Chelonia:
A Self-healing Storage Cloud

KnowARC

Presented by:

Salman Toor [salman.toor@it.uu.se] (Uppsala University, Sweden)

Group members:

Jon K. Nilsen [j.k.nilsen@fys.uio.no] (Oslo University, Norway)

Zsombor Negy [zsombor@niif.hu] (NIIF, Hungary)

Bjarte Mohn [Bjarte.Mohn@fysast.uu.se] (Uppsala University, Sweden)

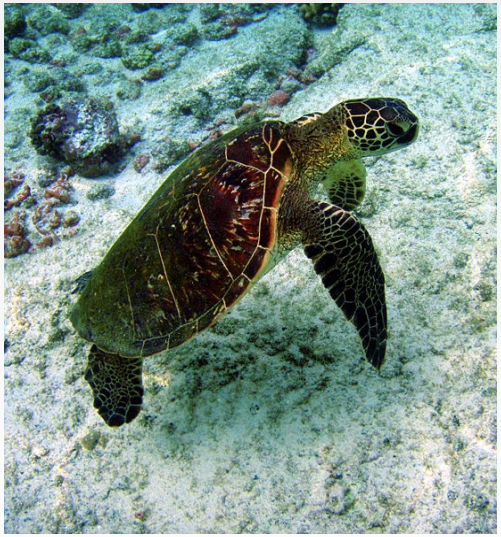

Chelonia

KnowARC

- What is Chelonia?
- Feature list of Chelonia.
- System components and Architecture.
  - ❖ Self-healing and High-availability.
  - ❖ Security framework
  - ❖ Accessibility of Chelonia
- Brief comparison with other systems.
- First proof-of-concept test results.
- Conclusion and Important links.

What is Chelonia?

KnowARC

- Chelonia is a storage cloud developed in the KnowARC project.
- It consists of a set of web-services, each has its well defined task.
- Chelonia services run in the HED (Hosting Environment Daemon) service container, also developed in the KnowARC project.

Feature List of Chelonia

KnowARC

- User-friendly interface.
- Global hierarchical namespace.
- Self-healing mechanism for available data.
- Non-intrusive architecture eliminate the single points of failure.
- Third party storage access.
- Flexible data access to the grid jobs.
- Secure third party file transfer.
- Replication of the metadata.

System Components

Bartender (B): Provides high-level interface for the clients.

Librarian (L): Catalog service, uses AHash to manage matadata.

AHash (A-H): Works as a metadata store. AHash can be for two types:

- Central AHash (For the small deployments)
- Replicated AHash (Berklay DB is used for replication)

Shepherd (S): Manages the storage node with the actual file.

