HO5604

Deploying MongoDB

A Scalable, Distributed Database with SUSE® Cloud

Alejandro Bonilla
Sales Engineer
abonilla@suse.com

Agenda

- SUSE Cloud Overview
  - Getting familiar with the Cloud way
- What is MongoDB?
  - The new workloads
- Hands On

A SUSE Cloud Overview

SUSE Cloud

Enterprise OpenStack distribution that rapidly deploys and easily manages highly available, mixed hypervisor IaaS Clouds

- Increase business agility
- Economically scale IT capabilities
- Easily deliver future innovations

Private Cloud Responsibilities

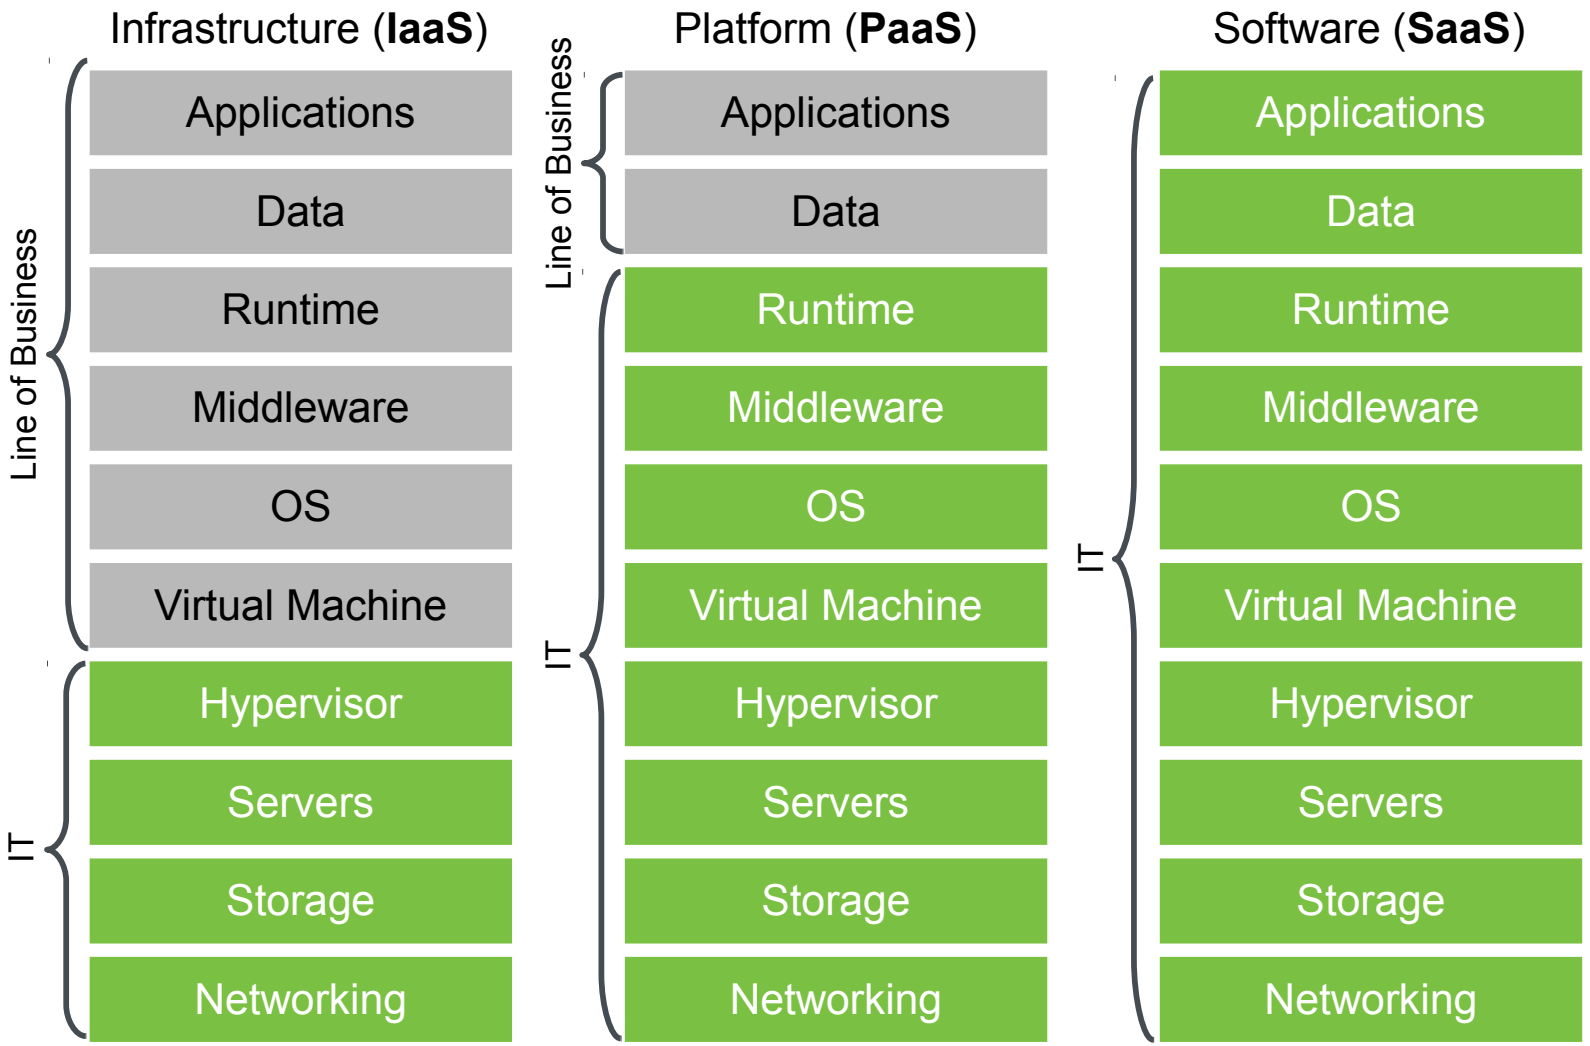

Platform for Your Data Center Evolution

Enterprise OpenStack Distribution

- Leading open source cloud project delivering fast innovation of advanced IaaS cloud services

Integration with Ceph Distributed Storage

- Reduce costs with a single software-defined storage solution for massively scalable and reliable block, object, and image storage

Driver of Organizational Innovation

- Develop and deploy new applications that take advantage of SUSE Cloud capabilities

Award-Winning Worldwide Support

- Backed by the excellence of SUSE engineering and the only organization with a 20-year history of supporting open source software

Fast Installation and Simplified Management

SUSE Cloud Administration Server

- Faster ROI through faster installation and easier management of OpenStack Cloud

Highly Available Cloud Services

- Maintain business agility and deliver enterprise-grade SLAs through continuous availability of cloud services

Standardized Product Life Cycle

- Packaged product integrated with SUSE update and maintenance processes to ensure simplified enterprise maintenance

