

Experiences with Cleversafe

Kaspars Krampis

About

- What is Cleversafe?
- How does it work?
- Our tests
- Results
- Conclusion

What is Cleversafe? (1)


- Dispersed storage technology
- Data are sliced using Information Dispersal Algorithms (IDAs)
- Slices are stored accross the network
- Subset of slices is neccessary to reconstruct the data
- Data are accessed using iSCSI

What is Cleversafe? (2)

- Consists of 3 components
 - Source computer (iSCSI initiator)
 - Accesser (iSCSI target)
 - Slicestors (actual storage)

What is Cleversafe? (3)

Dispersed Storage Infrastructure



What is Cleversafe? (4)

- OS: Linux (CentOS supported)
- Uses Java
- Dual licensed software
- GPLv2 and proprietary
- Additional components available only in commercial version

How does it work? (1)

- Typical configuration
 - 1 iSCSI initiator
 - 1 iSCSI target
 - 8 slice servers (6 needed to recreate the data)
- Can choose different configuration

How does it work? (2)

- Client (initiator) connects to target
- Gets “hard drive” (block device)
- Uses it as physical drive
 - Can partition
 - Use file system of choice

How does it work? (3)


- All the hard work does target machine
 - Slicing
 - Encrypting
 - Dispersing
 - Recreating
 - Managing

Our experience (1)

- Testing environment
 - 10 AMD Athlon64 3200+ machines
 - 100Mbit LAN
 - 32 and 64 bit Centos 5.2 Linux versions (for target and slice servers)
 - Linux and Windows iSCSI initiators

Our experience (2)

Cleversafe performance



Results (1)

- 64bit slice servers are faster
- Windows XP iSCSI implementation is slow
- Target machine is single point of failure
 - Both downloads and uploads data
 - Java process consumes 50-70% CPU
 - Load average 8 - 10

Results (2)

- ~40Mbit/s upload is max possible using 100Mbit/s network
 - 40Mbit/s upload from initiator results in 60Mbit/s upload from target
- Downloads should be faster than uploads but are not

1Gbit/s test (1)

- 1Gbit/s link from initiator to target
- 100Mbit/s link from target to slice servers
- Linux iSCSI initiator
- Target and slice servers: CentOS 5.2 64bit

1Gbit/s test (2)

- Client:

```
dd if=/dev/zero of=file bs=1M count=2000  
303.381s 6.9MB/s
```

- Target

```
vnstat -tr 180 -i eth0  
tx 9497.40 kB/s  
vnstat -tr 180 -i eth1  
rx 6781.19 kB/s
```

1Gbit/s test (3)

- Load on target: 10 – 15
- Java process consumes 80 – 90% CPU
- That's a lot!

Problems

- Some of advertised features doesn't work, most notably
 - Rebuilding of data
 - Encryption of data
- Should be implemented in upcoming version(s)

Conclusion

- Potentially usable technology
- Currently not ready for production use
- Not very fond of dual-licensing
- We'll keep an eye on Cleversafe
- Additional information: www.cleversafe.org

Thank you! Questions?