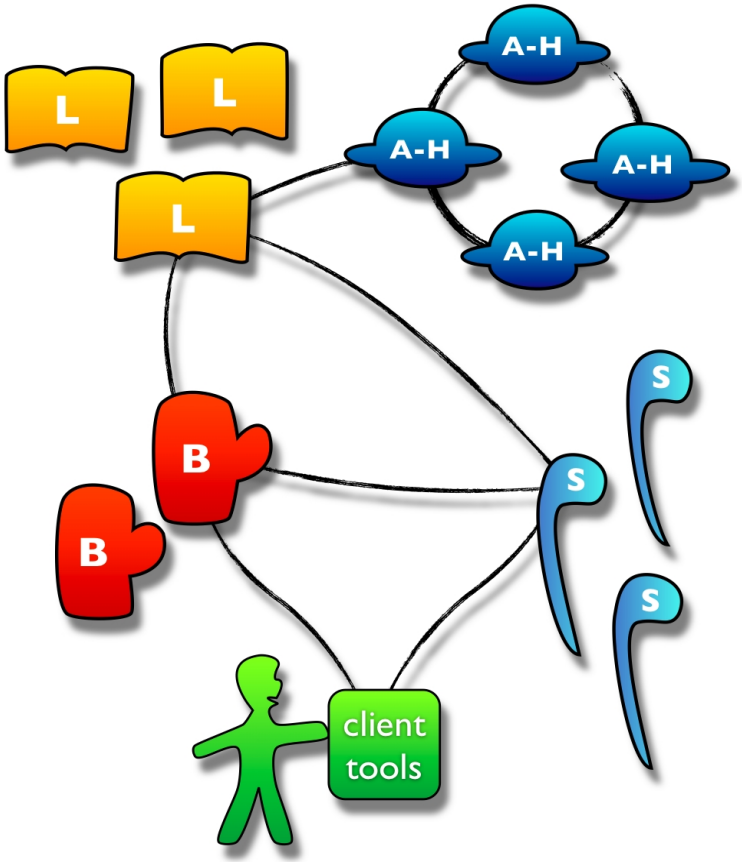

Self-healing and High availability

KnowARC

- Files in Chelonia are replicated, and broken replicas are repaired automatically by the system.
- Extra replicas are set to be ‘thirdwheel’, i.e. obsolete and will be removed from the store.
- The replication of the services gives high availability to the Chelonia cloud.
- A-Hash works in a master/slave model. If the master dies, available A-Hashes elect a new master. This feature allows self-healing/high-availability at the level of metadata.

Security Framwork

KnowARC

Three levels of security:

High level: Includes access policies on files, Collections and mountpoints.

Transfer level: A one-time Transfer URL (TURL) is generated and send to the client.

Service level: Services can be deployed anywhere. The integrity of the internal communication demands trust between the services in order to avoid attacks.

Accessibility of Chelonia

KnowARC

- Client tools
  - Command-Line-Interface for Linux, Windows and Mac.
  - FUSE module allows users to mount the global namespace into the local namespace.

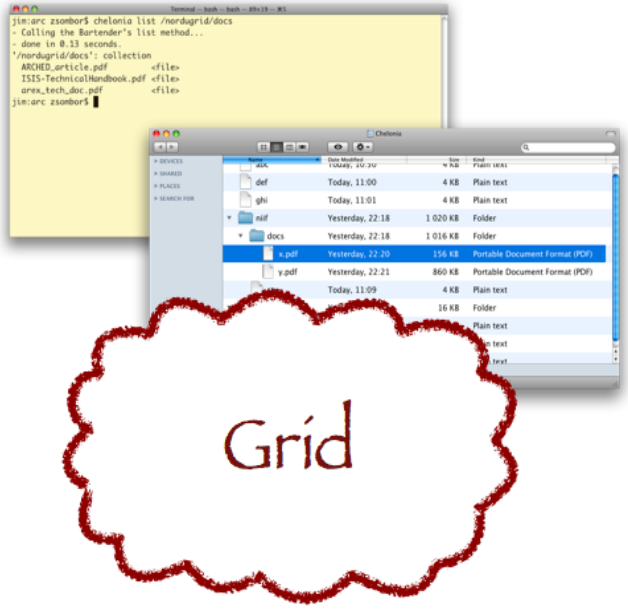

- Computing grids
  - Chelonia cloud can also be accessed from grid jobs.
- Third-party store
  - Chelonia allow Clients to mount the third party store in its global namespace.

Proof of concept test-cases

- Uploading 10 large size (1GB) and 100 small size (10KB) files with one replica each.
- Analyze the load distribution amongst the available shepherds.
- Increase the number of replicas for some files to 3.
- Kill one Shepherd and Librarian, try to download some of the files.
- End-to-End system response time.