SUSE® Cloud 4

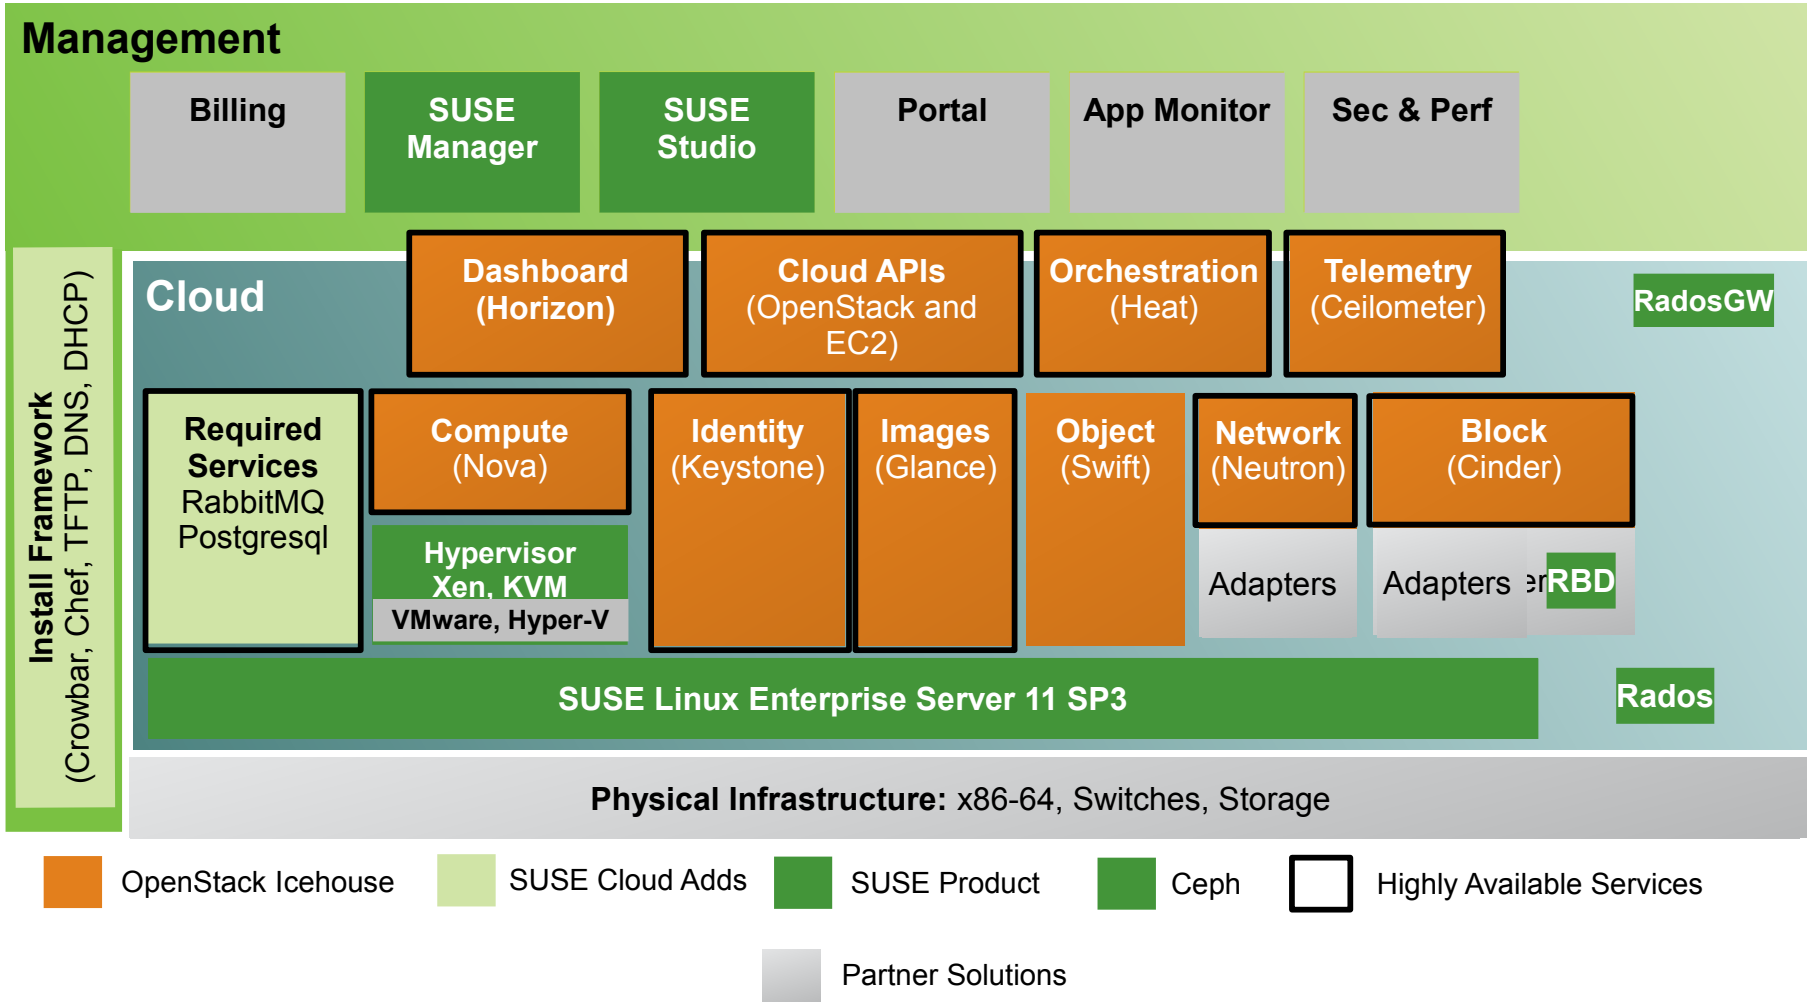
Management
Billing
SUSE Manager
SUSE Studio
Portal
App Monitor
Sec & Perf
Install Framework
(Crowbar, Chef, TFTP, DNS, DHCP)
Cloud
Dashboard (Horizon)
Cloud APIs (OpenStack and EC2)
Orchestration (Heat)
Telemetry (Ceilometer)
RadosGW
Required Services
RabbitMQ
Postgresql
Compute (Nova)
Hypervisor
Xen, KVM
VMware, Hyper-V
Identity (Keystone)
Images (Glance)
Object (Swift)
Network (Neutron)
Block (Cinder)
Adapters
Adapters
RBD
SUSE Linux Enterprise Server 11 SP3
Rados
Physical Infrastructure: x86-64, Switches, Storage
OpenStack Icehouse
SUSE Cloud Adds
SUSE Product
Ceph
Highly Available Services
Partner Solutions

What is MongoDB?

