

Deploy Scalable and Highly Performant NoSQL Databases with MongoDB and Diamanti Enterprise Kubernetes® Platform

Today's applications create massive volumes of new and rapidly changing data types. Enterprise application developers and IT architects have a strong need for a database system that can provide scale, performance and agility to manage this massive amount of rapidly changing data types. Application developers and IT architects depend on a database system that can scale and adapt itself to these dynamic changes and avoid re-architecting, re-designing and re-deploying their infrastructure and applications frequently.

Contrary to the monolithic relational databases, modern NoSQL databases are designed to allow the insertion of data without a predefined schema. NoSQL databases are highly scalable and provide superior performance compared to relational databases. NoSQL databases run on commodity, off-the-shelf storage and processors, versus relational databases which require large monolithic servers and storage infrastructure.

Typically, NoSQL databases support sharding, with which they natively and automatically spread data across an arbitrary number of servers, without requiring the application to be aware of the infrastructure composition. Data and query load are automatically balanced across servers, and when a server goes down, it can be quickly

and transparently replaced with no application disruption. This enables NoSQL databases to be deployed in a geographically distributed scale-out architecture with high availability. NoSQL databases support several data store use cases such as document databases, graph stores, key-value stores and wide-column stores.

While NoSQL databases provide the flexibility, scalability and availability at the application layer, containers and Kubernetes complement the same design principles by offering near instant scalability and inherent high availability at the infrastructure layer. Together, these factors enable the deployment of stateful database systems at massive scale in a continuous delivery model.

Diamanti and MongoDB: Delivering Production Grade, Highly Scalable and Highly Performant Document Databases

MongoDB and Diamanti bring together a scalable, highly performant and production grade NoSQL database solution on bare-metal Kubernetes infrastructure. The power of containers, Kubernetes, data volume abstraction and



Diamanti's unique quality-of-service guarantees enable the solution to easily scale on demand and provide unparalleled operational uptime and efficiencies across resource allocation and usage. The Diamanti platform offers the ability to backup and replicate data for high availability and disaster recovery, and it enables seamless movement of NoSQL database deployments across a hybrid cloud environment.

Deploy Containers on a Multi-Zone Kubernetes Cluster with Diamanti in Minutes

The Diamanti Enterprise Kubernetes Platform provides enterprises with turnkey operational infrastructure using standard virtualization protocols for storage and networking alongside open-source CNI and CSI plug-ins.

Each Diamanti cluster pools low-latency, high-performance NVMe flash storage and 40GbE networking, and comes with open-source Docker and Kubernetes pre-installed. Containers can be deployed minutes after racking and stacking the Diamanti cluster, where each pod is assigned a routable IP address due to Diamanti's innovative approach to network virtualization for containers.

Implementation of multi-zone clusters on the Diamanti platform allows a Kubernetes cluster to distribute nodes across different zones, ensuring application and infrastructure availability. Diamanti simplifies multi-zone cluster configuration and management with built-in capabilities to protect applications from failures.

Running MongoDB on Kubernetes with Diamanti

MongoDB Enterprise Kubernetes Operator deploys MongoDB as a set of microservices through the Ops Manager by making Kubernetes API calls. The Operator deploys and manages standalone, replica set, and sharded cluster deployments of MongoDB.

The Operator provides full instance configurability by allowing the end user control over resources, number of nodes, and data storage per available layer (data, log, etc). The operator supports cluster-native TLS for both agents running in the pod and applications connecting to the databases. This combined with MongoDB CRD for database users, provides the foundation for deploying performant and secure databases. Most importantly, it lets the end user upgrade an instance of MongoDB from one version to the next by simply editing the existing Kubernetes object, thus taking the pain of migrating data and version migrations away from the user.

