

Safe File Storage and Databases

From Research To Transfer: User-Controllable Cloud Storage

Josef Spillner

<mailto:josef.spillner@tu-dresden.de>
<xmpp:josef.spillner@jabber.org>

GÉANT3plus Datacenter IaaS Workshop
September 11-12, 2014, Helsinki, Suomi

Everybody claims big data technology

but do we understand
small data?

(didactic data: e.g. 1 byte)

Why a single byte may be worth high-quality storage:

Cool Imaginary Mobile App: Track Your Weight

```
unsigned char weight_kg = 85; // binary: 01010101
```

How to store 1 byte of data:

disk



1 byte store
0/0 bytes transfer

disk array




2 bytes store
(1.5 bytes store)
0/0 bytes transfer

cloud



1 byte store
1/1 bytes transfer

secure cloud array



2 bytes store
(1.5 bytes store)
2/1 bytes transfer
(1.5/1 bytes transfer)



<http://lab.nubisave.org/>

Short history:

- 2010 ideas and service modelling
- 2011 **NubiSave** storage controller started
UCC paper on Cloud storage controllers
- 2012 Usable software, Debian packaging
- 2013 **NubiVis**; π -Box personal Cloud VM
FGCS article on optimal storage
SUCRE summer school tutorial at KIT
- 2014 **NubiGate** VM, **StealthDB** database
+ a couple of interested companies

research

prototypical
software
development

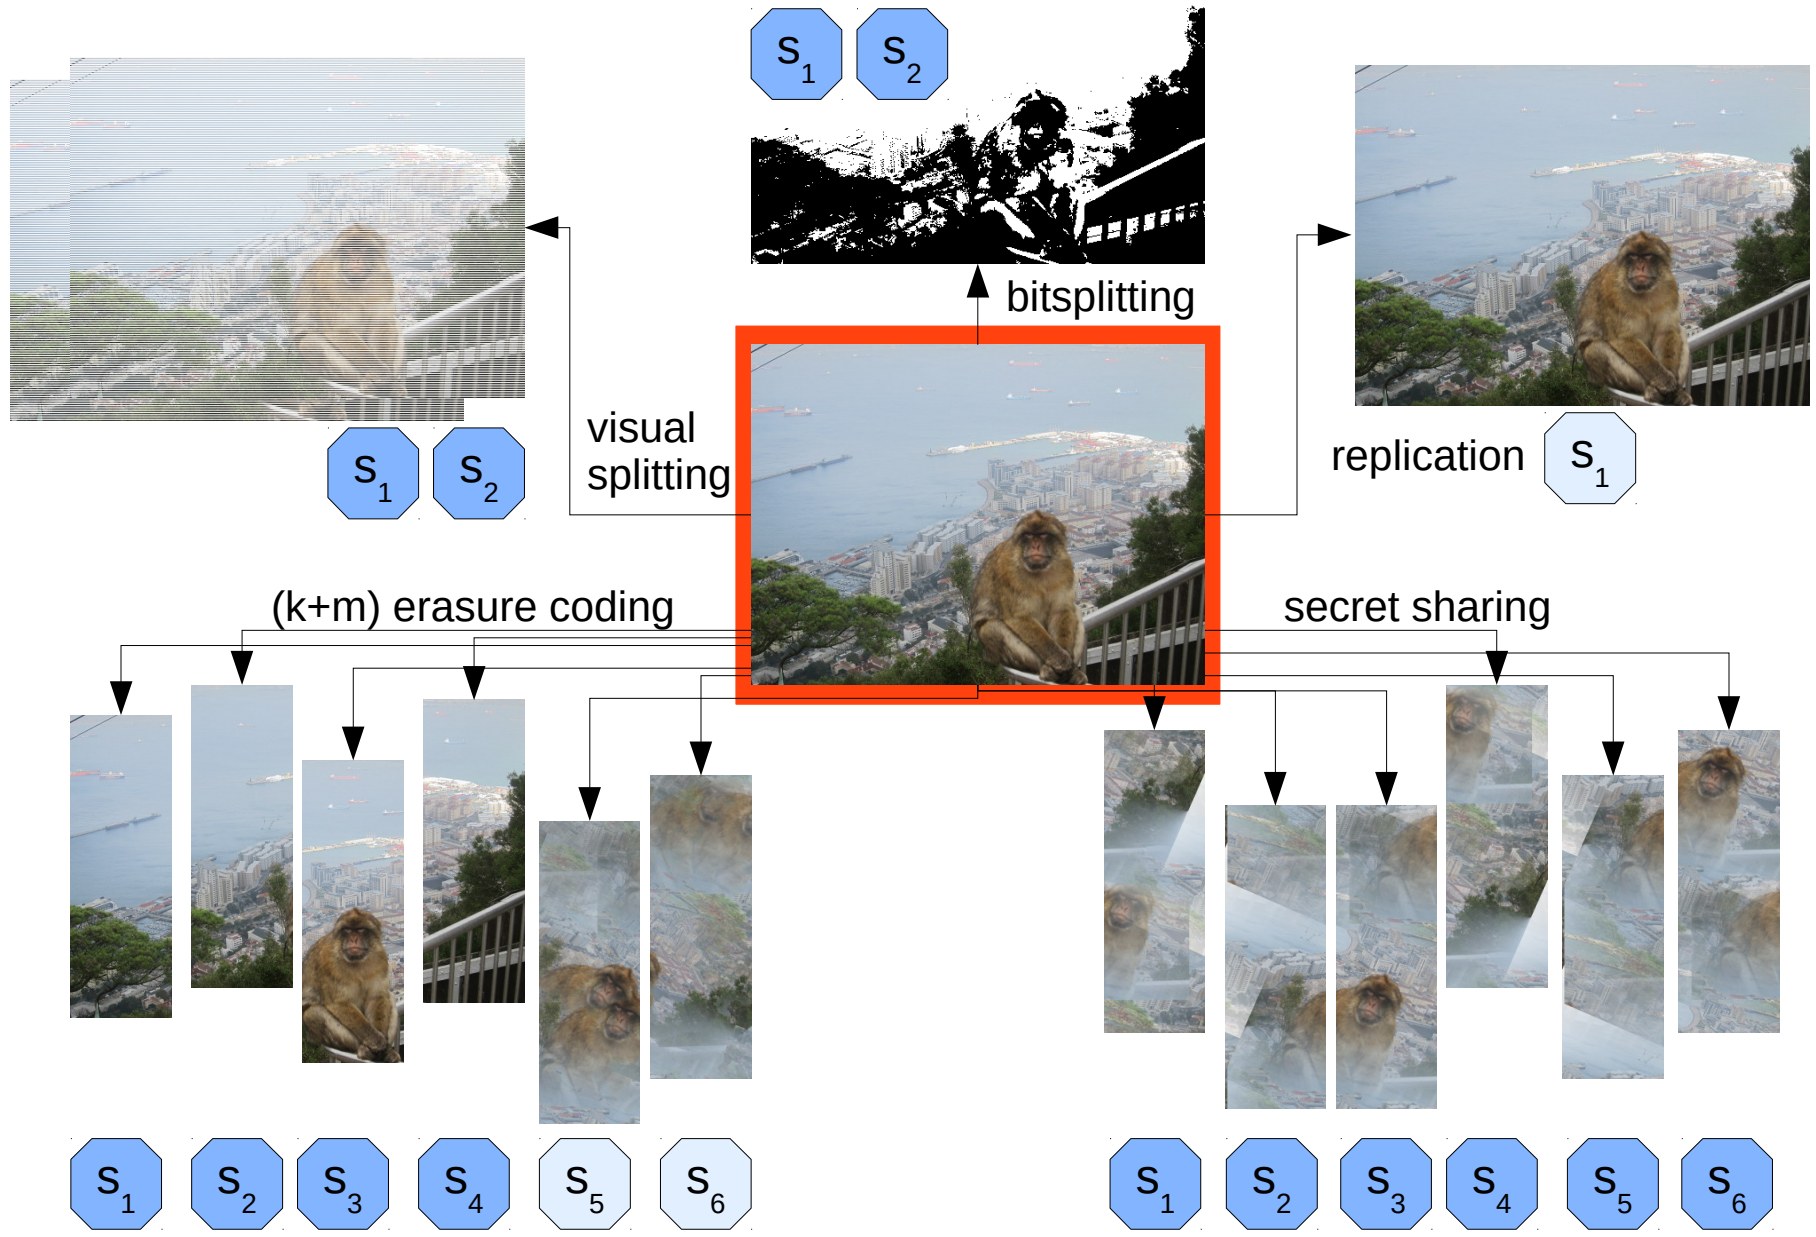
experiments
and pilot
installations

transfer

widespread
production
deployments

Research

Storage Service Combinations/Splitting



Base Ontology (WSML)

concept MeasurementUnit
conversionFactor impliesType **_double**

concept SpaceUnit subConceptOf
MeasurementUnit

instance GB memberOf SpaceUnit
conversionFactor hasValue 1024.0

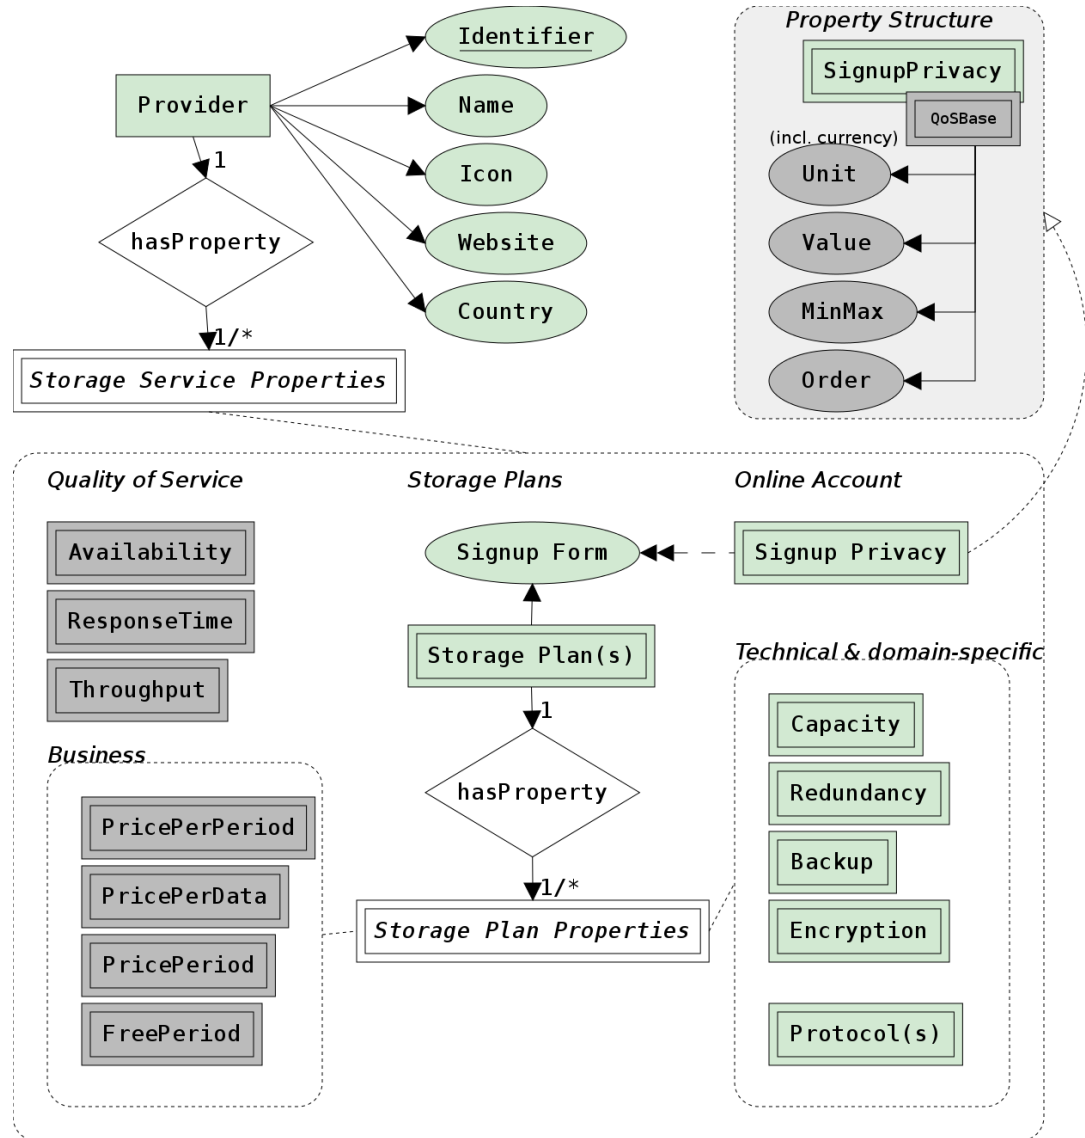
Instance Ontology

webService GoogleStorage
importsOntology {
 _"urn:ontology:conqo/CloudQoS.wsml#" }

