



<http://skytechnology.pl>

LIZARDFS

Distributed Parallel Geo Redundant File System



LIZARDFS

After FOSDEM, how do you manage your t-shirt overflow ?

- Get another closet module from IKEA ?
- Reuse that closet you still have in the basement ?
- Replace your closet with a new one ?
- Replace all your old t-shirts with FOSDEM ones ?

LIZARDFS

And storage ?

How do you cope with sudden storage requirement overflow ?

- Get a new shiny storage system ?
- Spend a lot of time extending your XXX solution ?
- Quit your job and emigrate to south america ?
- Use expensive and slow cloud storage ?

LIZARDFS

Think LizardFS ...

- Just add chunkservers ...
- Can be anything if used with replication goals
 - Use your old spare or spare systems as chunkservers
 - You can even use them temporary for a start and move to faster systems later
- Use different chunkserver groups for different performance requirements ...
 - Old stuff for archives
 - Hyperfast new rocketspeed for performance tier

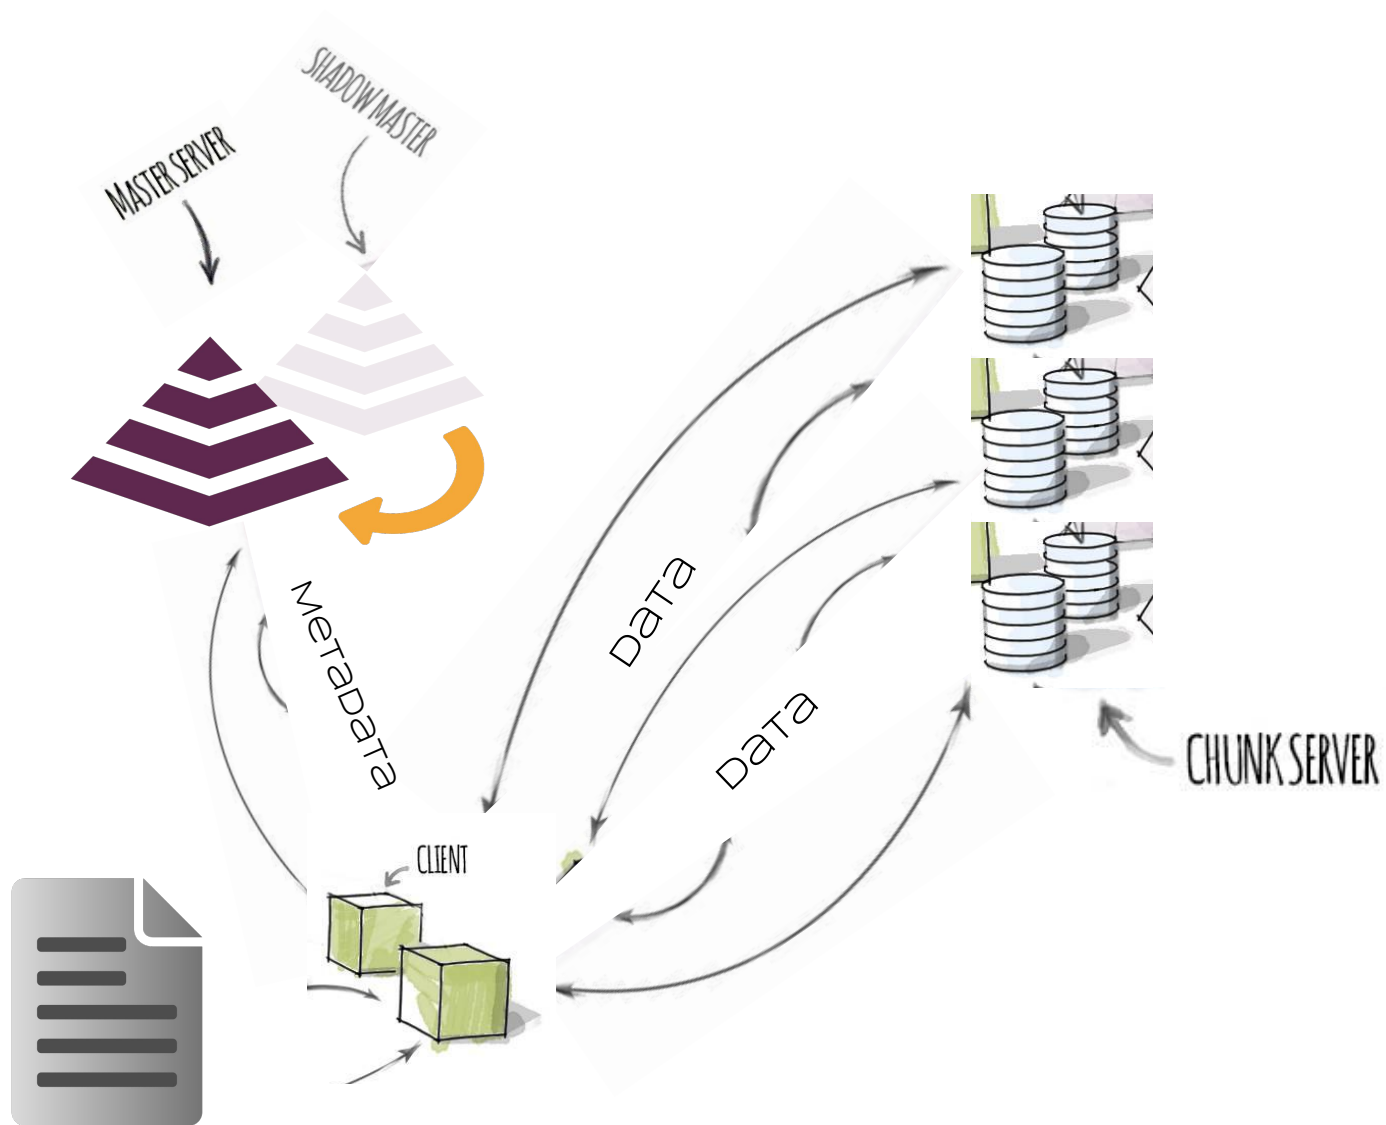
LIZARDFS

And more ...

- Simple to setup
- Simple to use
- Simple to manage
- Simple to grow .. or shrink ...
- Sooooo ... think legolike storage.
- There are advanced features too