Figure 1 shows one instance of a three-way MongoDB replica set created using a small amount of resources (3 CPU cores and 4 GB of RAM) to demonstrate the ability to push a large number of operations per second through this single cluster, achieving 10,000 database operations per second. Three containers were deployed: one Primary and two Secondaries. The Secondaries replicate the data from the Primary. Note that this was using MongoDB's existing technologies and practices. The key differences are the guarantees that Diamanti can assure with respect to the storage input/output operations per second (IOPS) and network latency.

Diamanti at a Glance

SIMPLICITY

- 15-minute bare-metal deployment
- Easy to manage and scale
- Kubernetes certified
- No vendor lock-in

PERFORMANCE

- Real-world 1,000,000 IOPS per 1U
- Consistent 100-microsecond latency
- Industry-leading application-level transactions per second

EFFICIENCY

- 70% lower TCO
- 100% host utilization
- 95% usable storage capacity
- No hypervisor needed
- Guaranteed QoS with no overprovisioning

ENTERPRISE READY

- Full-stack support
- Production-grade SLAs
- Secure multi-tenant isolation
- Advanced DR/DP
- On-premises availability zones

MongoDB at a Glance

BEST WAY TO WORK WITH DATA

- Work with data in a natural, intuitive way
- Get great performance without a lot of work
- Adapt and make changes quickly
- Supports a wide variety of data and queries

INTELLIGENTLY PUT DATA WHERE YOU NEED IT

- Deliver globally resilient apps through sophisticated replication and self-healing recovery
- Grow horizontally through native sharding
- Run operational and analytical workloads in the same cluster
- Place data on specific devices and in specific geographies for governance, class of service, and low-latency access

FREEDOM TO RUN ANYWHERE

- Database that runs the same everywhere
- Leverage the benefits of a multi-cloud strategy with no lock-in
- Available as a service in 50+ regions

The Diamanti console provides detailed view of the deployment with network and storage usage statistics. Figure 2 shows the MongoDB cluster running in a stable state with a light load.

Yahoo! Cloud Serving Benchmark (YCSB) was utilized as our testing tool for benchmarking. In this case, we used Workload C, which simulates 50% read and 50% write requests. The test was setup so that all reads were sent to the Secondaries and all writes were directed to the Primary.

Using a real-world example to showcase the power of MongoDB and Diamanti, we inserted and updated more than 11 million records, which was approximately 40 GB of data, and this accounted for 30 million total database operations. We utilized the asynchronous YCSB driver in order to achieve maximum throughput, but it must be noted that using the regular driver would also perform at similar levels.

Figure 3 shows the database metrics in MongoDB Ops Manager Console for Primary and Secondaries under YCSB load test. As the number of connections increased, the Operation Execution Times consistently remained below one millisecond on the Primary which serviced all writes. The Operation Execution Times usually remained below two milliseconds on the two Secondaries which serviced the read and replication requests. The three-way MongoDB replica set was instantiated with a limited set of 3 CPU cores and 4 GB of RAM.