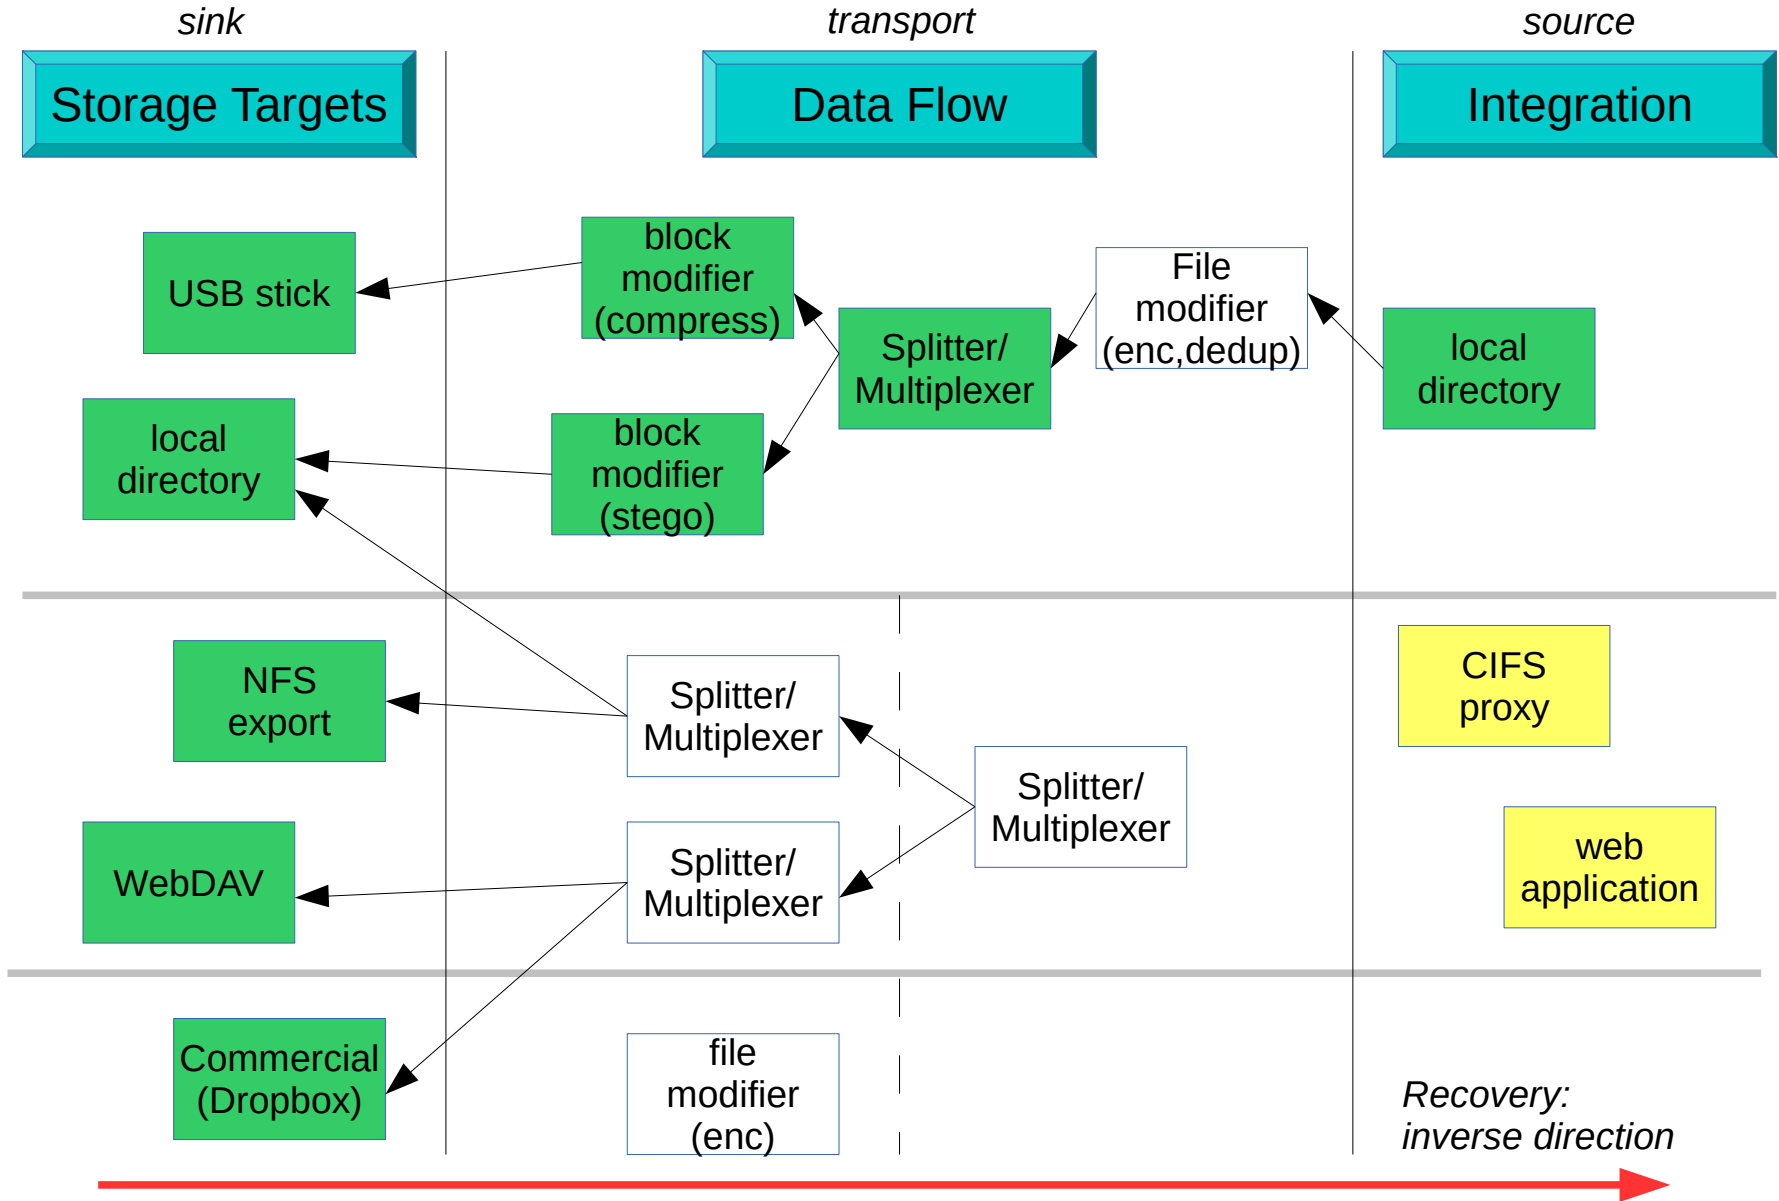
capability ServiceCapability
postcondition definedBy ?serviceType memberOf
cloud#CloudStorage .

instance PricePerData memberOf {
 cloud#PricePerData,
 qos#ServiceSpec }

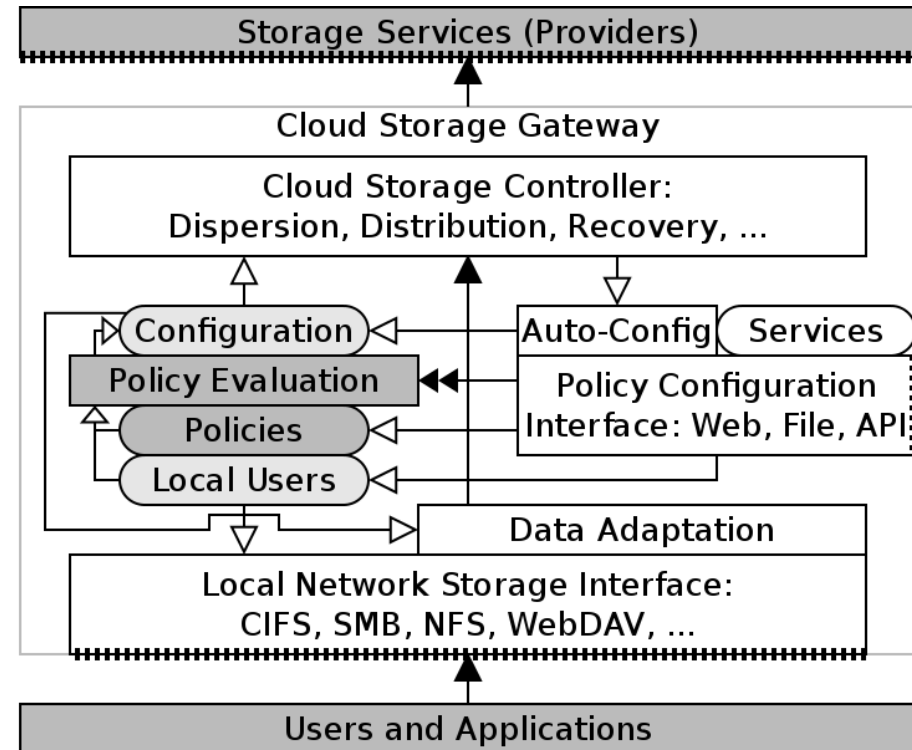

qos#value hasValue 0.17
qos#unit hasValue qos#Euro



Storage & Recovery Flows



Multi-User Storage Gateways



FlexDDPL: Flexible Data Distribution Policy Language

scopes → applied in contexts (e.g. user database)

~user

!negation

T:22:00-06:00

@group

fragment

mime:message/rfc822

rules → applied to targets (e.g. storage provider)

store

control

adapt

Background

- secret sharing, forward error correction / replication, visual dispersion, bitsplitting

Striping (→ RAID0)

- capacity 1.0, safety 1.0, performance 2.0

Mirroring (→ RAID1)

- capacity 0.5, safety 2.0, performance 1.0

XOR parity (→ RAID5, RAID6)

