

# Large Scale OLTP

Databases are used in various domains and have more rich features than when first invented. But most serious scenarios still require a robust Online Transaction Processing (OLTP) capability and, with the evolution of digitalization, the database must be able to scale to support the business growth.

## Key Benefits

A large scale OLTP system meets an array of requirements

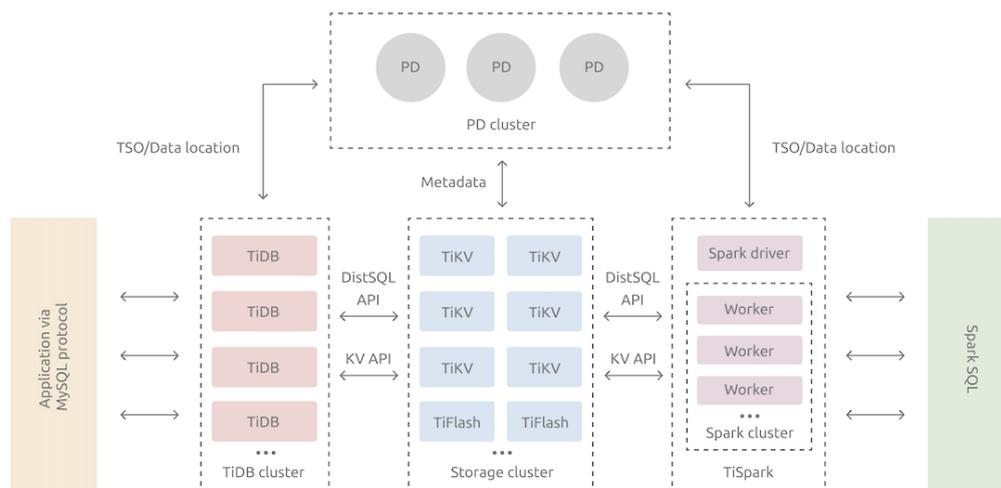
- Scalability enables the environment to adapt to changing needs
- Availability keeps the data always available and accessible
- Low latency means that users are able to access the information in a timely manner
- Throughput delivers the write and read capabilities needed to meet your business requirements

## Example Architecture

Designed as a distributed database, TiDB is composed of 3 main components:

- TiDB Server
- Placement Driver (PD) Server
- Storage Servers

Currently, there are 2 types of storage servers: TiKV and TiFlash. The compute and storage components are separated and you can scale both without worrying about downtime or causing issues to other components.



## TiDB Server

TiDB Server is a stateless SQL layer that accepts the connection. The TiDB server receives SQL requests, performs SQL parsing and optimization, generates an execution plan, then executes it.

## Placement Driver (PD) Server

The PD server is the brain of the entire TiDB cluster. The main functions are:

- Storing the topology structure of the entire TiDB cluster
- Storing the metadata of storage servers
- Providing the TiDB Dashboard management UI
- Allocating transaction IDs to distributed transactions
- Managing the Timestamp Oracle (TSO)

## Storage Servers

### TiKV Server

TiKV is a distributed transactional key-value storage engine. [Region](#) is the basic unit to store data. All data in TiKV is automatically maintained in multiple replicas (three replicas by default).

### TiFlash Server

The TiFlash Server synchronizes data from TiKV. TiFlash stores data by column, and is mainly designed to accelerate analytical processing.

## Reference Architecture

Depending on the scale and availability requirements, you can run your cluster over one availability zone (data center) as in figure 1, or multiple availability zones as in figure 2. In figure 3, two clusters are deployed separately in different regions to support better latency and running in bi-directional binlog synchronization to keep everything in sync.

