

LizardFS

DISTRIBUTED, PARALLEL, FAULT-TOLERANT, HIGHLY AVAILABLE,
GEO-REDUNDANT FILE SYSTEM

Customer:

Tinkoff



TINKOFF

Solution:

*LizardFS with
Greenplum Database,
2PiB of data*



Benefits:

*Scalability,
performance,
open-source, ease of
use, SDS*

Overview

Tinkoff is an online financial ecosystem built around client's needs. It provides a full range of financial services for individuals and businesses. Tinkoff pays special attention to the development of lifestyle banking: the ecosystem provides customers the opportunity to analyze and plan personal spending, invest savings, receive bonuses from loyalty programs, book travel, buy movie tickets, book restaurant tables, and much more.

In all communications, Tinkoff actively uses artificial intelligence and machine learning technologies; more than 40% of customer calls in chats are processed without the participation of bank employees. All products and most of Tinkoff's internal IT systems are developed by the company itself. 80% of the headquarters staff are IT professionals.

Business Challenge

Tinkoff is a data-driven company and has a heavily used and fast-growing DWH platform. The core of Tinkoff's Data Warehousing is the Greenplum Database, which processes and stores several hundreds of terabytes of data each day.

It has developed a system called Dual ETL to backup and restore data to the Greenplum Disaster Recovery cluster. One of the main parts of that system is data storage. Initially, it was an NFS storage which once stopped to fulfill our growing speed requirements.

Tinkoff needed scalable, fast, efficient and easy managed open source SDS to meet our growing requirements.



+48 537 598 137

contact@lizardfs.com

www.lizardfs.com

LizardFS

DISTRIBUTED, PARALLEL, FAULT-TOLERANT, HIGHLY AVAILABLE,
GEO-REDUNDANT FILE SYSTEM

Solution

LizardFS became a fast and easy-managed SDS for day-by-day usage. LizardFS installation is now being used for synchronization of two Greenplum clusters.

Fail tolerance is also one of the key features which Tinkoff likes in LizardFS – even if the whole DC is lost – they still have access to all their data.

Currently Tinkoff's LizardFS cluster resides in two data centers, 8 machines in each. The data amount is now more than 2 PiB, every day they write and read several hundreds of TiB of data in more than 5000 data streams.

LizardFS proved itself to scale very well horizontally – Tinkoff has already had 2 expansions and plans to expand the cluster several more times.

“We managed to create a good and working solution. When we see that we need more speed – we just add several more servers and run further!”

Dmitriy Nemchin
Data Engine Team Lead
Tinkoff

“We are proud to make the data storage better for enterprises from all over the world. Individual approach is the only approach we have to our customers!”

Marcin Konarski
Chief Technology Officer
LizardFS

Contact us: contact@lizardfs.com

Subscribe us on social media



+48 537 598 137

contact@lizardfs.com

www.lizardfs.com