing the variability of the performance graphs.

## Why does this matter?

For performance sensitive applications, database performance can be a key requirement. Customers may have a wide range of performance, security, and price requirements across all their databases and having hybrid options to distribute instances across both public cloud and dedicated on-premises hardware according to those requirements can help customers match the workload to the best and most cost effective deployment type. The simplicity of management also reduces the staff time required for the MongoDB database, and also lessens the risk of human error interfering with the application.

## Pricing considerations

Although the list prices for the configurations for this test are listed in the below tables, customers need to consider the following points when building cost comparison models:

- The cloud instances are available per hour and while we've calculated the yearly price, not all workloads may run 24x7 and scalability can also be dynamically adjusted with various auto scaling functions.
- The configuration used for the on-premises Diamanti tests was a 3 physical node configuration. While we've listed the costs for the initial acquisition and 1 year of any support and software subscriptions, the hardware costs would be amortized over the life of the server, which typically ranges from 3 to 5 years. Additionally, in our tests the nodes were only running our benchmark, leaving spare capacity on the nodes for other workloads. Unlike the public cloud, on-premises deployments have dedicated physical hardware that customers can pack and consolidate workloads as they see fit and this density will vary greatly depending on what workloads are being consolidated on the hardware.
- Staffing is also a key factor as a public cloud and on-premises deployments have very different mixes of what is managed by a provider and what is customer managed. However, on-premises deployments have more options available today that can offload some of the customer responsibility. For example in the Diamanti case, the solution consists of pre-integrated hardware and container software, which can eliminate some deployment work. Increasingly available SaaS based options, such as the MongoDB Cloud Manager, can shift database management into a cloud model, while keeping the database on-prem.

# Small Instance Performance Test

## Validation Process

IDC observed the performance of two 3 node MongoDB clusters, one deployed using MongoDB Atlas on the public cloud and the other on a 3 physical node Diamanti cluster. The specifications of each (per node) are:

	MongoDB Atlas (Cloud)	Diamanti
Compute	M40 VM instance with 4 CPU cores, 16GB RAM	Bare metal Docker container with 4 CPU cores, 16GB RAM
Storage	80GB (block storage volume, 240 IOPS)	80GB (virtualized NVMe SSDs through CSI driver)
Network	10GbE	10GbE (NIC through CNI driver using SR-IOV)

The list prices for the provisioned clusters are:

	MongoDB Atlas (Cloud)	Diamanti	
	3 node replica set with M40 instances	3x D20 Small nodes each with 20 CPU cores, 192GB RAM and 4TB of NVMe storage with Diamanti Spektra software	\$111,000 with one year standard software subscription
MongoDB Enterprise Advanced	Included	MongoDB Enterprise Advanced subscription - pricing is variable depending on configuration	Annual license per node or in a RAM pool model
MongoDB Cloud Manager	Included	MongoDB Cloud Manager Standard subscription	\$39/server/mo
Operations per second YCSB Benchmark	372 - 389	10,395 - 11,918	
Pricing	\$ 1.04 per hour or \$9110.40 per year (assuming 24x7 operation)	\$111,000 for one year + MongoDB Enterprise licensing	

Each MongoDB cluster was observed by IDC to run the Yahoo! Cloud Serving Benchmark (YCSB), an open source benchmarking software used to evaluate the performance of data serving systems, such as NoSQL databases. 2 runs were executed with different workload characteristics, a YCSB Workload A run (50-50 read/write test) and a YCSB Workload B run (95-5 read heavy test). The operations per second were monitored in real time using the MongoDB Atlas or MongoDB Cloud Manager console.

## Workload A (50-50 read/write)

Cloud



372 ops/sec

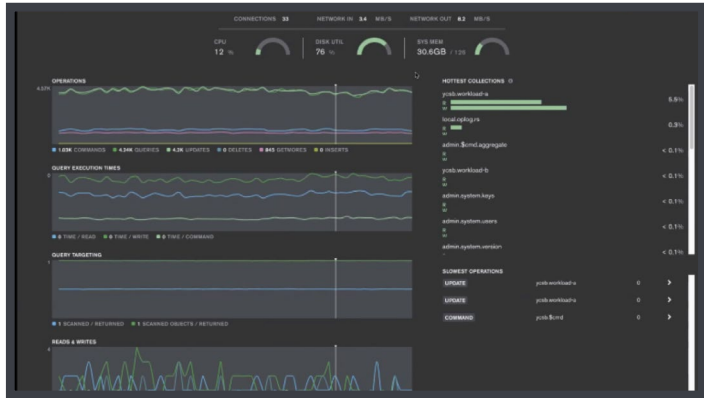
## Workload B (95-5 read heavy)