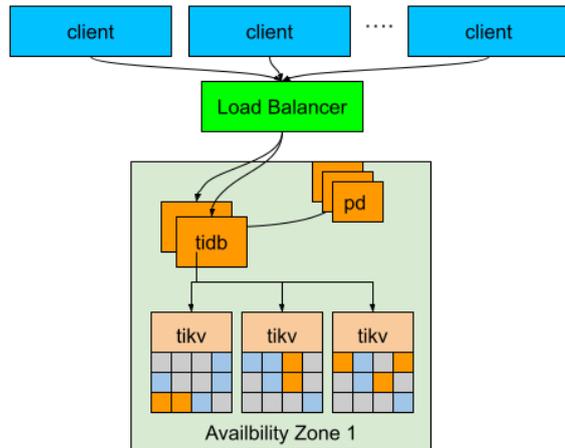


Fig.1 Large scale OLTP reference architecture - Master Cluster in single AZ

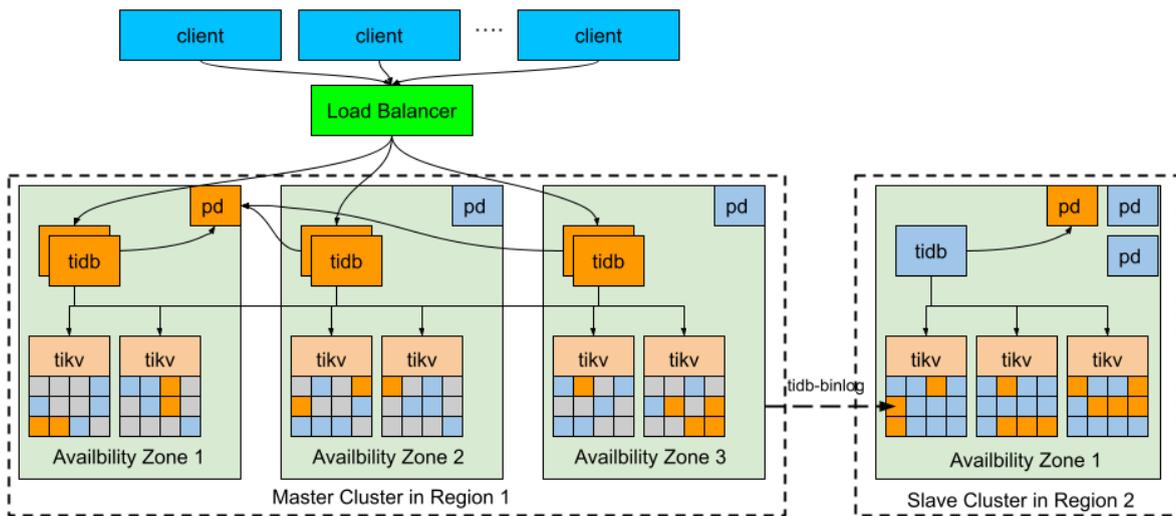


Fig.2 Large scale OLTP reference architecture - Master Cluster cross 3AZ

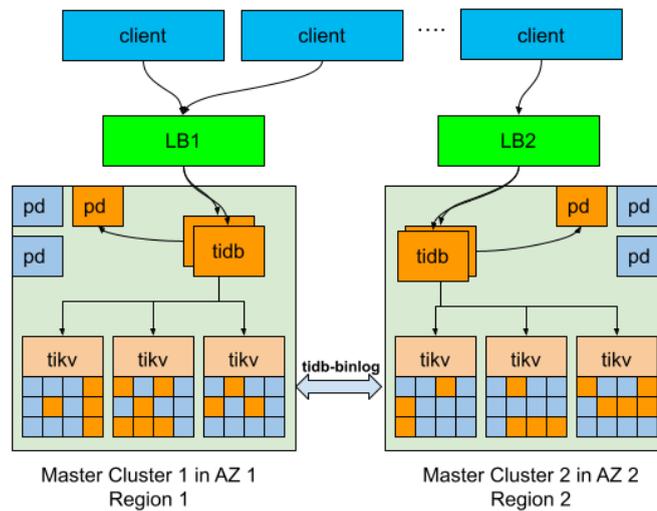


Fig.3 Large scale OLTP reference architecture - Master-Master Cluster in different Region via bi-directional synchronization