- 4 disks: capacity 0.75, safety 1.33, performance 1.0

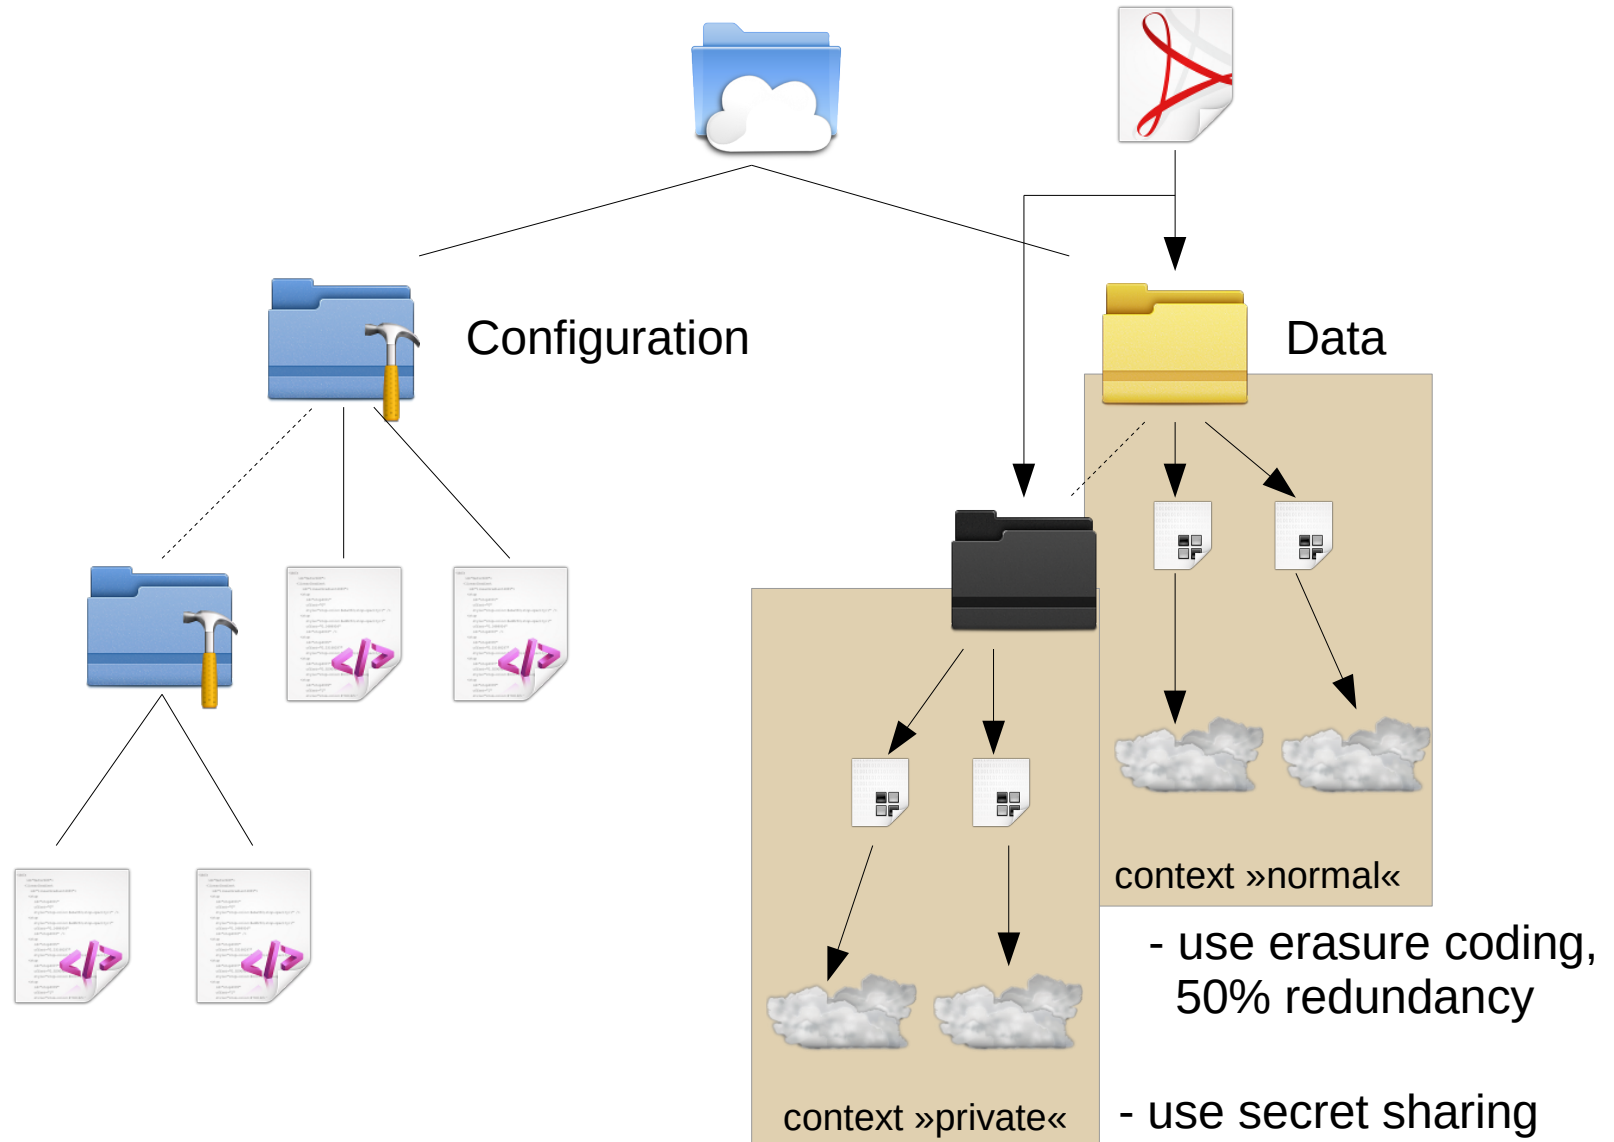
Erasure coding (→ RAIDn)

- n storage targets, k significant, m redundant: capacity k/m , safety $1+m/k$
- (Rotated) Reed-Solomon:
maximum-[Hamming-]distance separable (MDS) property
- Cauchy-Reed-Solomon, Vandermonde-Reed-Solomon:
practically secure; CRS faster due to XOR operations compared to $GF(2^w)$
- LRS: Locally Repairable Codes, e.g. Xorbas (Facebook) 10+4+2 coding

AONT: information-theoretically secure


AONT-RS: blend of AONT with (C)RS; alternative: encrypted fragments

Feature: Nested Contexts



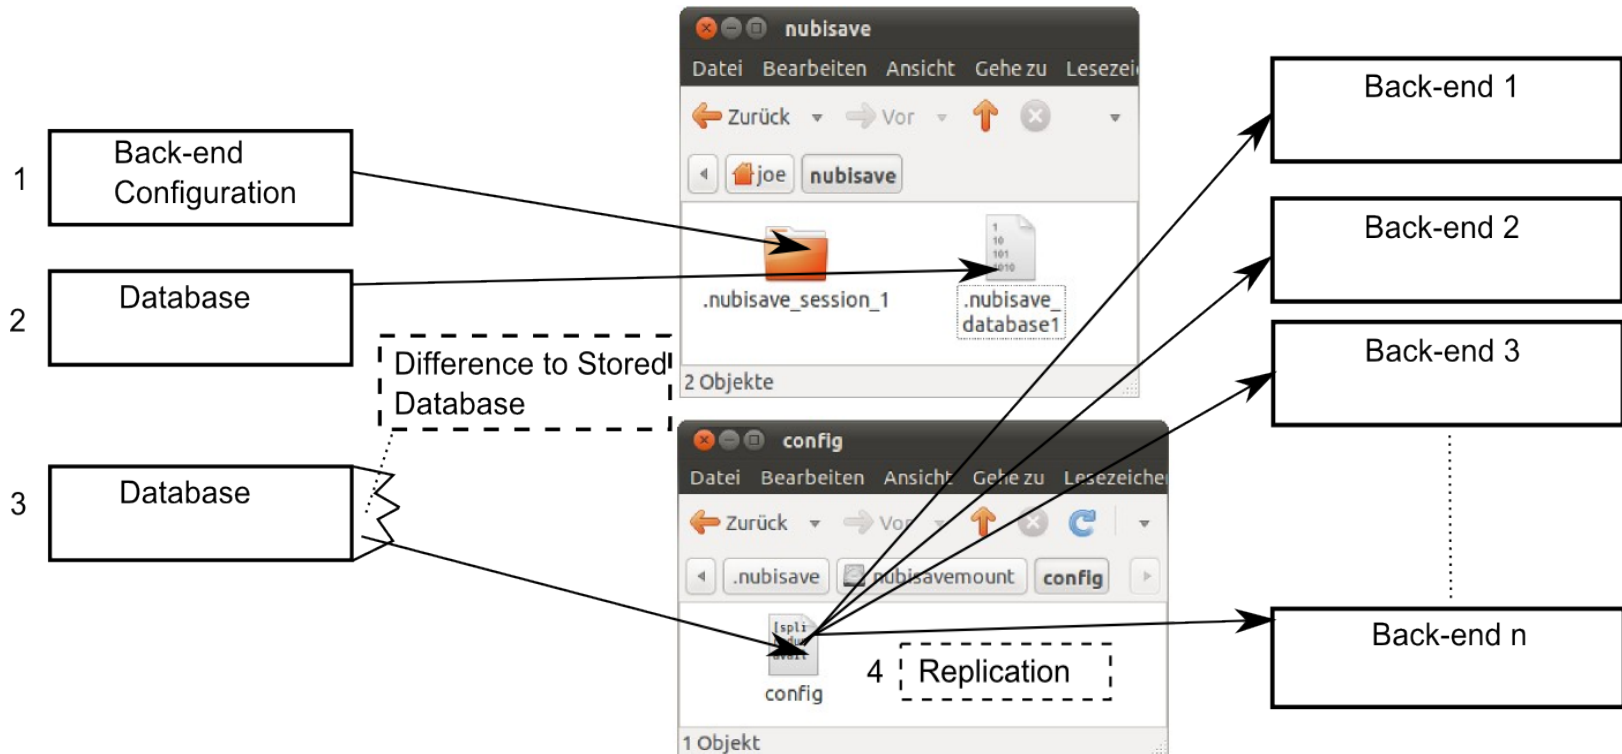
Assigning weights to the storage targets

- to fully utilise the capacity
- to exploit faster upload connections



Saving the database which contains all file and fragment metadata...

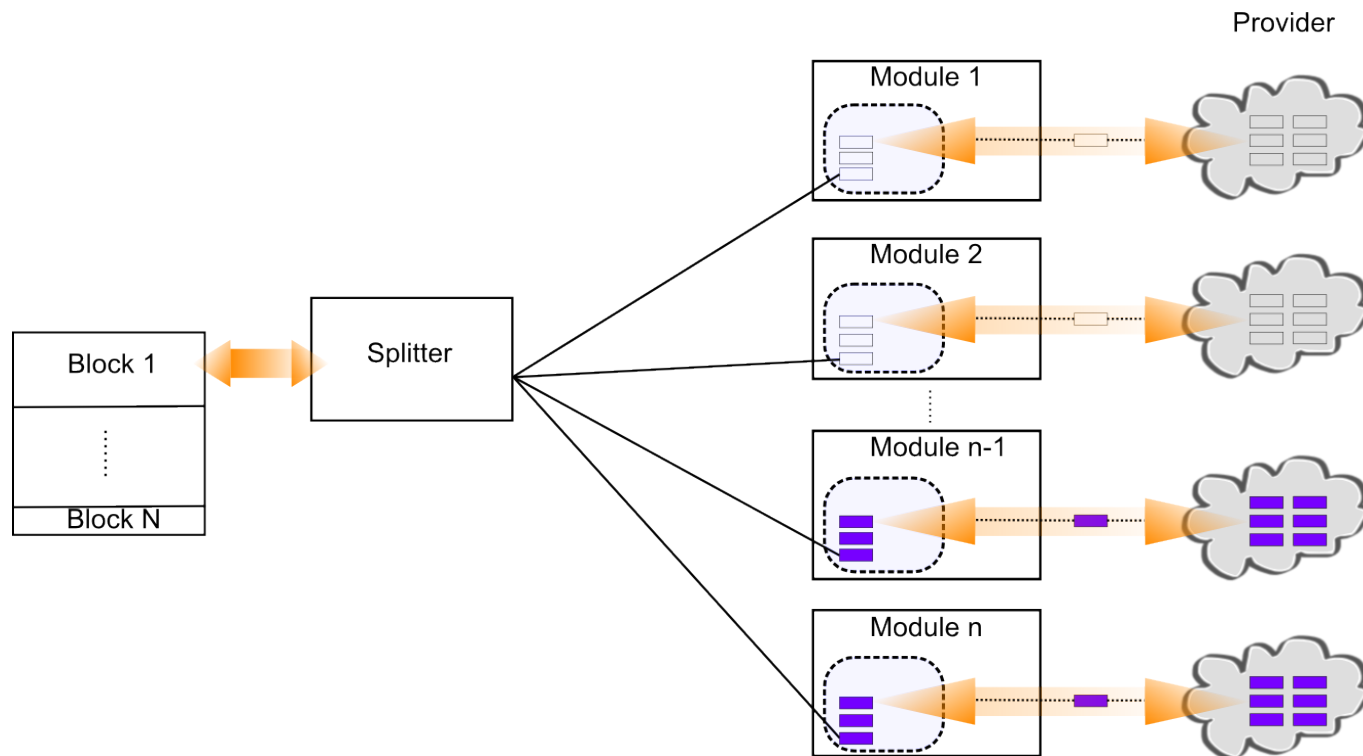
- to avoid single point of failure
- to allow for selective sharing