MongoDB

The document-oriented NoSQL database

General Purpose

Document Database

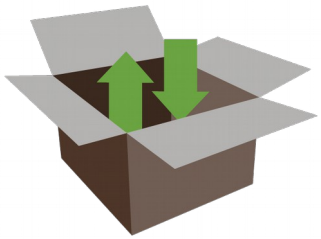

Open-Source

mongoDB

MongoDB Overview

Agile

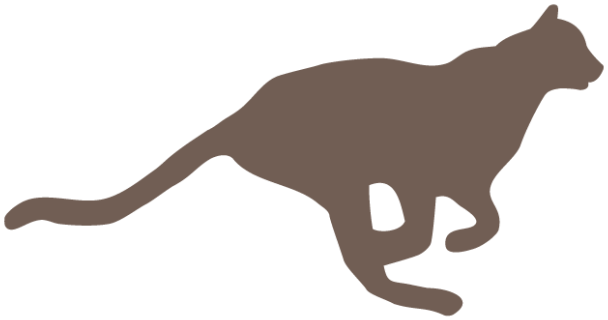

Scalable

mongoDB

MongoDB Features

- JSON Document Model with Dynamic Schemas
- Auto-Sharding for Horizontal Scalability
- Text Search
- Aggregation Framework and MapReduce
- Full, Flexible Index Support and Rich Queries
- Built-In Replication for High Availability
- Advanced Security
- Large Media Storage with GridFS

mongoDB

Drivers and Connectivity

Drivers

Drivers for most popular programming languages and frameworks


Java


Ruby

JavaScript


Perl


Python


Haskell

Shell

Command-line shell for interacting directly with database

```
> db.collection.insert({product:"MongoDB",
type:"Document Database"})
>
> db.collection.findOne()
{
        "_id"         : ObjectId("5106c1c2fc629bfe52792e86"),
        "product"     : "MongoDB"
        "type"        : "Document Database"
}
```

mongoDB

Document Oriented

Relational

MongoDB

```
{
  first_name: ‘Paul’,
  surname: ‘Miller’,
  city: ‘London’,
  location: [45.123,47.232],
  cars: [
    { model: ‘Bentley’,
      year: 1973,
      value: 100000, … },
    { model: ‘Rolls Royce’,
      year: 1965,
      value: 330000, … }
  ]
}
```

mongoDB

Flexible, Powerful Querying

Rich Queries	• Find Paul's cars • Find cars in London built 1970 - 1980
Geospatial	• Find all of the car owners within 5km of Trafalgar Sq.
Text Search	• Find all the cars described as having leather seats
Aggregation	• Calculate the average value of Paul's car collection
Map Reduce	• What is the ownership pattern of colors by geography over time? (is purple trending up in China?)

MongoDB

```
{
  first_name: 'Paul',
  surname: 'Miller',
  city: 'London',
  location: [45.123,47.232],
  cars: [
    { model: 'Bentley',
      year: 1973,
      value: 100000, … },
    { model: 'Rolls Royce',
      year: 1965,
      value: 330000, … }
  ]
}
```

mongoDB

Document Model Benefits

- Expressive, Flexible, Simplified Data Modeling
  - A single document can express and encompass a wide variety of notions
  - No need to migrate for simple extensions
  - Fewer collections as most data can be encapsulated in a single document
- Easier Development
  - Developers understand documents as it maps well to their data structures
- Faster Time to Market
  - Agile development means faster results