Cloud



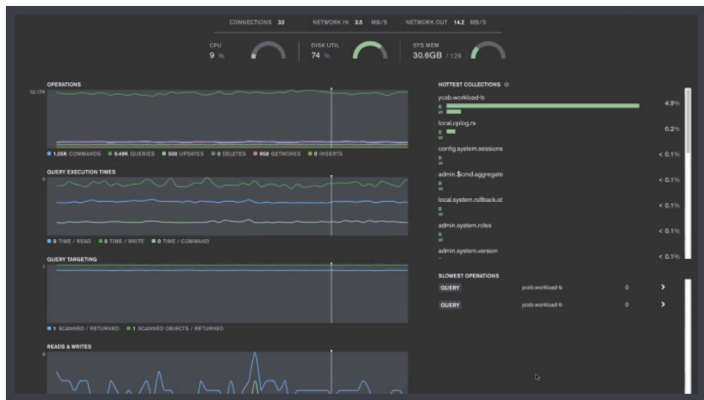
389 ops/s

Diamanti



10,315 ops/sec (27.7x improvement)

Diamanti



11,998 ops/s (30.8x improvement)

# Large Instance Performance Test

## Validation Process

IDC observed the performance of two 3 node MongoDB clusters, one deployed using MongoDB Atlas on the public cloud and the other on a 3 physical node Diamanti cluster. The specifications of each (per node) are:

	MongoDB Atlas (Cloud)	Diamanti
Compute	M140 VM instance with 48 CPU cores, 192GB RAM	Bare metal Docker container with 48 CPU cores, 192GB RAM
Storage	1TB (block storage volume, 3000 IOPS)	1TB (virtualized NVMe SSDs through CSI driver)
Network	10GbE	10GbE (NIC through CNI driver using SR-IOV)

The list prices for the provisioned clusters are:

	MongoDB Atlas (Cloud)	Diamanti	
	3 node replica set with M40 instances	3x D20 Small nodes each with 20 CPU cores, 192GB RAM and 4TB of NVMe storage with Diamanti Spektra software	\$111,000 with one year standard software subscription
MongoDB Enterprise Advanced	Included	MongoDB Enterprise Advanced subscription - pricing is variable depending on configuration	Annual license per node or in a RAM pool model
MongoDB Cloud Manager	Included	MongoDB Cloud Manager Standard subscription	\$39/server/mo
Operations per second YCSB Benchmark	372 - 380	12,398	
Pricing	\$96,272.40 per year (assuming 24x7 operation)	\$111,000 for one year + MongoDB Enterprise licensing	



It should be noted that both in this and the prior prices quoted in this document, these are only the prices associated with these tests, and are not necessarily reflective of other experiences. These prices can be used as a starting point. Organizations should also take into account factors such as volume discounting and, in the case of comparing cloud with on-prem costs, staffing costs.

Each MongoDB cluster was observed by IDC to run the Yahoo! Cloud Serving Benchmark (YCSB), an open source benchmarking software used to evaluate the performance of data serving systems, such as NoSQL databases. 2 runs were executed with different workload characteristics, a YCSB Workload A run (50-50 read/write test) and a YCSB Workload B run (95-5 read heavy test). The operations per second were monitored in real time using the MongoDB Atlas or MongoDB Cloud Manager console.

## Workload A (50-50 read/write)

### Cloud



372 ops/sec

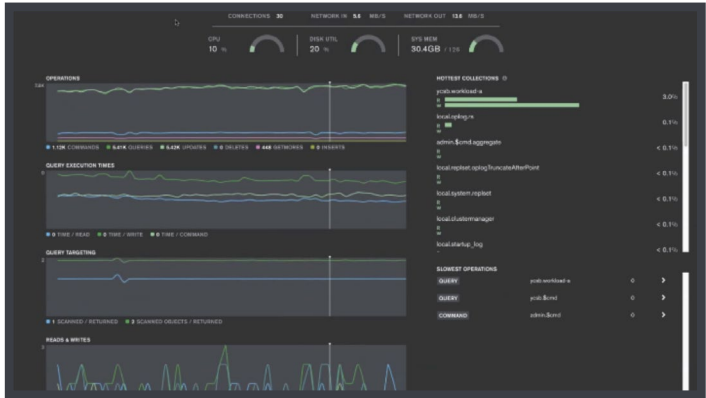
## Workload B (95-5 read heavy)