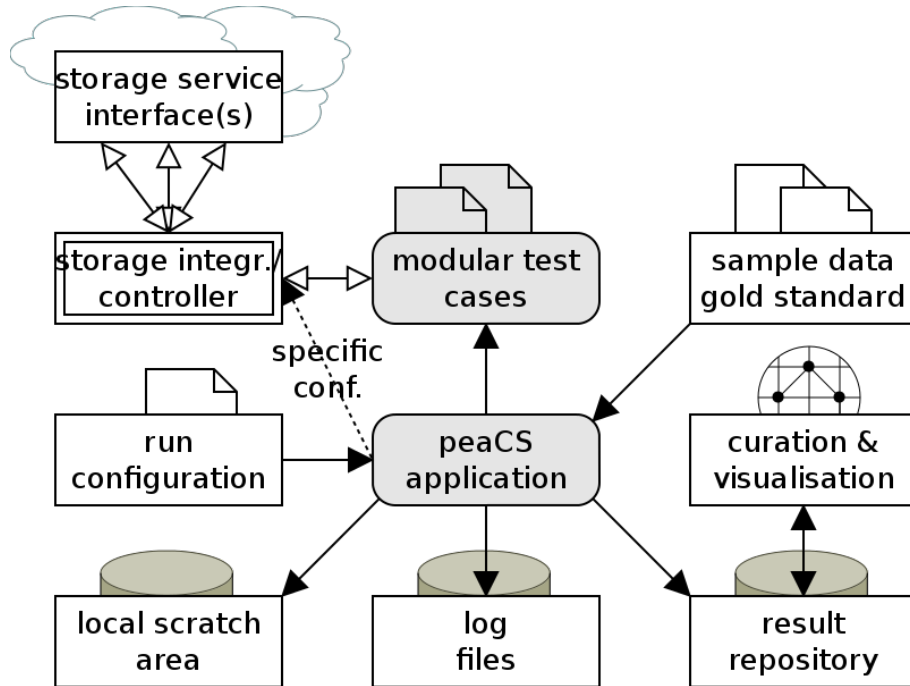


Streaming: Continuous operation during read or write of large file

- buffers on controller device can be kept small
- improved performance through parallel coding and transmission