mongoDB

Performance

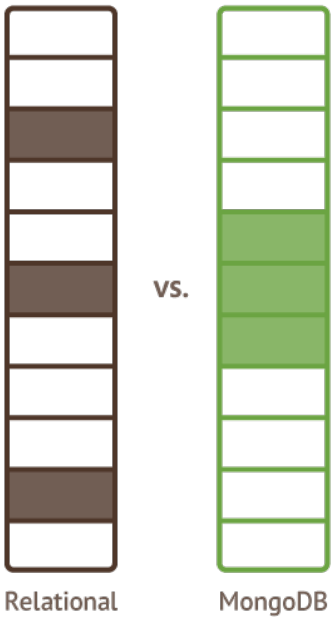

Better Data Locality

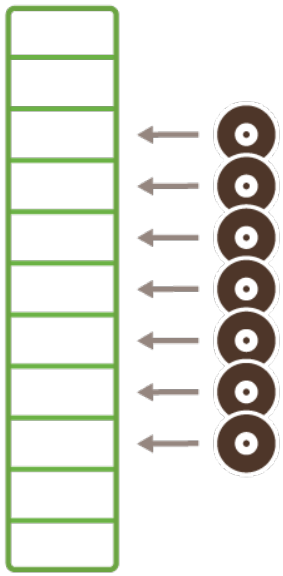

In-Memory Caching

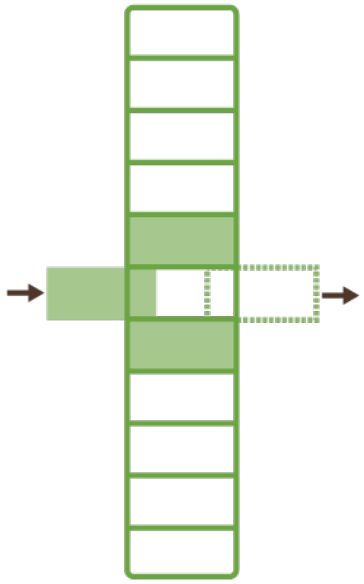

In-Place Updates

mongoDB

Scaling MongoDB

- Replica Sets
  - Redundancy, failover, high availability
- Sharding
  - Auto-paritions data, read/write scalability
- Multi-datacenter deployments
- Tunable durability, consistency
- Engineered for zero downtime

mongoDB

High Availability

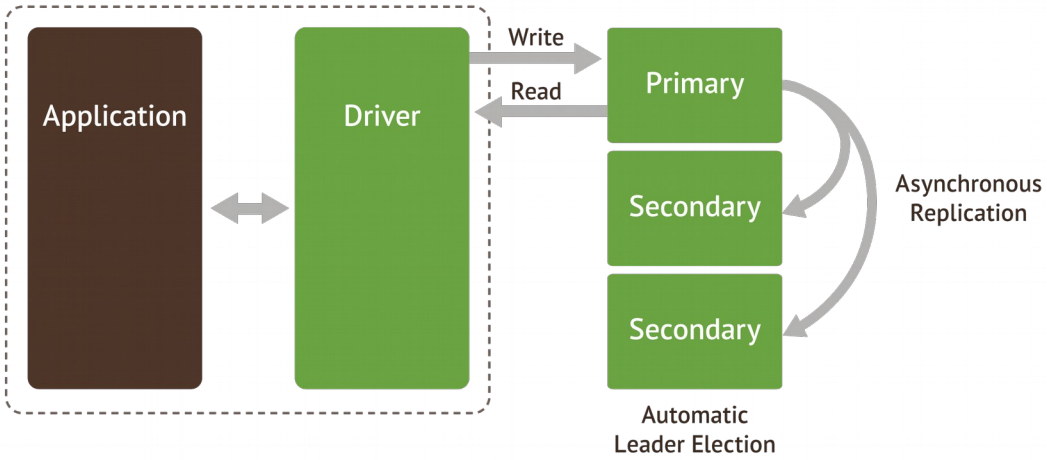

- Automated replication and failover
- Multi-data center support
- Improved operational simplicity (e.g., HW swaps)
- Data durability and consistency