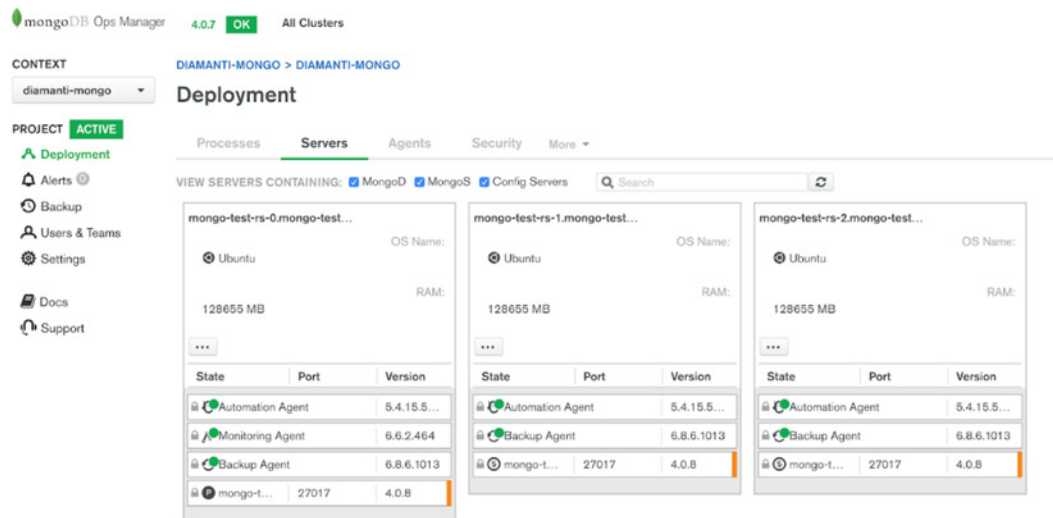


FIGURE 1: MongoDB Replica Set Running on Diamanti Platform

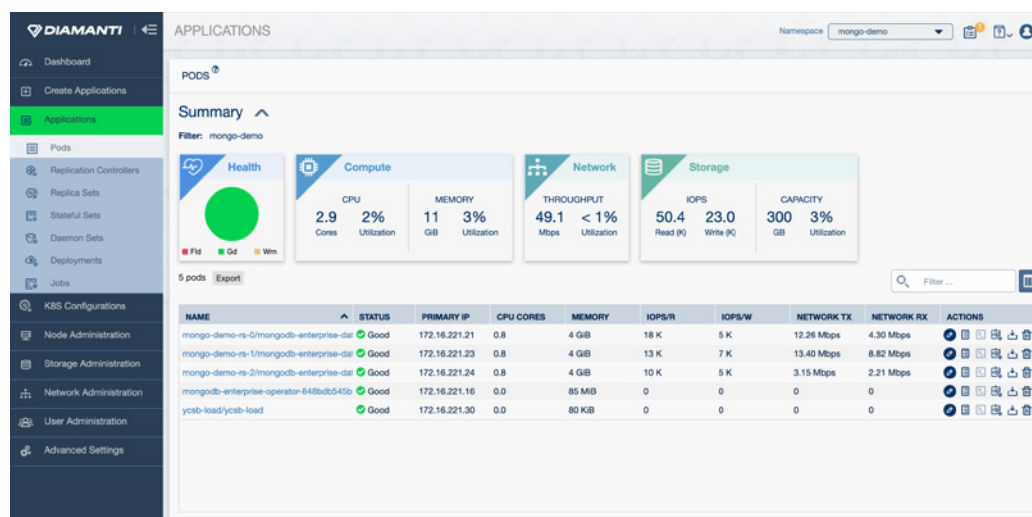


FIGURE 2: Diamanti console showing MongoDB replica set with detailed statistics

As the read/write IOPS on the MongoDB instance increased, the system was able to provide guaranteed consistent performance due to Diamanti’s unique storage and network quality-of-service guarantees. Furthermore, using the MongoDB Ops Manager, the database instance can be seamlessly expanded without having to destroy and rebuild the instance. As the footprint of the database instance grows, there is no impact to performance of the system due to the unique differentiations of the Diamanti platform.

Summary

Using the power of Kubernetes on bare-metal, users can tap into performance, ease of use and flexibility of deploying MongoDB in containers. In this use case, we showcased how an extremely low-resourced instance of MongoDB is able to produce highly sustained throughput, while also maintaining stable and low latencies. With minimal resources on the Diamanti Enterprise Kubernetes Platform, users are able to push MongoDB's performance and achieve 10,000 operations per second which can be directly translated into significant efficiency benefits and cost savings. With the power of Diamanti, users are able to harness the simplicity and elasticity of deploying MongoDB, while also leveraging guaranteed high performance.