Batching: Joining read and write request for many small files



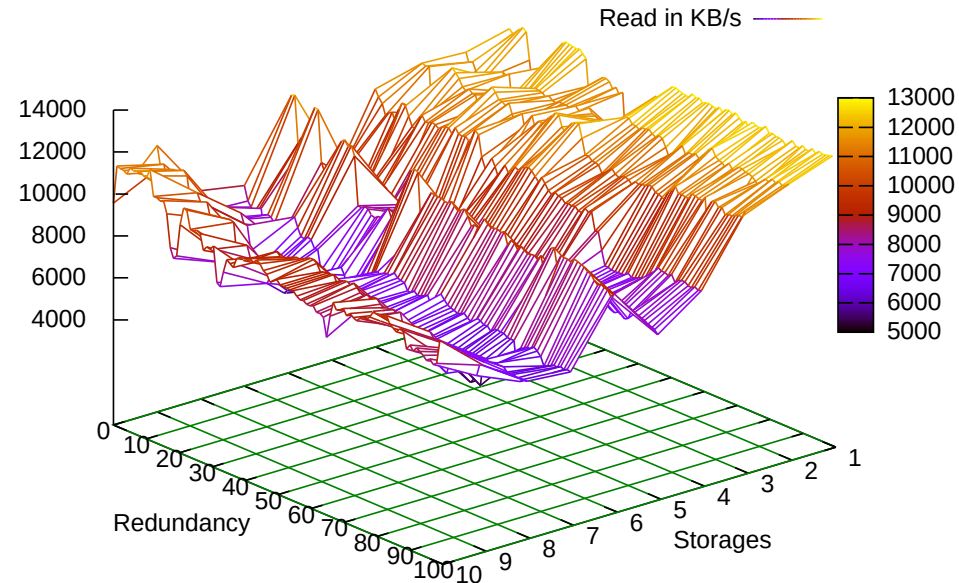


peaCS: Performance and Efficiency Analysis for Cloud Storage




Talk video available:
<http://www.tele-task.de/archive/lecture/overview/7682/>

Nubisave Read Performance - 1MB File - UseAllInParallel

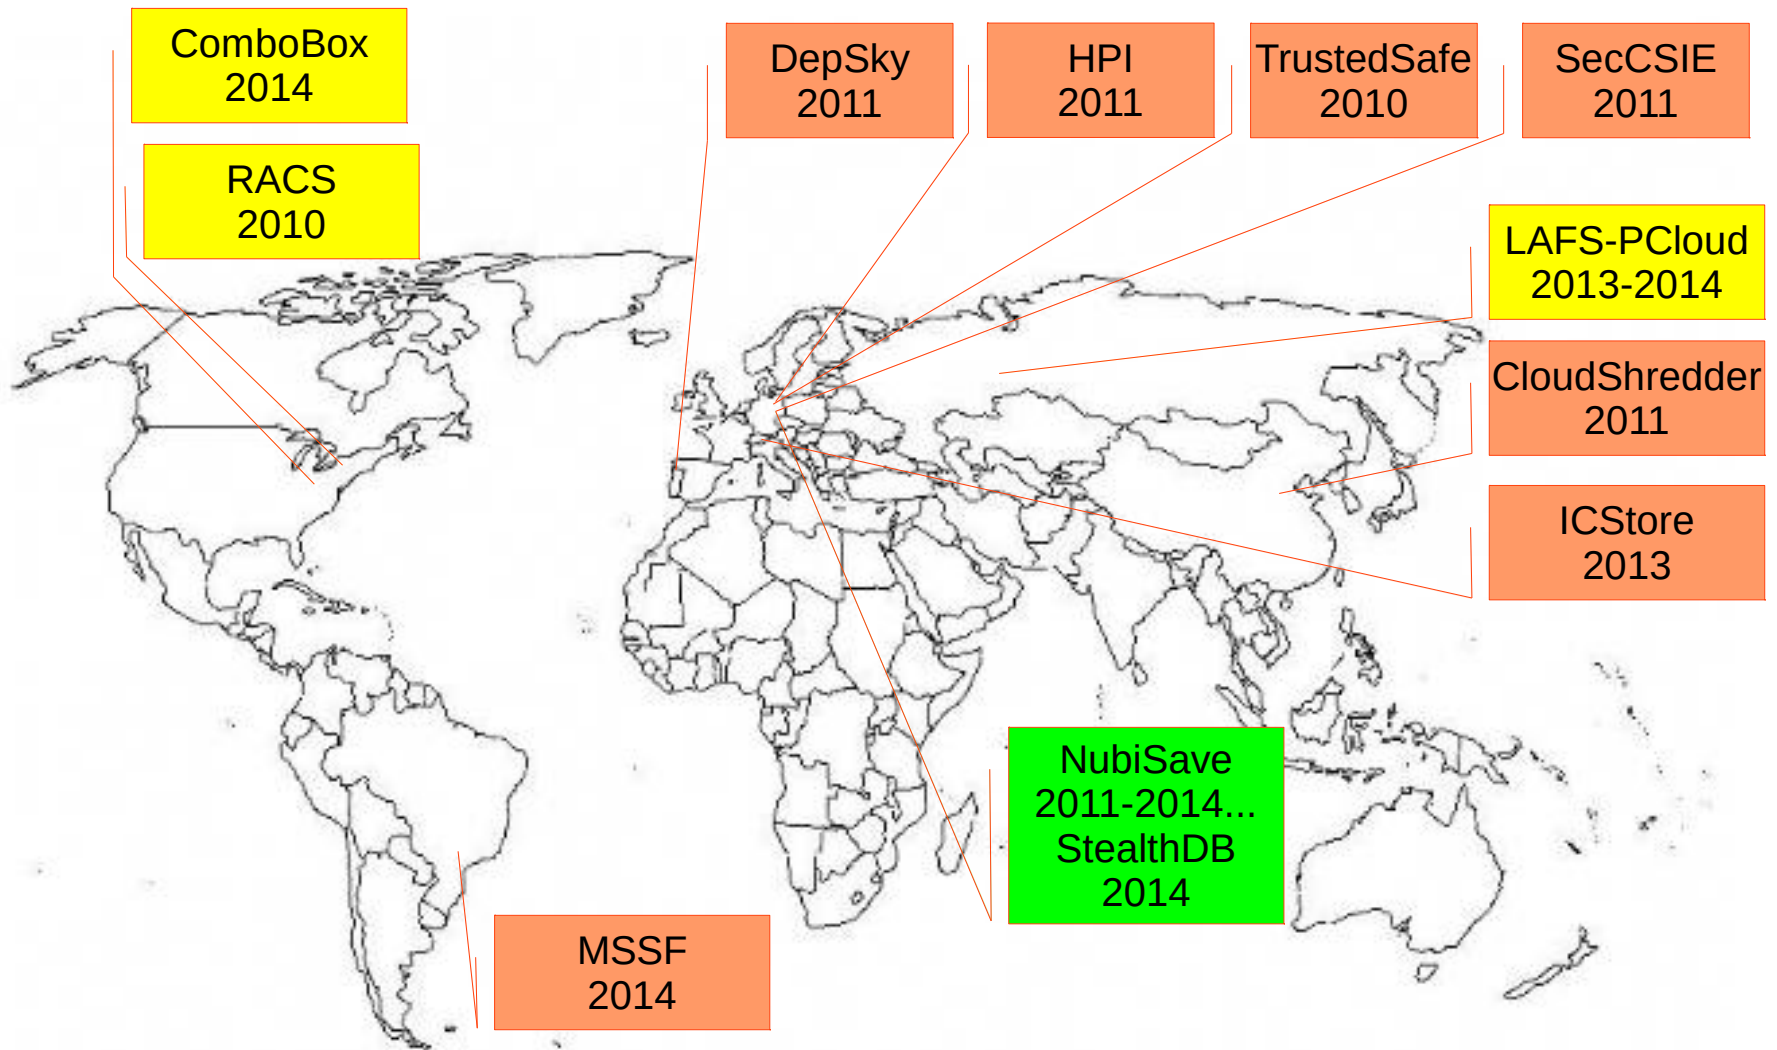


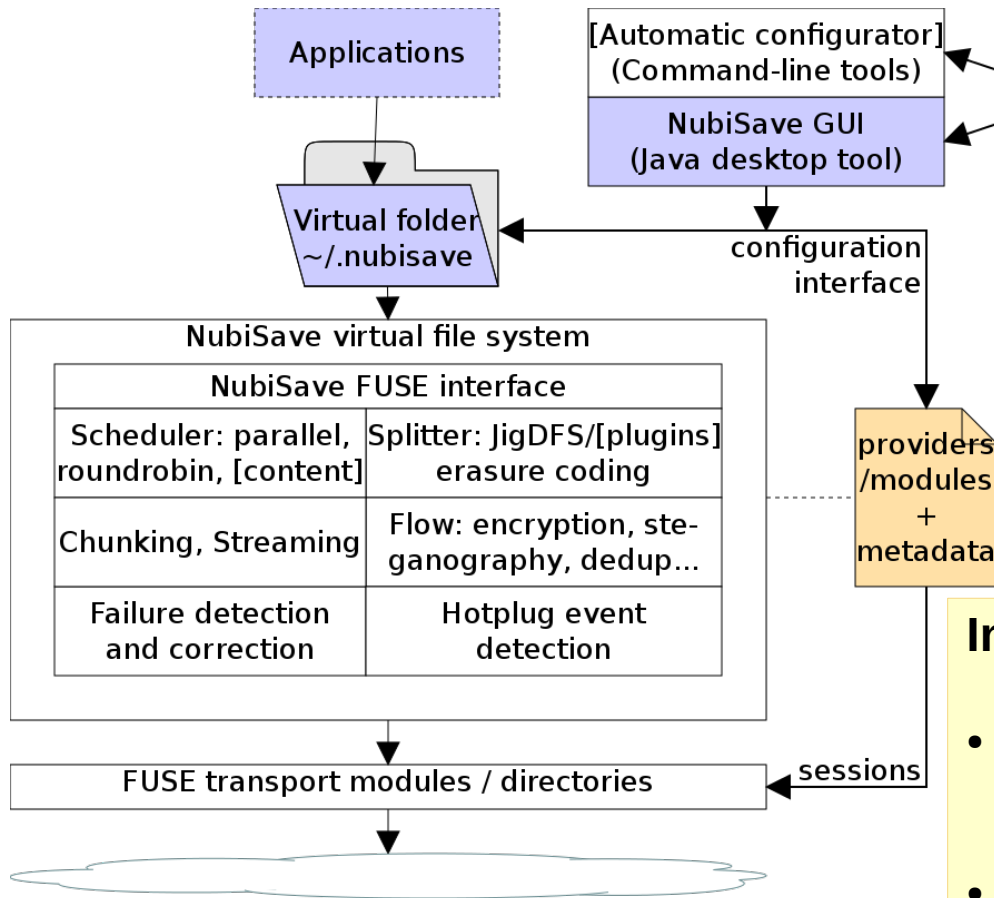
Storage controller power consumption



Software

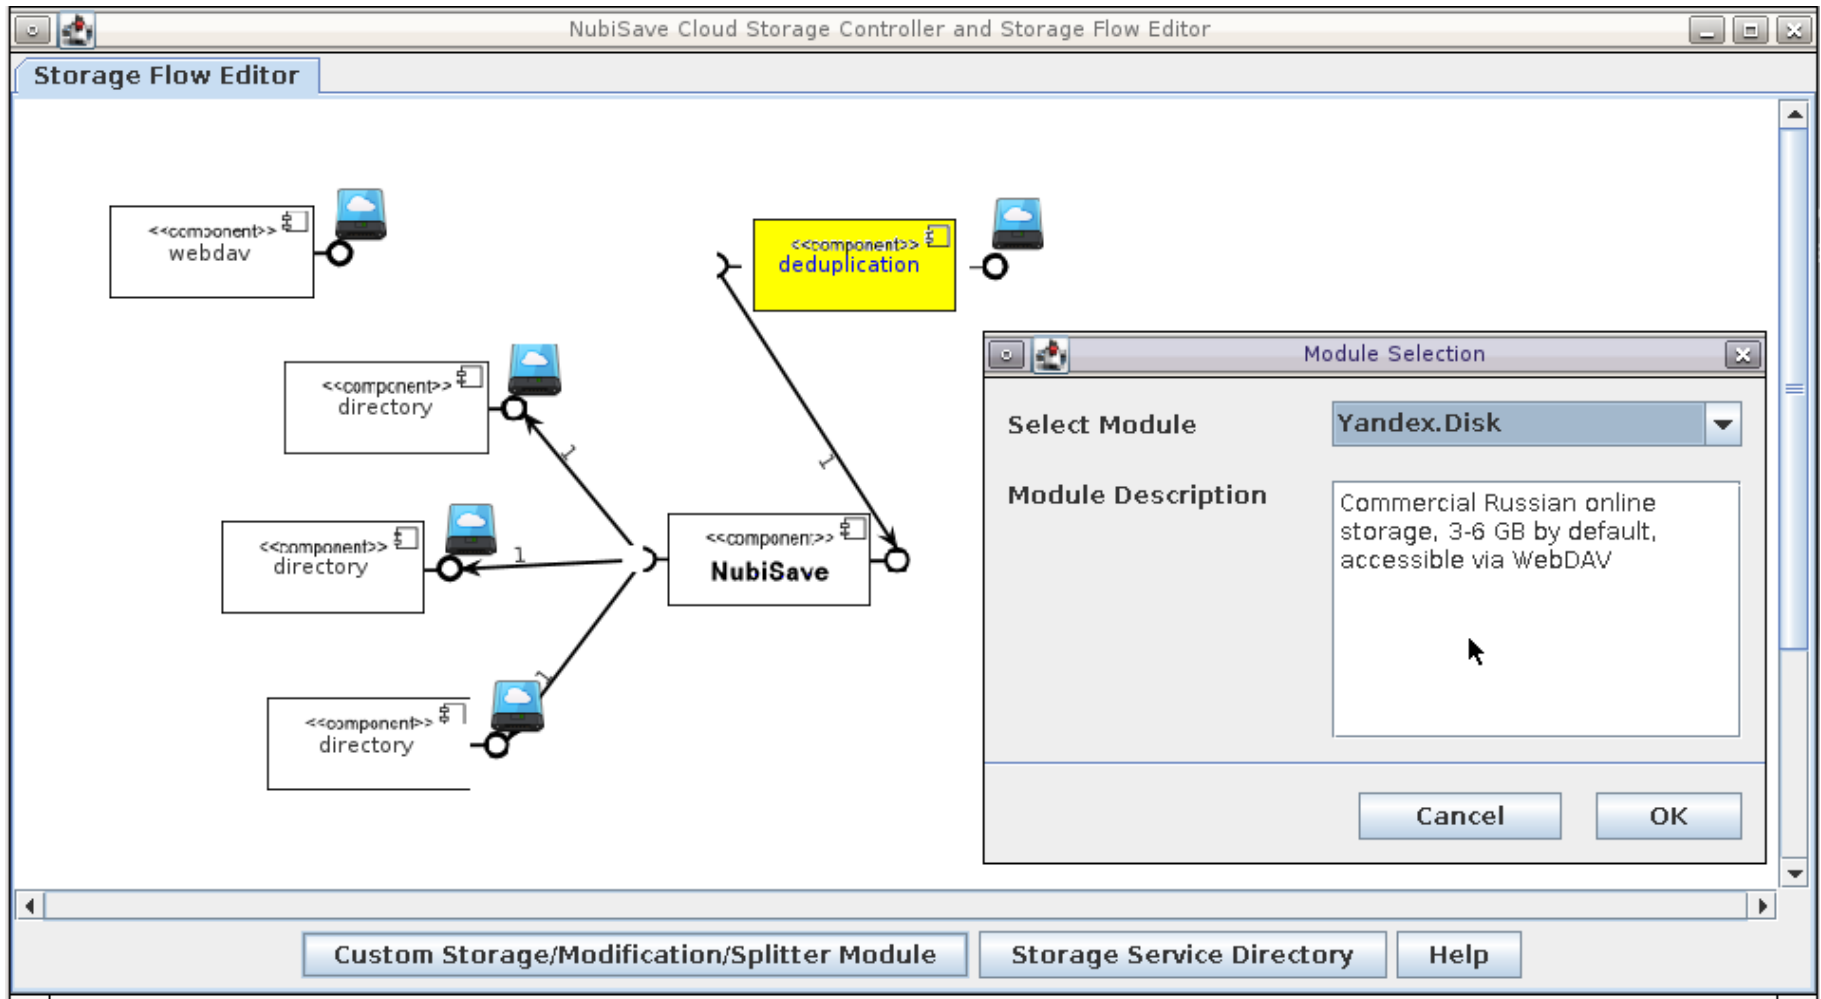
World Map of Storage Controllers





Installation of NubiSave

- Git repository
 - git://nubisave.org/git/nubisave
- Debian packages
 - <http://nubisave.org/packages/>
- Integrated into π -Box or NubiGate VMs
 - <http://nubisave.org/downloads/>



Profile selection:

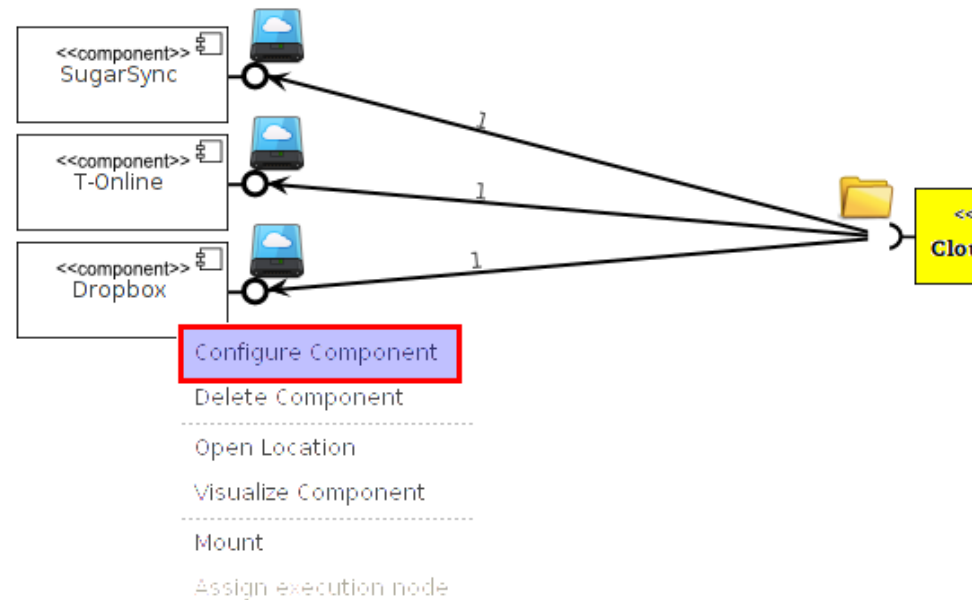
Free of Charge

Profile description:

Free of charge profile that costs nothing;
Sugarsync will only work free of charge for 30 days
Then you need to switch to a paid account.

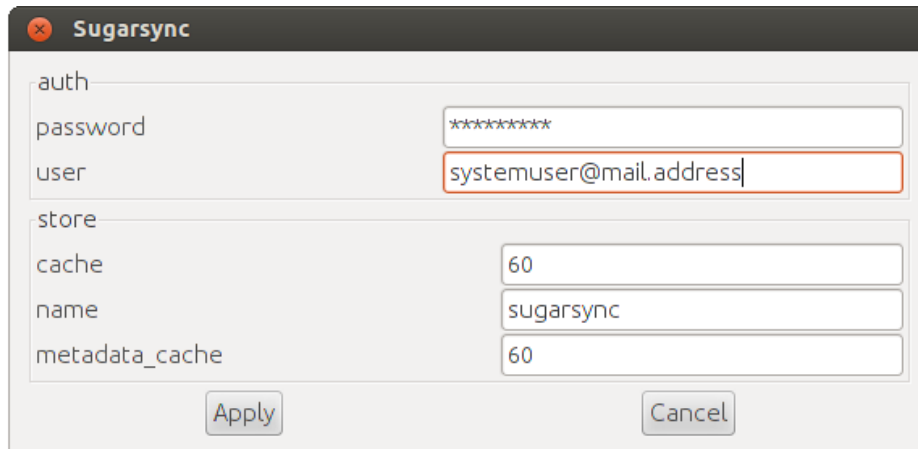
☒ SugarSync
☒ T-Online
☒ Dropbox

Apply Cancel



Scenario:
Pervasive storage on all devices

Credentials may be auto generated

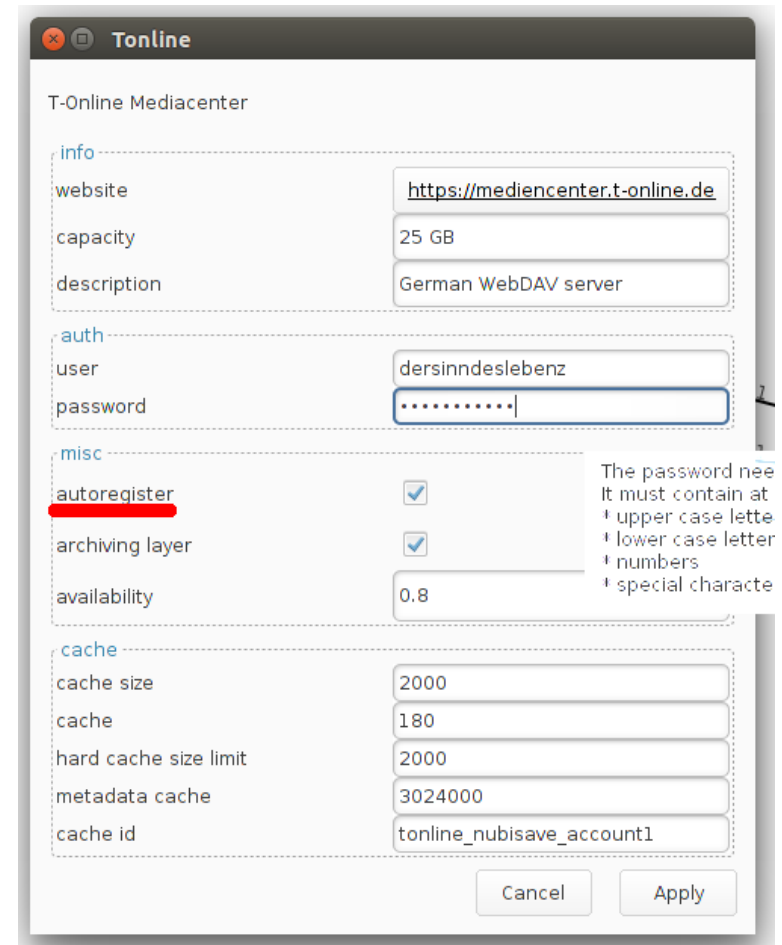


The Sugarsync configuration window shows fields for authentication and storage. The 'user' field is highlighted with a red border, indicating auto-generated credentials.