### Cloud



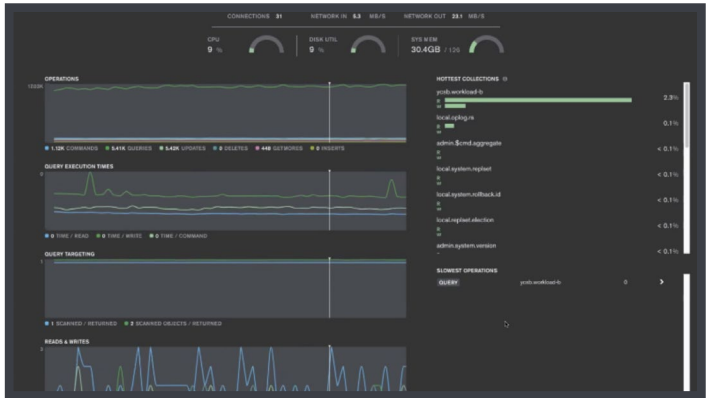
380 ops/s

### Diamanti



12,398 ops/s (33.3x improvement)

### Diamanti



12,398 op/s (32.6x improvement)

# Essential Guidance: Advice for Buyers

**Containers are the computing foundation for the next generation and enterprises are beginning to adopt containers en masse.** Kubernetes is a powerful and modern platform for containers, but it is complex and difficult to deploy. Kubernetes is also not a solution in and of itself; it requires many supporting elements to make it into a complete and operational container platform for the enterprise. Enterprises considering Kubernetes should look to remove as much complexity as possible through pre-built and commercially supported container platforms.

**Kubernetes is not only for stateless workloads.** Stateful databases, such as MongoDB, are also now supported in Kubernetes, but require persistent storage and high performance I/O systems. While containers and Kubernetes can abstract away much of the infrastructure, the underlying physical I/O systems are still important, particularly for I/O heavy workloads like MongoDB. Thus, the underlying infrastructure is still an important consideration when containerizing high performance workloads such as databases. Enterprises should also look to leverage Kubernetes operators, which will make provisioning and management of complex, stateful workloads easier.

IDC concludes that the Diamanti platform, an integrated solution designed for modern containerized applications, can offer an easy way for enterprises to deploy Kubernetes and host stateful workloads like MongoDB with high levels of performance because of its optimized storage and networking subsystem that feature I/O offload. Additionally, MongoDB Cloud Manager offers an easy to use and ready to go management platform that can offer insights into on-premises MongoDB clusters, which can also complement cloud instances with its sister counterpart, MongoDB Atlas. This provides the user with a hybrid cloud experience and allows the user to leverage both deployments types, which can be optimized for each workload's security, cost, and performance requirements.

## IDC Conclusion

IDC concludes that the Diamanti platform, an integrated solution designed for modern containerized applications, can offer an easy way for enterprises to deploy Kubernetes and host stateful workloads like MongoDB with high levels of performance because of its optimized storage and networking subsystem that feature I/O offload. Additionally, MongoDB Cloud Manager offers an easy to use and ready to go management platform that can offer insights into on-premises MongoDB clusters, which can also complement cloud instances with its sister counterpart, MongoDB Atlas. This provides the user with a hybrid cloud experience and allows the user to leverage both deployments types, which can be optimized for each workload's security, cost, and performance requirements.

# IDC Lab Validation Methodology

This Lab Validation Report provides a summary of an extensive validation process performed by IDC in collaboration with the supplier's teams. IDC relied on the supplier's equipment, facilities and their configuration to perform this validation. All of the tests were conducted during the presence of one or more IDC Analysts.

This report is meant to provide a quick set of inferences and insights for IT professionals and business decision makers seeking to perform further due diligence on the capabilities of the product and/or services that have been validated in this Report. However, the goal of this Report is not to supply detailed hands-on test plans and validation jobs. It is not meant to replace the evaluation process that most businesses will conduct before making any decision to purchase the product and/or services.

It is for this reason that this Report is not designed to be an all-inclusive document on all the capabilities of the product, but rather as a concise document that highlights select features/functions of products, their relative performance, and the value these features bring to businesses.

Finally, even though this Report is a sponsored document, it is not meant to be an IDC endorsement of the product, service or the sponsoring supplier. IDC's opinions are its own and not influenced by the production of this document.

## IDC Custom Solutions

IDC Corporate USA  
5 Speen Street  
Framingham, MA  
01701, USA  
T 508.872.8200  
F 508.935.4015  
Twitter @IDC  
idc-insights-community.com  
www.idc.com

**This publication was produced by IDC Custom Solutions.** The opinion, analysis, and research results presented herein are drawn from more detailed research and analysis independently conducted and published by IDC, unless specific vendor sponsorship is noted. IDC Custom Solutions makes IDC content available in a wide range of formats for distribution by various companies. A license to distribute IDC content does not imply endorsement of or opinion about the licensee.

External Publication of IDC Information and Data — Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.

Copyright 2018 IDC. Reproduction without written permission is completely forbidden.

Sponsored by:

