

Elasto

Cloud Storage Library and Client

David Disseldorp

ddiss@samba.org



Cloud Storage

Introduction

- Infrastructure as a service
 - Hardware exists but only the interface is important
 - Public or private
- Flexible
 - Scalable
 - Provision instantly
 - Pay for what you use
 - Multiple performance and reliability tiers

Cloud Storage

Introduction

- Access anywhere
 - Public gateway

- Mostly reliable
 - Local and geographic redundancy
 - 99.9% service-level agreements

- Low maintenance

Cloud Storage

Pitfalls

- Insecure
 - Personal data transferred over the internet
 - Where does it land, who has access?
- Slow
 - Internet connectivity bottleneck
 - Distance from datacenter
- Vendor lock-in
 - Non-standard interfaces
 - Proprietary implementations

Cloud Storage

Pitfalls – continued

- Vulnerable to outage
 - Distributed system
 - Company insolvency
- Not cheap
 - \$30 per-month for first 1 TB of storage*
 - Extra for requests

* Amazon S3 “Standard Storage”, Region: US Standard

RESTful Interface

- Representational state transfer
- HTTP or HTTPS
 - GET, PUT, POST, DELETE, HEAD, etc.
- XML API
 - URL query parameters
 - Request / response headers
 - Authentication and Authorization
 - Body
 - Object data or XML

RESTful Interface

Entity Addressing

- Amazon S3
 - Bucket
 - Container with access and location policies
 - Object
 - Actual data

- Azure
 - Account
 - Container
 - Blob

RESTful Interface

Example

- Amazon S3 Get Object request

```
> GET /test.txt HTTP/1.1
```

```
Host: ddiss.s3.amazonaws.com
```

Object

```
Accept: */*
```

Bucket

```
Date: Mon, 19 Aug 2013 19:53:08 +0000
```

```
Authorization: AWS
```

```
AKIAIEUERTYQRFAYKWZQ:AeNuMcFDIZxw5lvnqF62tqn  
EgRk=
```

Access Key ID

Signature

```
< HTTP/1.1 200 OK
```


Elasto Cloud

Elasto

What is it?

- Cloud storage client
 - Library and command line utility
- Written in C
- Curl HTTP(S) transport
- Speaks Azure and Amazon S3 storage protocols
- SUSE Hackweek project

Elasto

Library

- Comprehensive support for storage requests
 - Amazon S3
 - Bucket create/list
 - Object upload/download/copy/delete/list
 - Microsoft Azure
 - Account and container create/list
 - Blob upload/download/copy/delete/list
- Request marshalling and unmarshalling
- Authentication and signing

Elasto

File IO Library

- POSIX style API
- Currently only supports Azure
 - Uses Azure's page-blob interface
 - Sparse file IO
 - Requires 512-byte aligned requests
 - S3 PUT requests do not accept object offset parameters
- Cmocka test suite

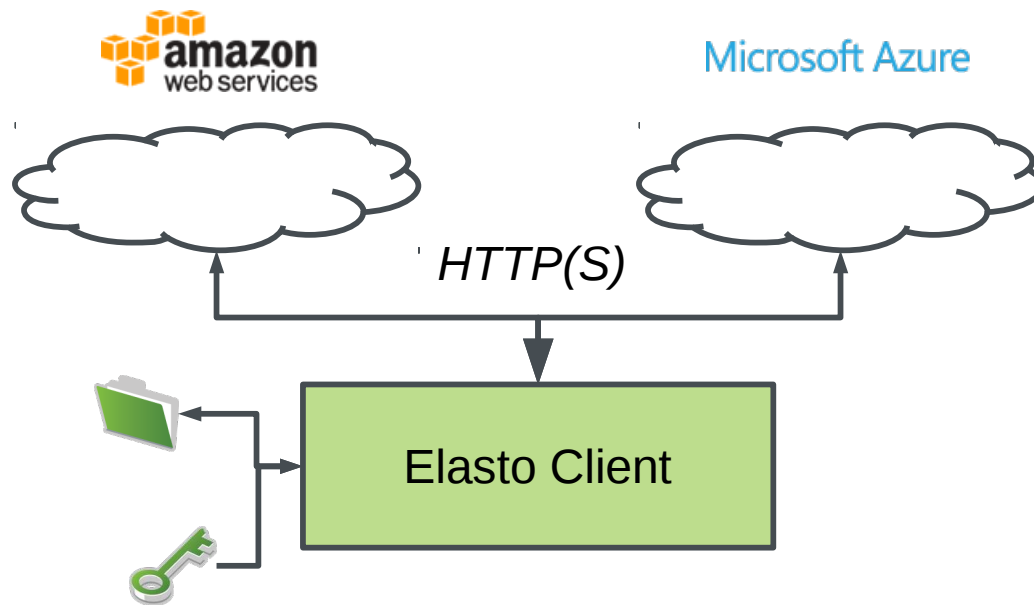
Elasto

Command Line Utility

- smbclient style syntax
- Abstracts away cloud differences
 - Authentication information
 - Amazon S3 Access Key ID and Secret Access Key
 - Publish Settings file for Azure
 - SSL Certificate
 - Storage entity paths
 - Amazon S3 /Bucket/Object
 - Azure /Account/Container/Blob