Category	Field	Value
auth	password	*****
	user	systemuser@mail.address
store	cache	60
	name	sugarsync
	metadata_cache	60

Buttons: Apply, Cancel

Auto-registration vs. Captchas

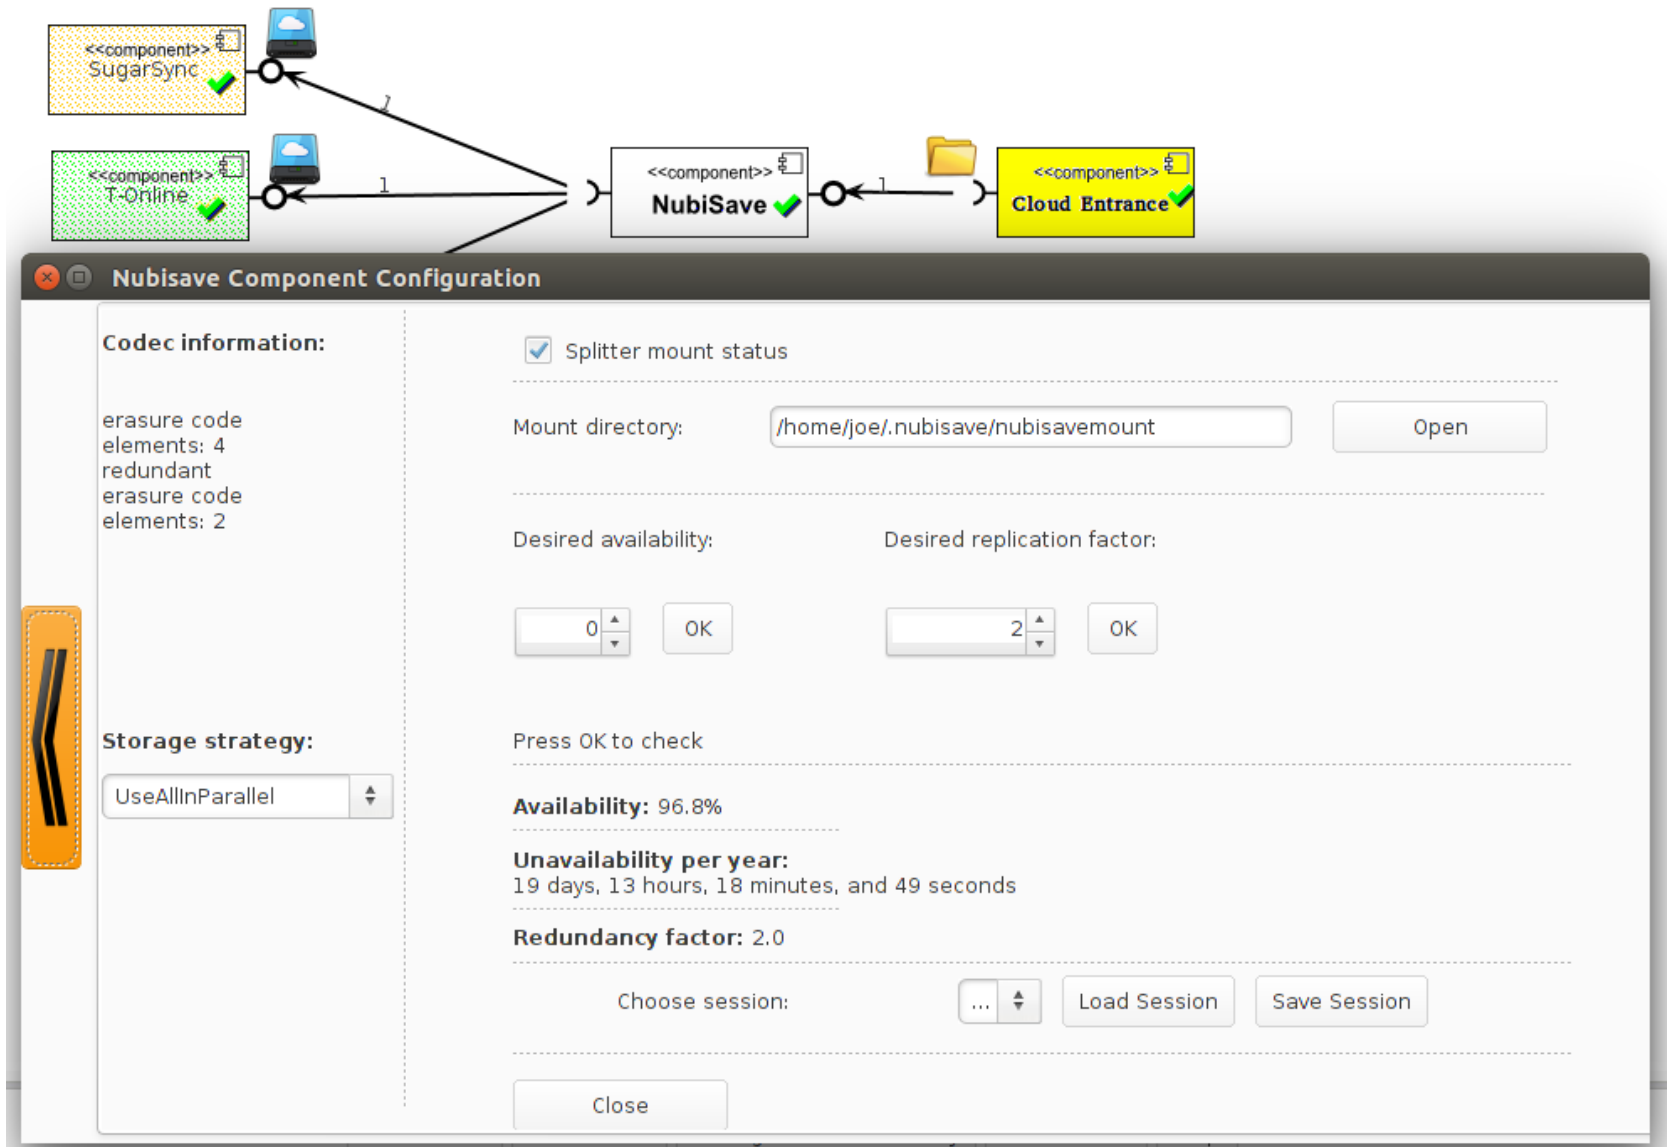


The T-online configuration window shows fields for account information, authentication, and storage. The 'autoregister' checkbox is checked, indicating auto-registration.

Category	Field	Value
info	website	https://mediencenter.t-online.de
	capacity	25 GB
	description	German WebDAV server
auth	user	dersinndeslebenz
	password
misc	autoregister	<input checked="" type="checkbox"/>
	archiving layer	<input checked="" type="checkbox"/>
	availability	0.8
cache	cache size	2000
	cache	180
	hard cache size limit	2000
	metadata cache	3024000
	cache id	tonline_nubisave_account1

Buttons: Cancel, Apply

The password needs to contain at least:
* upper case letter
* lower case letter
* numbers
* special character



Storage Service Selection

Property	Value	Unit	Priority
<input checked="" type="checkbox"/> Price per Data/GB	0.00	Euro	High
<input checked="" type="checkbox"/> Price per Period/Month	0.00	Euro	High
<input type="checkbox"/> Throughput	0.00	Mbps	Low
<input checked="" type="checkbox"/> Availability	99.00	%	Low
<input checked="" type="checkbox"/> Downtime	20.00	MilliSecond	Low
<input type="checkbox"/> Response Time	0.00	MilliSecond	Low
<input type="checkbox"/> Signup Privacy	0.00	%	Low
<input checked="" type="checkbox"/> Capacity	10.00	GB	Medium
<input type="checkbox"/> Redundancy	0.00	%	Low
<input type="checkbox"/> Encryption	0.00	0/1	Low

Name	Module	Pr/GB	Pr/Mo	Av	Th	Dt	Rt	SP	Ca	Re	En
(r:212.0) 4Shared (via CloudFusion) Fourshared		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
(r:210.0) GMX Mediacenter (via Clo... GMXMediacenter		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
(r:202.0) Box Free (via CloudFusion) Box		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
(r:200.0) Dropbox Free (via CloudF... Dropbox		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
(r:200.0) IDriveSync Free	IDriveSync	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
(r:112.0) Box (via CloudFusion) Box	Box	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
(r:112.0) Amazon S3 (via CloudFus... AmazonS3		<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Add selected services
☐ Update service database

Search

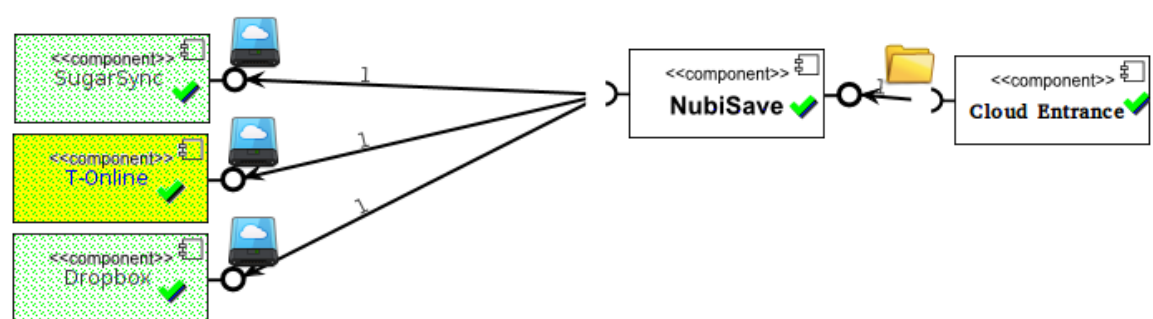
Cancel

NubiSave Cloud Storage Controller and Storage Flow Editor

Component Information

Downloaded: 78.70464MB
Uploaded: 2559.355MB
Download rate: -0.0MB per seconds
Upload rate: 0.26720738MB per second
Cached data: 1744MB
Last heartbeat: 0s ago
Number of errors: 2
Last error StoreAccessError has occurred 7998 seconds ago.
Description: HTTP/1.0 423 Locked
Files to synchronize: 1

Some files that still need to be uploaded:



```
graph LR; SugarSync[<<component>> SugarSync] -- 1 --> NubiSave[<<component>> NubiSave]; T-Online[<<component>> T-Online] -- 1 --> NubiSave; Dropbox[<<component>> Dropbox] -- 1 --> NubiSave; NubiSave -- 1 --> CloudEntrance[<<component>> Cloud Entrance];
```

Load Profile Add Module Storage Service Directory Visualization Help

Controller operation

\$ nubisave [<instance>]

Master script which starts both of the below combined.

\$ nubisave headless

Starts a new instance of the splitter/dispersion file system.

\$ nubisave gui

Starts the storage flow editor with storage integration configuration.

\$ nubisave stop

Stops the splitter.

Storages and database

\$ nubisave-status

Lists all splitter instances and storage targets attached to them or as part of any storage flow.

