# LIZARDFS

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

## BUILD YOUR STORAGE FROM COMMODITY HARDWARE, USE YOUR CURRENT INFRASTRUCTURE

SCALE UP TO 12 EXABYTES

INCREASE YOUR CAPACITY 4 TIMES

ELIMINATE STORAGE RELATED DOWNTIME

REDUCE YOUR STORAGE COST BY 30% - 70%

REDUCE TIME SPENT ON STORAGE MAINTENANCE BY 90%

KEEP YOUR DATA WHERE IT IS NEEDED ... CLOSE TO YOUR PROCESSING POWER

OFFICIALLY PART OF :

ubuntu fedora debian

### FINANCIAL

Big Data Elastic search
SAS Analytics

### DATA CENTER

Backup Big Data
Cloud drive Virtualization

### MEDIA

Archive Animation Render Farm Storage
Special Effects (SFX) Visual Effects (VFX)
Reconstruction Postproduction

### TELECOMS

Billing Big Data
SAS Analytics CCTV Recordings

### SCIENCE

Satellite Images Seismologic Data
Microscope Images Genome Sequencing

All data is distributed among multiple chunk servers.

Erasure Coding allows for parallel writes to multiple chunk servers for increased performance.

Chunk servers are build on standard commodity platforms and layered on top of any POSIX filesystem like ZFS so all the performance tuning options can be used.

Lizardfs can be containerized, which means you can run chunk servers within a container like LXC, openVZ or LXD. That allows for servers that have low I/O requirements to be utilized as chunk servers at the same time.

Automatic handling of tape drives.

Tiering between groups of different chunk servers.

Decreases RTO and RPO by securing your metadata on multiple layers via HA setup and multiple metadata loggers.

Performance depends on how you build your chunk servers so if you want extra high performance, you just improve your chunk servers or use a small group of chunk servers as your high performance tier.

LizardFS is 100% Hardware agnostic.
You can run it on any hardware you want as long as it can be managed by a Linux or Unix system.

## PLUG AND SCALE

scale vertically as well as horizontally by simply adding or removing a single drive or node

## EASY TO INSTALL AND MAINTAIN

get your cluster up and running within 2 hours (current record 28 minutes)

## GEO REPLICATION

allows you to create geo stretched clusters stretching virtually any distance

Native Clients for Linux, MacOSX, FreeBSD and Windows. NFS support (pNFS).
Hadoop connector.