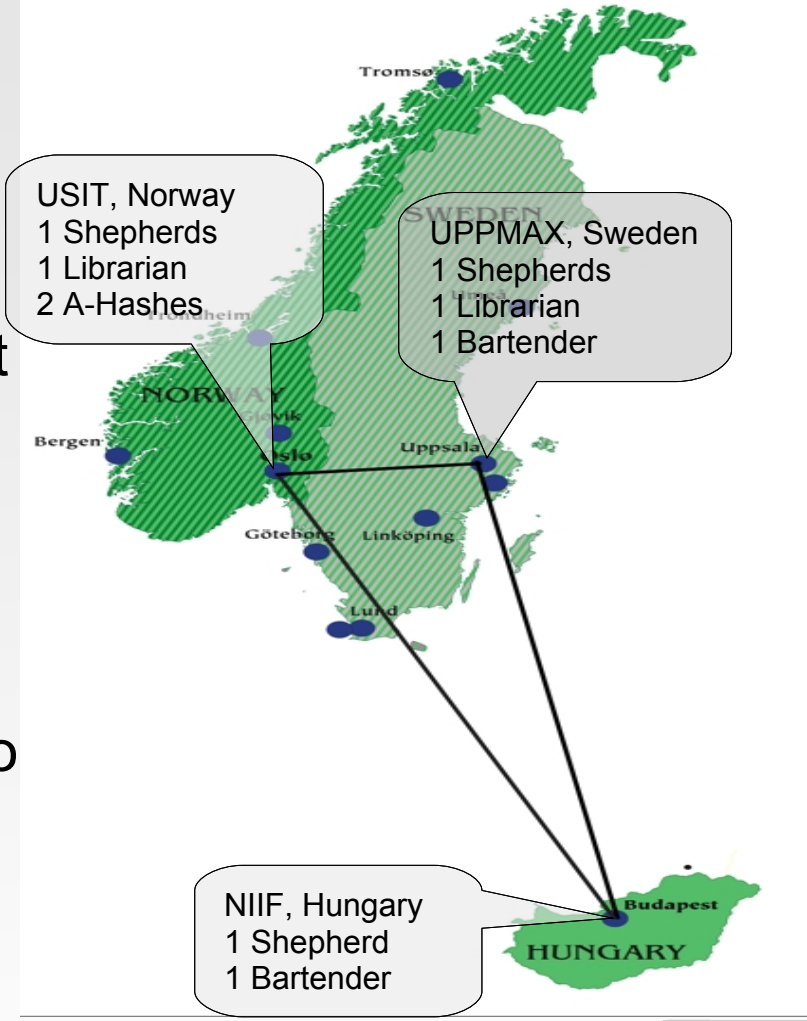

Results

KnowARC

- Upload response time.