\$ nubisave-mounter [<module>]

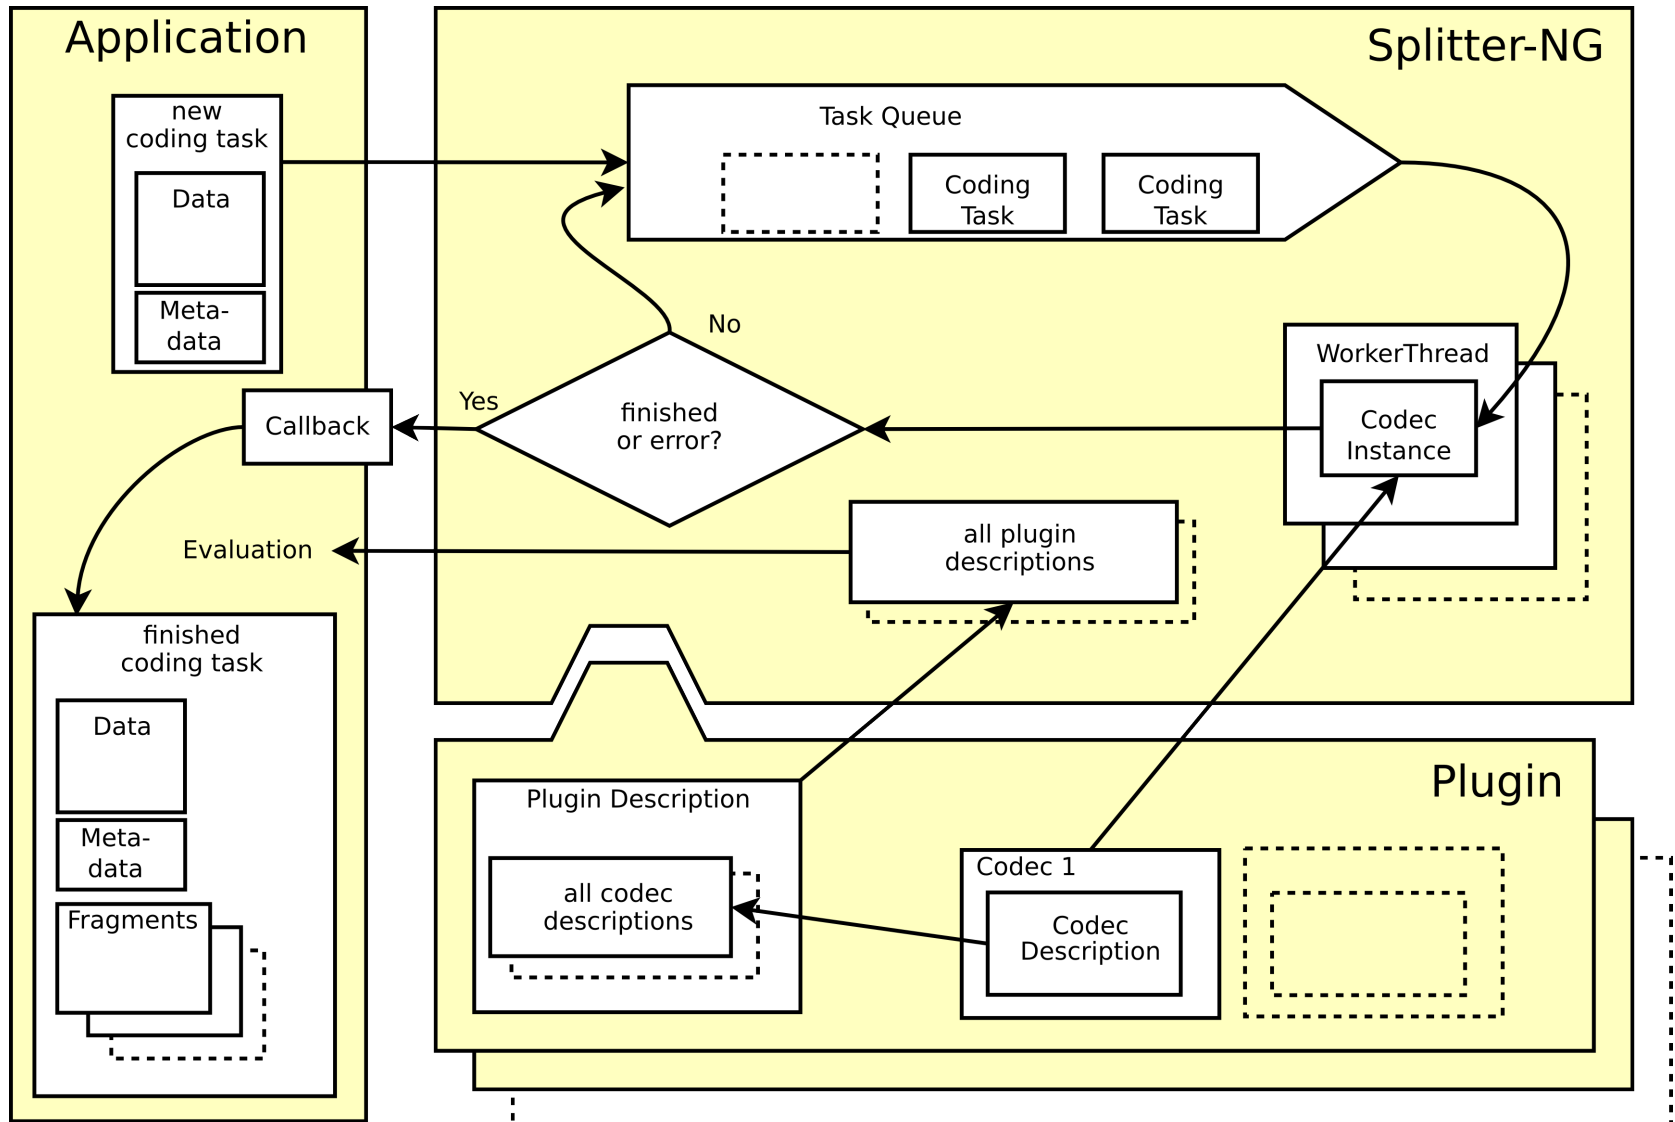
Mounts all splitters and/or storage targets.

\$ nubisave-unmounter [<module>]

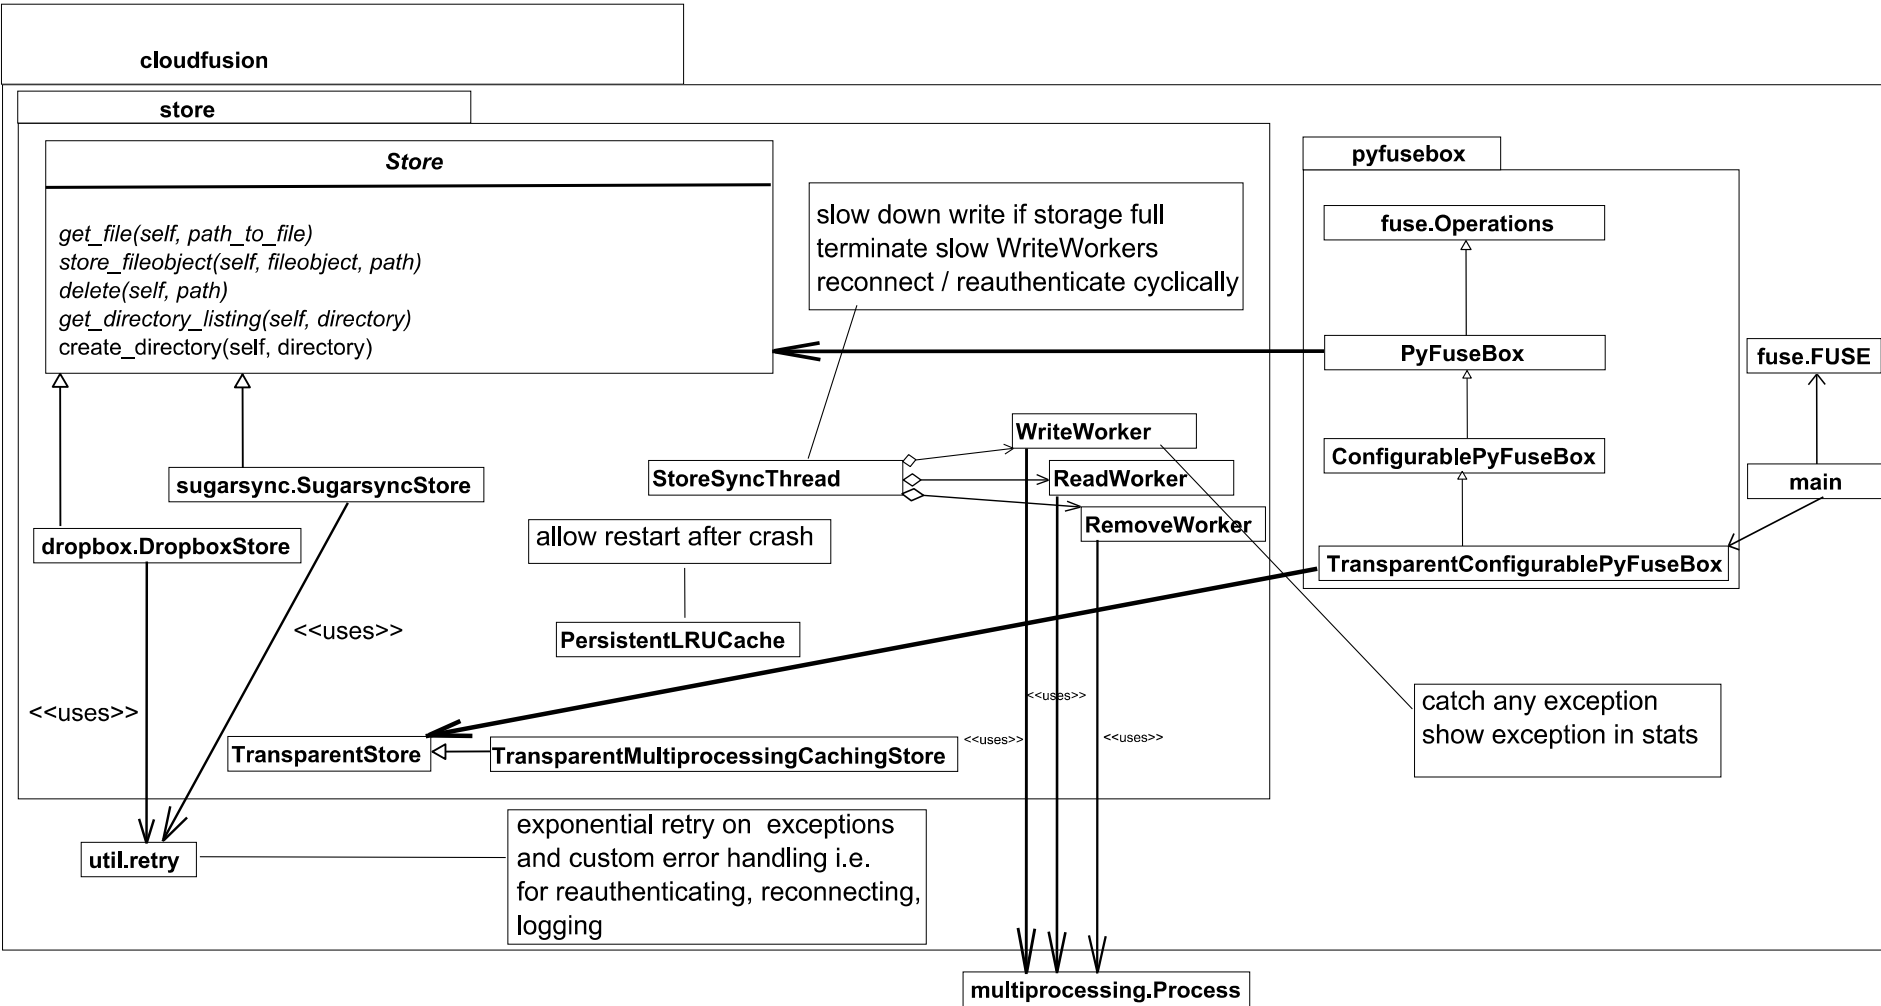
Unmounts; inverse of the above.

\$ nubisave-database [...]

Inspection of storage metadata.



Plugins: Jerasure, JSharing, RAID-1, Bitsplitter, more coming...



Anwendungen Orte Do, 11. Jul, 16:12

NubiVis - Visualization of d x Neuer Tab x Einstellungen x

localhost:8080/webapplication/src/index.html

View Menu

Start View

Tree View

New Tree View

Settings...