## Example Applications

The need for massive OLTP exists in many industries and for a variety of reasons. Payment processing is one industry where the ability to scale to meet user needs is paramount. Delivering a fully ACID-compliant environment is important in this area so that users are assured of consistent query results. Buyers can transfer money through cell phones to sellers via mobile apps. Users can submit payments to credit cards or other companies. The business can be growing at a speed unheard of in the past. There is a need for a database that can easily scale to support such hyper-growth. The design of TiDB, with compute and storage separation, is crucial to support scalability.

Users have an expectation of availability. Whether it is the desire of customers to get in touch with you through any channel at any time or simply the belief that a company's application must be available at all times, maintaining your application's availability is important. You miss great opportunities if your system cannot maintain availability. TiDB supports deployment across different machines, racks, or data centers. Using a consensus protocol to synchronize data between leader and follower nodes, your data replicas are distributed across different nodes to maintain high availability. Together, this works to ensure continuous availability to your user base.

It is also key that you provide a smooth and responsive experience to customers during the interaction. With the increasing complexity of modern programming languages and frameworks, overall performance benefits from the low latency of the database engine. With a distributed SQL optimizer, your SQL statement gets the best response time without concern for which storage engine is responding.

Modern architecture design usually involves a consideration of handling high throughput during peak time. TiDB makes it easy since you can get higher Queries Per Second (QPS) by adding more nodes to the stateless TiDB SQL layer, or larger storage by simply adding more nodes to the TiKV layer. These actions can be taken independently, so your environment is consistently right-sized for the current situation.

## Customer Examples

### Zalopay

ZaloPay is a mobile payment application based in Vietnam, serving user's daily life and business needs. The service launched in 2017. Merchants are looking for quick settlements and value-added services and consumers are looking for a better experience. Meanwhile, new regulations for data privacy, data governance, and payment response time are adding volatility to the already uncertain industry. To stay competitive in the payment industry, Zalopay needed to update its data architecture, especially its database. The key reasons for moving to TiDB include:

- The flexibility to run on cost-effective commodity hardware or directly in the cloud
- A strong performance and availability
- The capability to provide new customer insights that can help businesses develop differentiating services

### GAEA

GAEA is a mobile gaming provider in China and aims to develop high-quality games for international players. GAEA uses its GaeaAD system to support the cross-platform real-time advertising system. GaeaAD performs a real-time match between the advertising data and the information reported by the game SDK. In other words, GaeaAD conducts a real-time analysis based on the data of the advertisements on different advertising channels and the number of players brought by the corresponding channels, with the purpose of displaying and optimizing the conversion effects of advertising within minutes.

GAEA was encountering some limitations of its legacy MySQL application. Queries for a match had increased to over 2 minutes from about 10 seconds. They were also reaching data limits on the number of rows in some tables, requiring them to delete data to maintain performance. After implementing TiDB, they found acceptable query response times and are now able to store all of the historical data they want, providing enhanced matching based on customer data.

### Shopee

Shopee is the leading e-commerce platform in Southeast Asia and Taiwan. It provides users with an easy, secure, and fast online shopping experience through strong payment and logistical support. To continuously improve the customer experience, both platform and

sellers want to understand how consumers make their decisions by analyzing the user journey, not only from their usage pattern on the site, but also from customer support, comments, and social media posts. They also need to stay 7X24 ready for any burst of trending or unplanned social market events. A well-managed real-time inventory helps them to better manage their supply chain and P&L.

## Additional Resources

[Easily Build Your Mission-Critical Applications at Any Scale](#)

[Deep Dive into TiKV Transactions: The Life Story of a TiKV Prewrite Request](#)

[TiFS, a TiKV-Based Partition Tolerant, Strictly Consistent File System](#)