LizardFS Architecture






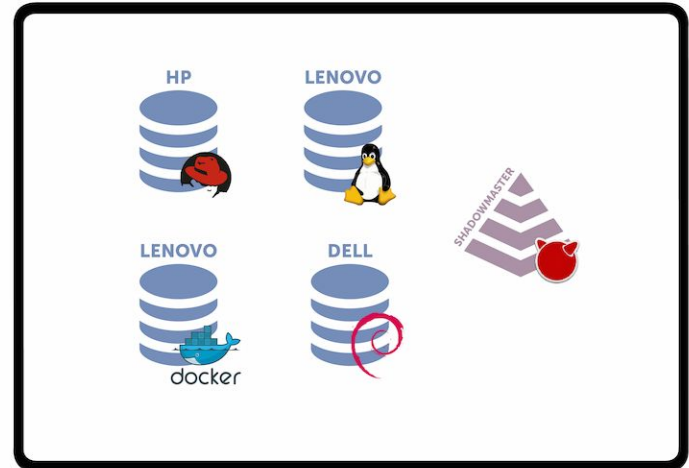
INSTALLATION TIME ~2H / RECORD 28 min




DATACENTER 1




DC 2

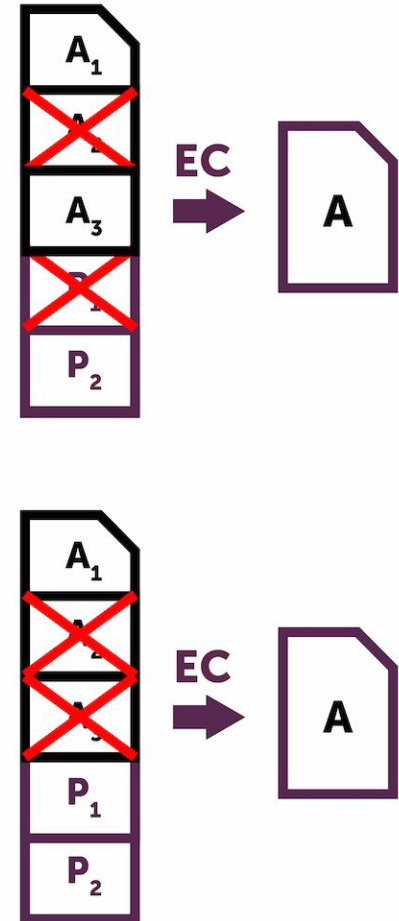
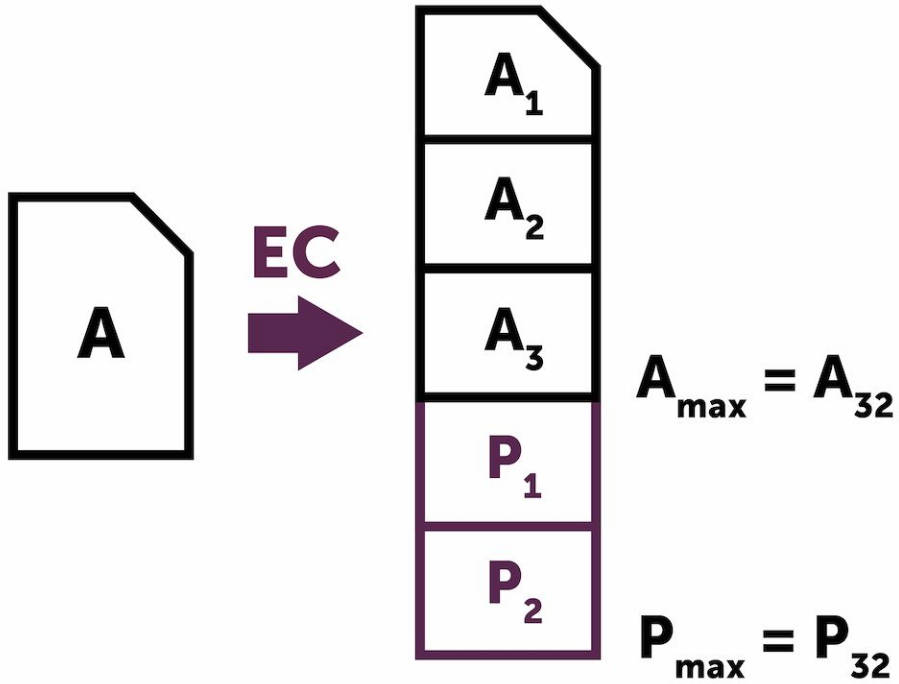