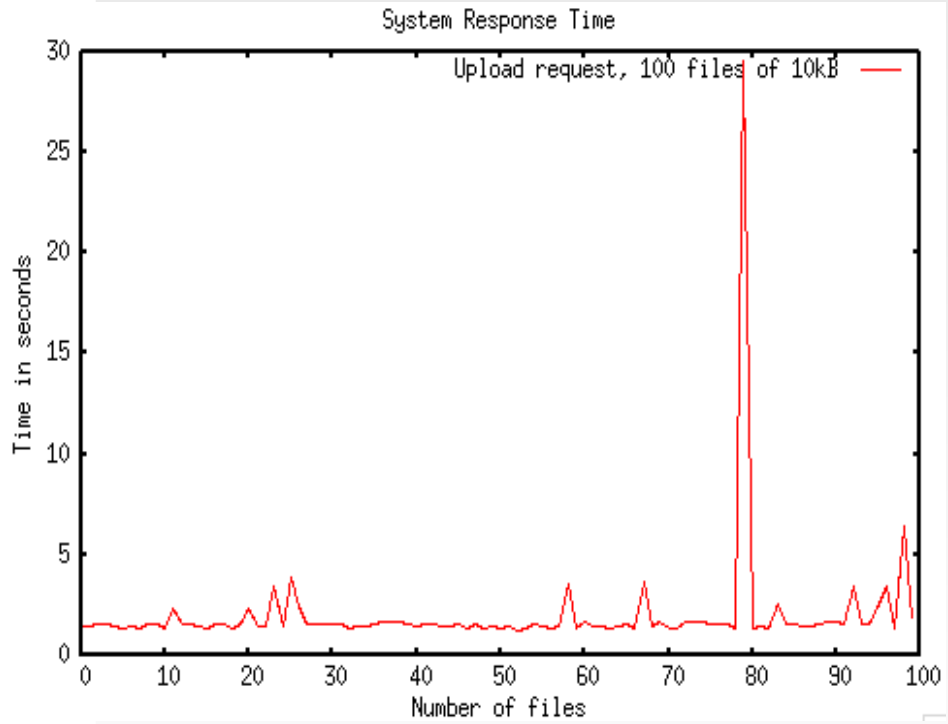

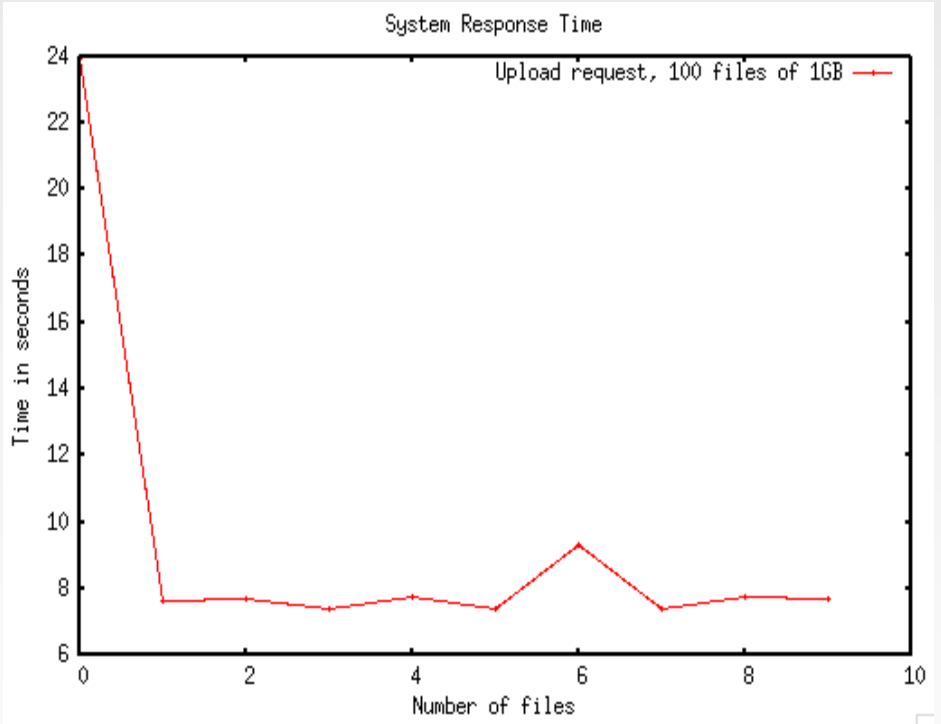

Results

- Distribution of data.

	NIIF	USIT	UPPMAX	Total
Load (GB)	2.79711	2.79718	3.72924	9.32353

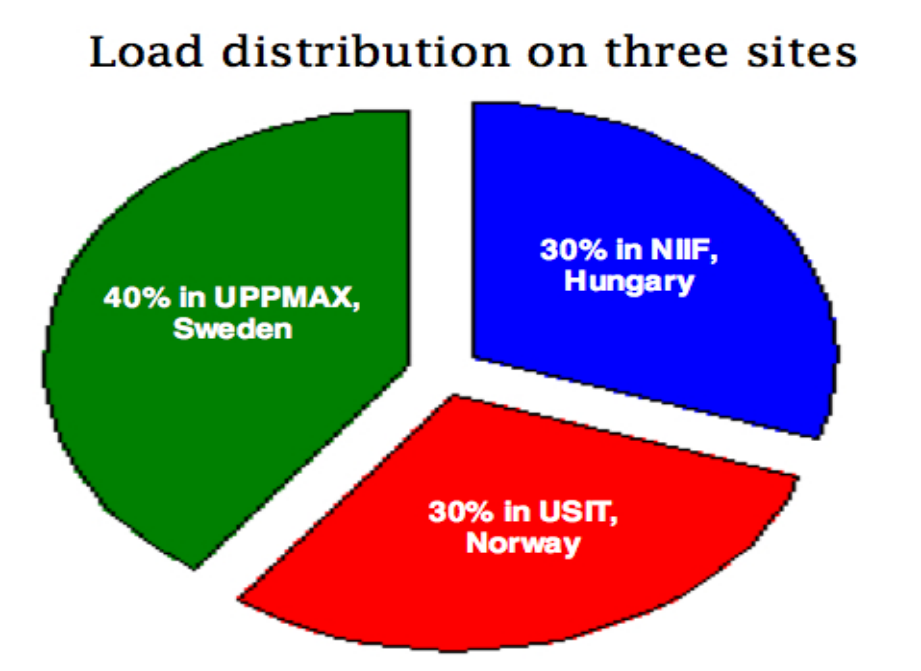

Results

- Increase the number of replicas to 3.
- USIT Librarian and UPPMAX Shepherd is offline.
- Download 100 small files.