Demonstration

Elasto Command Line Client



iSCSI Cloud Gateway

Elasto + istgt

istgt and Elasto

Cloud Backed Block Device

- istgt
 - Userspace iSCSI target
 - Pluggable disk backend
- Elasto file IO backend
 - Provisions logical unit cloud object
 - SCSI requests mapped to cloud IO

istgt and Elasto

Cloud Backed Block Device

- Why?
 - Improved flexibility
 - Security
 - Reliability
 - Performance
 - Cost savings

istgt and Elasto

Cloud Backed Block Device

- Flexible
 - Standard interface
 - All enterprise operating systems provide an iSCSI initiator
 - No reliance on Cloud provider specific application or website
 - No application changes required
 - Integrates into existing work-flow
 - Use existing device and filesystem building blocks

istgt and Elasto

Cloud Backed Block Device

- **Reliable**

- RAID between local and cloud storage devices
- Fail-over between cloud storage providers (only Azure ATM)
 - Cloud-to-cloud backup

- **Secure**

- Transparent encryption
 - Dmccrypt, BitLocker, etc.
- No reliance on transport or endpoint security

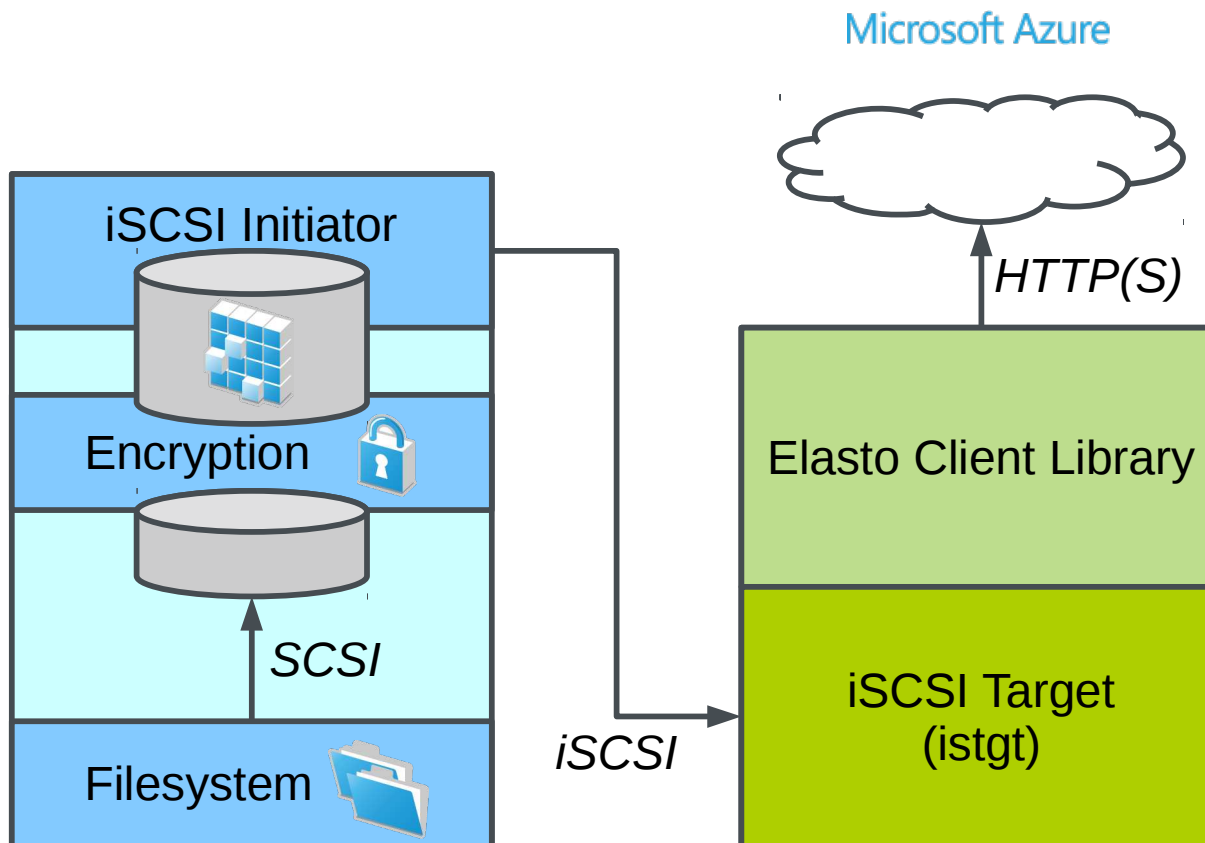
istgt and Elasto

Cloud Backed Block Device

- Performance
 - Use read or write-back cache
 - Bcache, dmcache, etc.
 - HSM integration
- Save money
 - Transparent compression
 - Btrfs, NTFS, etc.
 - Reduced IO
 - Optimised backup using Btrfs send / receive
 - Make the most of price fluctuations

Demonstration

Elasto Cloud iSCSI Gateway





Elasto

Future Challenges

- Support more cloud storage providers
 - Amazon S3 interface becoming an industry standard
- Pipeline download / upload operations
 - Curl support sub-par, GET & HEAD only
- File IO library support for Azure
 - S3 does not allow for write requests at specific ranges
- Testing
 - Expand existing cmocka unit test suite
- Distribution submission

<http://elastocloud.org>

Thank you.