DC 3






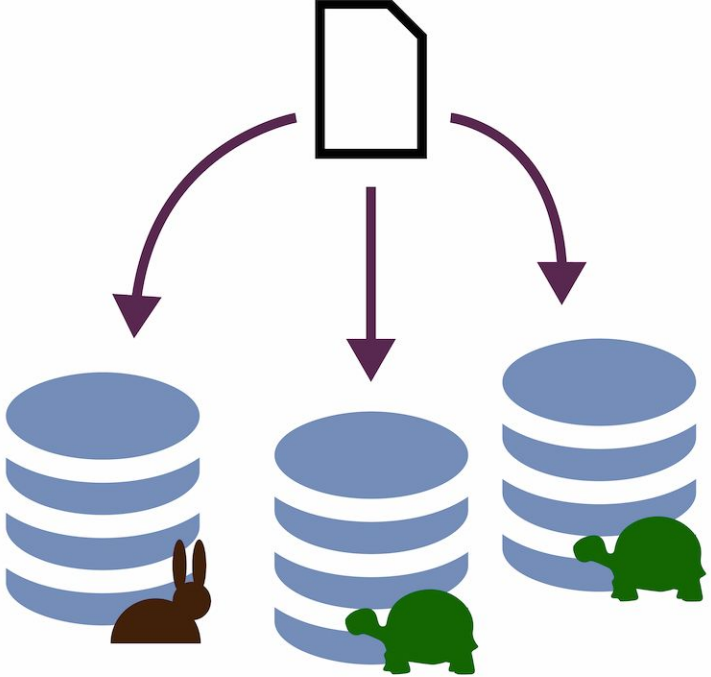
ERASURE CODING



PERFORMANCE





TIERING

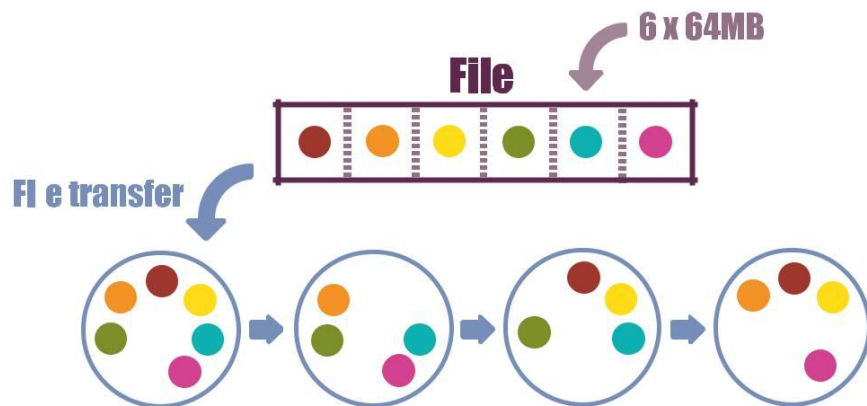


LIZARDFS

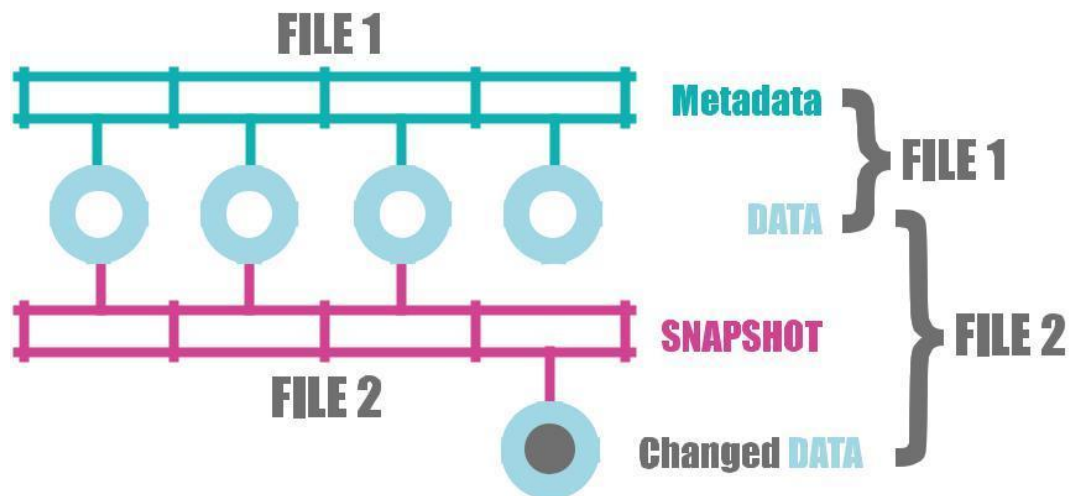
High availability



Data replication



Snapshots



Autobalancing



Portability



DEMO

Enough theory ...



Q&A

Do you have any questions ?

Regarding how things work ?

Regarding implementation ?

Why we did it ?

Other ?

For more Answers just grab one of us during
the session breaks ...

Thank you

Michal Bielicki

m.bielicki@lizardfs.com

<http://www.lizardfs.org/>