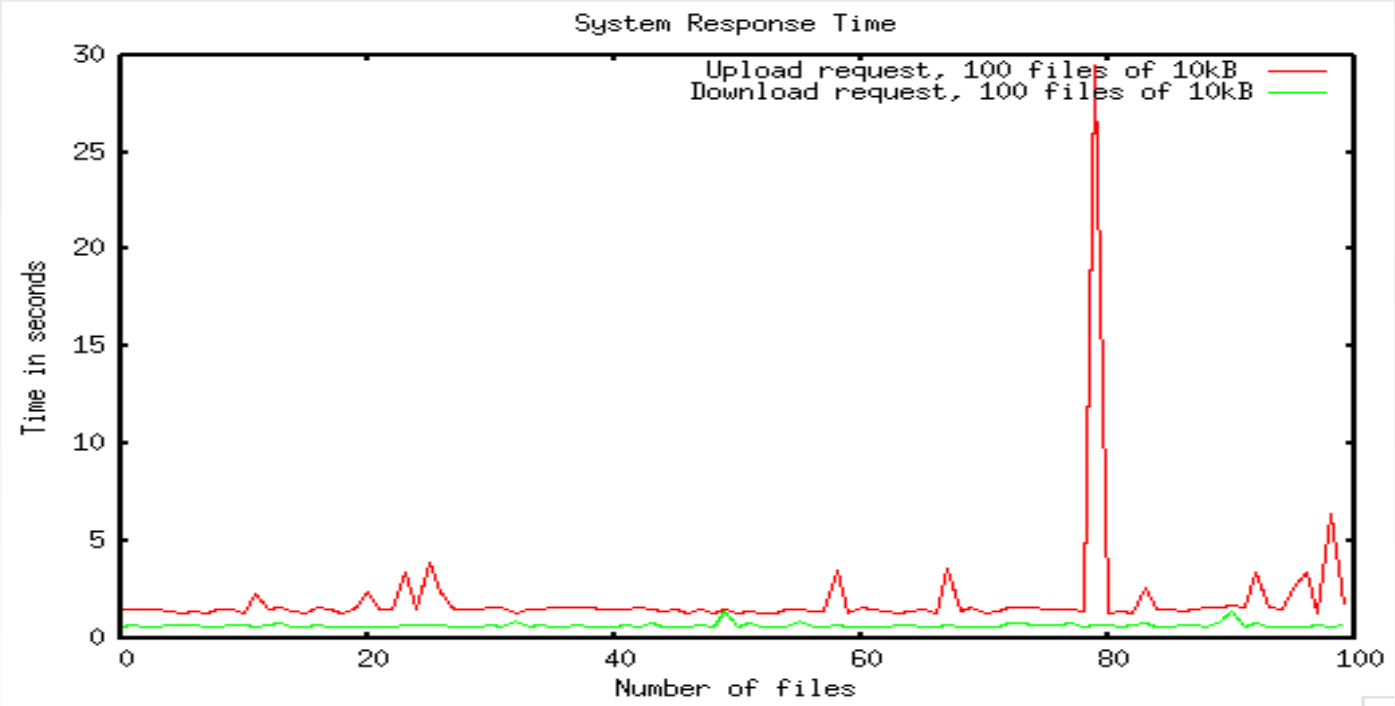

Comparison with other systems

KnowARC

Capabilities\Storage	dCache	DPM	Scalla	iRODS	Hadoop	XtreemFS	Amazon S3	Chelonia
Grid enabled (X.509)	Yes	Yes	No	Yes	No	No	No	Yes
Distributed metadata	No	No	No	No	No	No	N/A	Replicated
Can store on other Ses	No	No	No	Yes	No	No	No	Yes
Global namespace	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Access to other	No	No	No	Yes	No	No	No	Yes
Locally mountable	No	No	Yes	No	Yes	Yes	Yes	Yes
SRM	Yes	Yes	3rd party	No	No	No	No	No
Own storage	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes

Needs checking

Conclusion and Important links

KnowARC

Conclusion

Chelonia is a grid-enabled, self-healing storage cloud. The extensive stability and performance testing is currently in process. But the flexibility in the architecture and the features provided by Chelonia, make it a comparable solution with the other production level systems.

Important links

- Chelonia http://www.knowarc.eu/chelonia.html
- NorduGrid http://www.nordugrid.org/
- KnowARC project http://www.knowarc.eu/
- NGIn project http://www.nordugrid.org/ngin/
- Chelonia Demo http://www.youtube.com/watch?v=NEUWzGHHGhc