StorageView

Map View

Cover View

Bubble View

File Size View

Settings View

Search

Nubisave >

refresh

Path	File Size	Fragments
Musik	135,9 M	
Stairway To Heaven Unplugged Acoustic.mp3	8,1 M	
+ Galloglass	46,3 M	
+ Yngwie Malmsteen	23,0 M	
+ Bruce Dickinson	58,4 M	
+ Bilder	5,6 M	
+ Dokumente	655,9 K	
sicherheitsrisiken bei clouddiensten.pdf	139,1 K	
kuendigungsschreiben-wohnung2.pdf	27,6 K	
Musterbrief Bing Maps Streetside.pdf	21,5 K	
btwf.ppt	457,5 K	
WordDokument.docx	10,1 K	
+ Videos	157,3 M	

Filter Menu

Storage Provider

☒ WebDav

☒ TUD ZIH sshfs

☒ /home/mydata

☒ Windowsfreigabe

☒ FTP Server

nmo

0.9

nexif

22-rdf-syntax-ns

nfo

nie


nco







nmm

Elements Resources Network Sources Timeline Profiles Audits Console

Errors Warnings Logs Debug

martin@de... webapplica... wimd Dokumente NubiVis - V... [bubbleVie... [gitk: uploa... Nubisave Musterbrie...

NubiVis - Visualization of distr... 

localhost:8080/nubivis/    Google   

View Menu

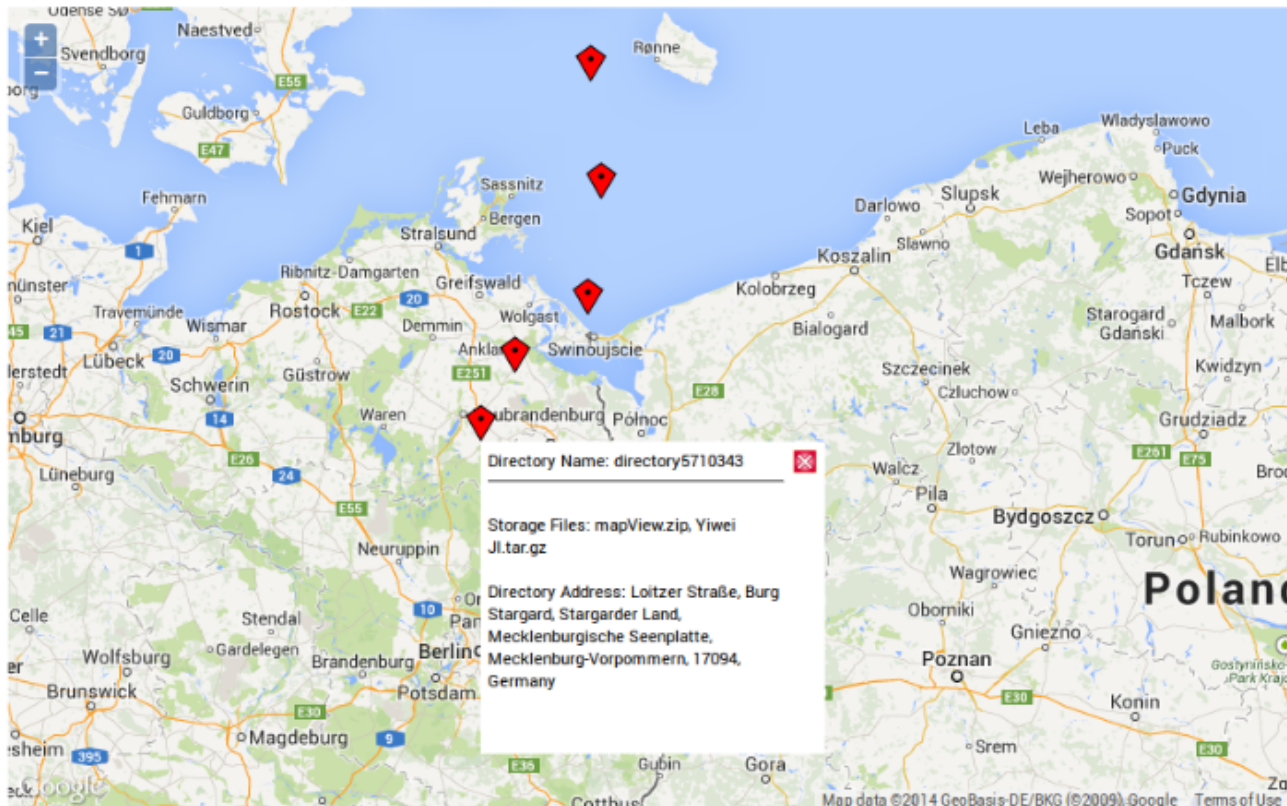
Home

Map

Settings...

Search

Nubisave >



Directory Name: directory5710343

Storage Files: mapView.zip, Yiwei JI.tar.gz

Directory Address: Loitzer Straße, Burg Stargard, Stargarder Land, Mecklenburgische Seenplatte, Mecklenburg-Vorpommern, 17094, Germany

Filter Menu

Storage Provider

- ☒ directory7609990
- ☒ directory5710343
- ☒ directory8306599
- ☒ directory3720047
- ☒ directory6221748

Globe

Tree

Bubbles

Size Chart

Storages

Discover

Documents

Settings





Benutzername

Passwort

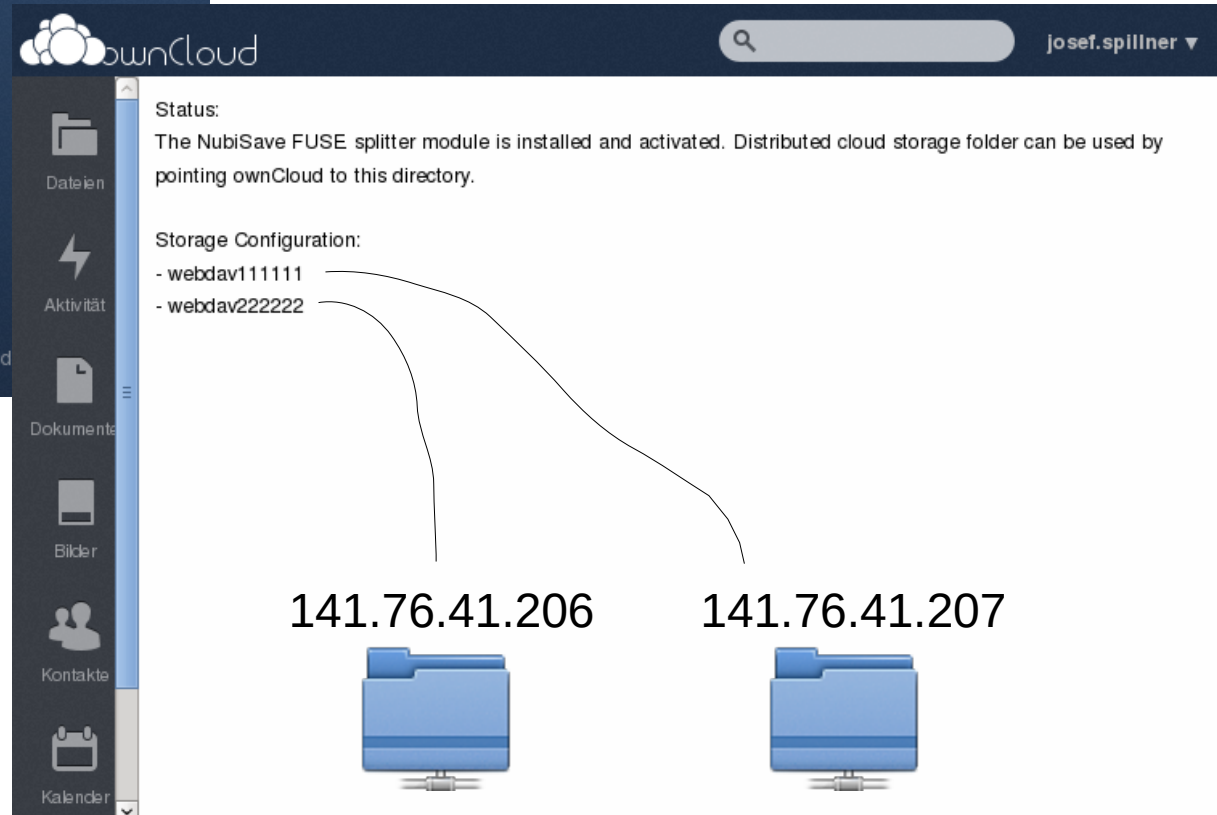
☒ merken

Einloggen

ownCloud – web services under your control | NubiSave – optimal cloud

Our instance „RN-Dropbox“
141.76.41.232/owncloud

Your instance...?

A screenshot of the ownCloud web interface. The top bar shows the ownCloud logo, a search bar, and the user name "josef.spillner". The left sidebar contains icons for "Dateien", "Aktivität", "Dokumente", "Bilder", "Kontakte", and "Kalender". The main content area displays a status message: "Status: The NubiSave FUSE splitter module is installed and activated. Distributed cloud storage folder can be used by pointing ownCloud to this directory." Below this, the "Storage Configuration" section lists two entries: "- webdav111111" and "- webdav222222". Two lines connect these entries to folder icons below. The first line connects "- webdav111111" to a folder icon labeled "141.76.41.206". The second line connects "- webdav222222" to a folder icon labeled "141.76.41.207".

ownCloud

Status:
The NubiSave FUSE splitter module is installed and activated. Distributed cloud storage folder can be used by pointing ownCloud to this directory.


Storage Configuration:

- webdav111111
- webdav222222

141.76.41.206

141.76.41.207

A fragment in the cloud...




... what can we do with it?

Dispersed Processing:

- structure-preserving bitsplitting
 - => search (any data)
 - => arithmetics, statistics (structured data)

Encrypted Dispersed Processing:

- homomorphic encryption
- order-preserving encryption




```
josef@rumba (git) dispersedalgorithms/db$ ./stealthdb
~~ StealthDB ~~
Using database stealthdb.
Storing all data and performing all procedures on ['mem://localhost'].
>>> HELP;
StealthDB Quickhelp
SHOW DATABASES|TABLES
CREATE TABLE <table> [(column type, ...)]
DESCRIBE <table>
DROP TABLE <table>
CREATE DATABASE <database>
USE DATABASE <database>
DROP DATABASE <database>
[EXPLAIN ANALYZE] SELECT */SUM(*)/AVG(*)/column, ... FROM <table> [ORDER BY <column> [ASC|DESC] [OPTIMIZE FOR <goal>]]
INSERT INTO <table> (column, ...) VALUES (value, ...)
USE CLOUDS <cloud> [AND <cloud>...][WITH <distribution>]
MODE <mode>
>>> █
```

[illegible]

Summary

Prototypes:

<http://nubisave.org/> - NubiSave Cloud Storage Controller

<http://lab.nubisave.org/stealthdb/> - StealthDB database

What's next in our lab...

- StealthDB performance optimisation
- StealthDB security: „proof of possession“ protocol
- increased NubiVis/NubiSave integration, refactored metadata handling

What's next from your side...

- test and give feedback, please :)

- [STS14] Josef Spillner, Sebastian Tilsch, Alexander Schill:
NubiVis: A Personal Cloud File Explorer.
Submitted to 11th International Conference on Mobile and Ubiquitous Systems: Computing, Networking and Services (MobiQuitous), London, UK, December 2014.
- [SM14] Josef Spillner, Johannes Müller:
PIcav: Precise, Iterative and Complement-based Cloud Storage Availability Calculation Scheme. Submitted to *7th IEEE/ACM International Conference on Utility and Cloud Computing (UCC), London, UK, December 2014.*
- [SS14] Josef Spillner, Alexander Schill:
Towards Dispersed Cloud Computing.
2nd IEEE International Black Sea Conference on Communications and Networking (BlackSeaCom), Chişinău, Moldova, May 2014.
- [SS13] Josef Spillner, Alexander Schill:
Orchestration of Distributed Storage Targets through Storage Flows.
5th IEEE International Conference on Cloud Computing Technology and Science (CloudCom), Bristol, UK, December 2013.
- [SQ+13] Josef Spillner, Maximilian Quellmalz, Martin Friedrich, Alexander Schill:
peaCS - Performance and Efficiency Analysis for Cloud Storage.
Workshop on Cloud Storage Optimisation (CLOUSO) @ 2nd ESOCC, Málaga, Spain, September 2013.
- [SMS13] Josef Spillner, Johannes Müller, Alexander Schill:
Creating Optimal Cloud Storage Systems.
Elsevier Future Generation Computer Systems (FGCS), Issue 29(4), p. 1062-1072, June 2013. DOI 10.1016/j.future.2012.06.004.