mongoDB

Scalability

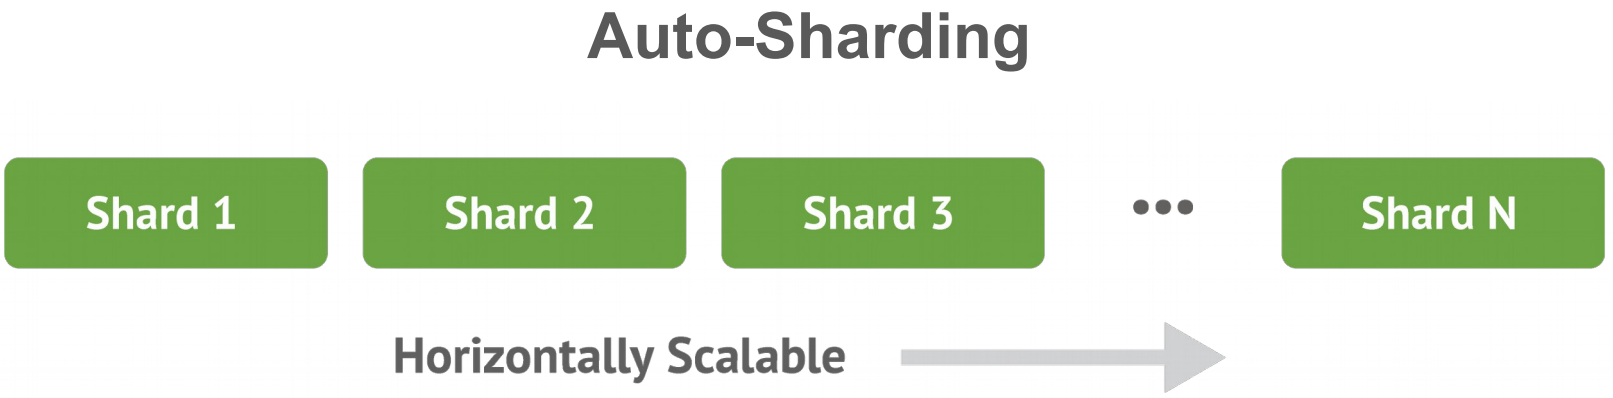

- Increase capacity as you go
- Commodity and cloud architectures
- Improved operational simplicity and cost visibility