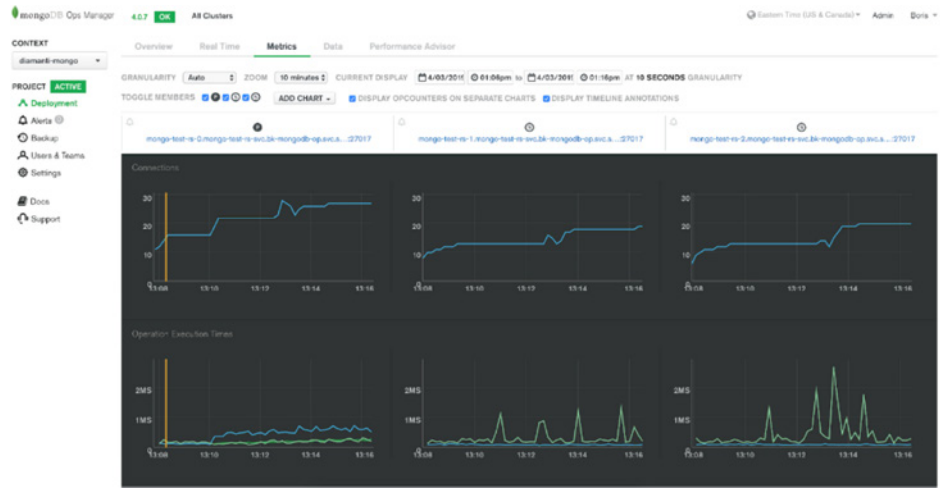


FIGURE 3: MongoDB Replica Set under YCSB Load

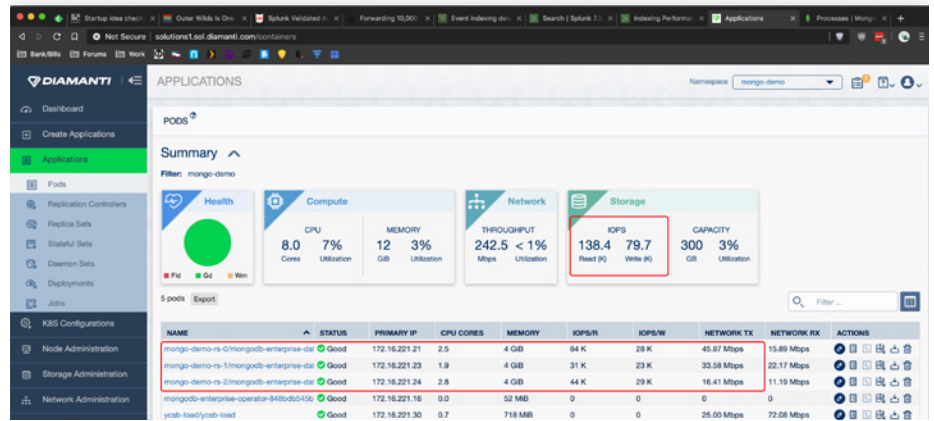


FIGURE 4: Diamanti Console with MongoDB ReplicaSet under YCSB Load

ABOUT DIAMANTI

Diamanti delivers the industry's only purpose-built, fully integrated Kubernetes platform, spanning on-premises and public cloud environments. We give infrastructure architects, IT operations, and application owners the performance, simplicity, security, and enterprise features they need to get cloud-native applications to market fast. Diamanti provides the lowest total cost of ownership to enterprise customers for their most demanding applications. Based in San Jose, California, Diamanti is backed by venture investors CRV, DFJ, Goldman Sachs, GSR Ventures, Northgate Capital, and Translink Capital.

ABOUT MONGODB

MongoDB is the leading modern, general purpose database platform, designed to unleash the power of software and data for developers and the applications they build. Headquartered in New York, with offices across North America, Europe, and Asia-Pacific, we are close to where you do business. MongoDB has more than 13,000 customers in more than 100 countries. The MongoDB database platform has been downloaded over 60 million times and there have been more than 1 million MongoDB University registrations.