mongoDB

Sharding and Replication

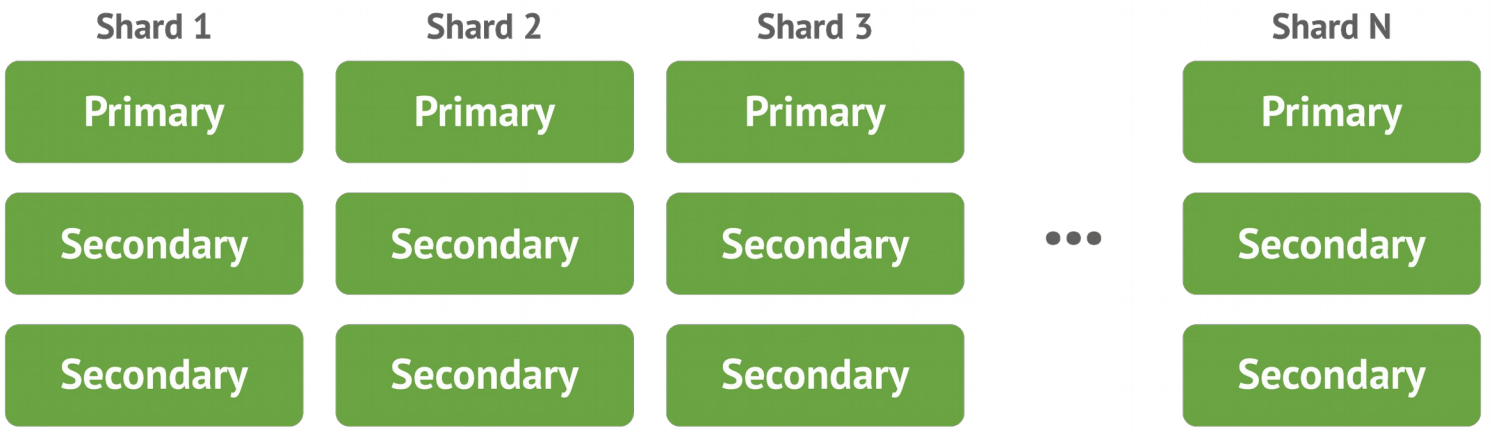

mongoDB

MongoDB Architecture

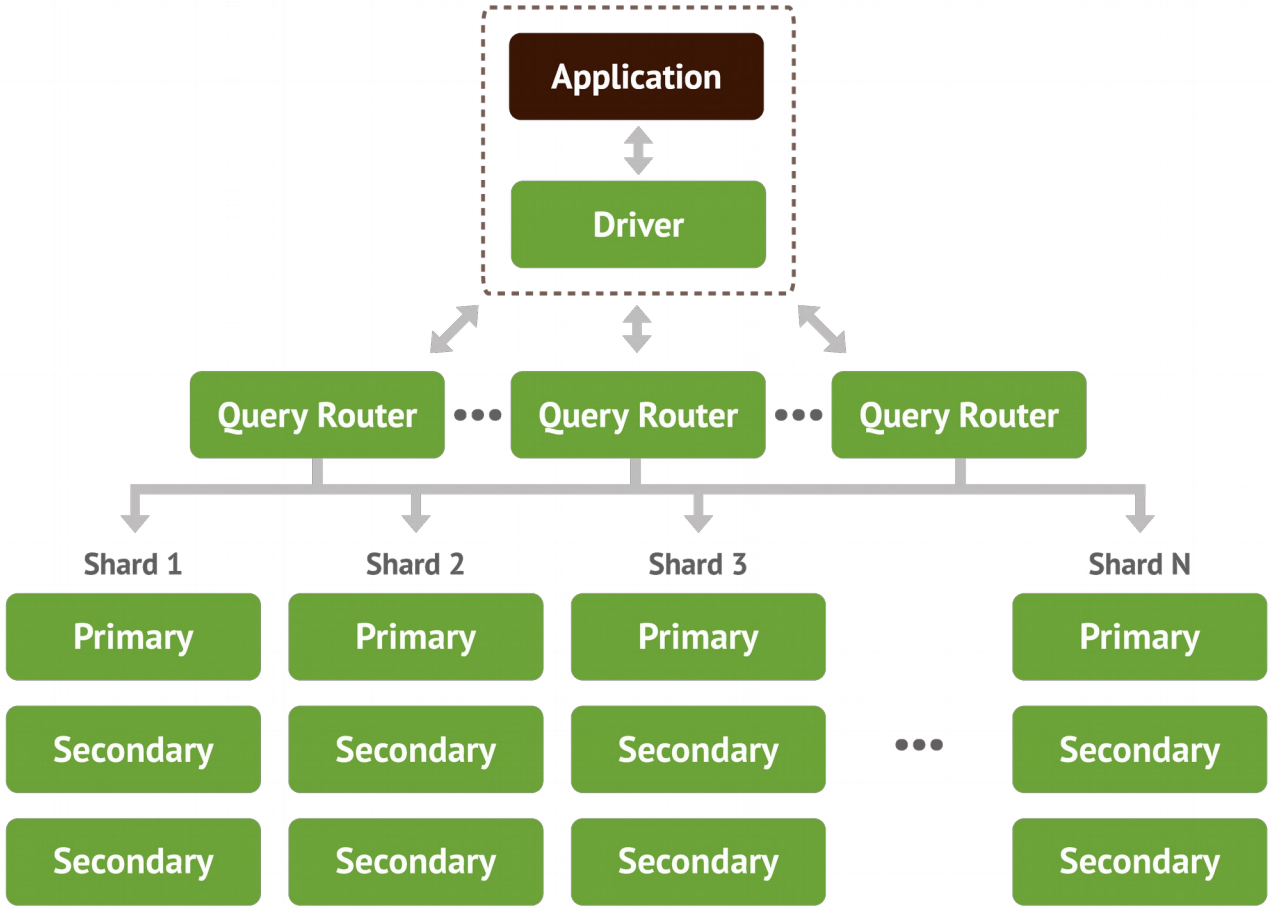

mongoDB

Hands On!

Thank you.

SUSE®
We adapt. You succeed.

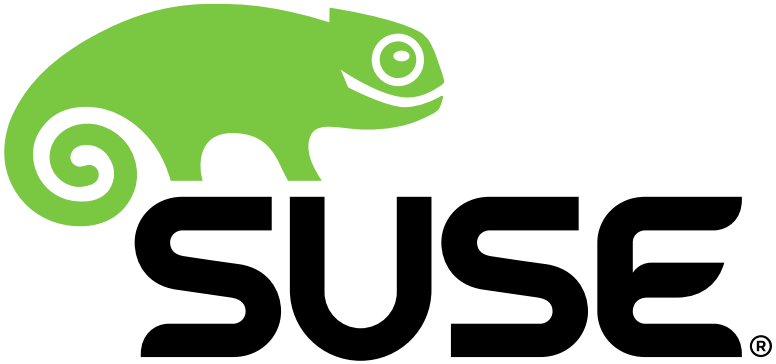
SUSE®
We adapt. You succeed.