

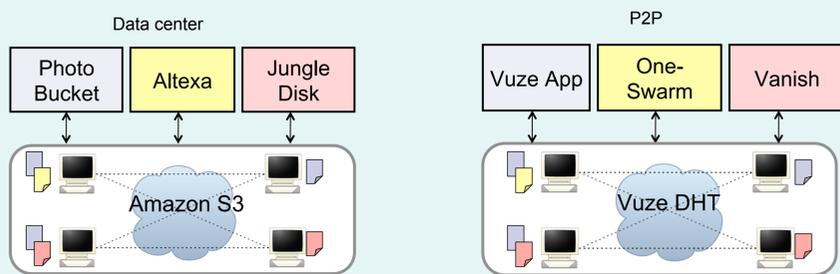
Comet: An Active Distributed Key-Value Store



Amit Levy Roxana Geambasu Tadayoshi Kohno Arvind Krishnamurthy Hank Levy

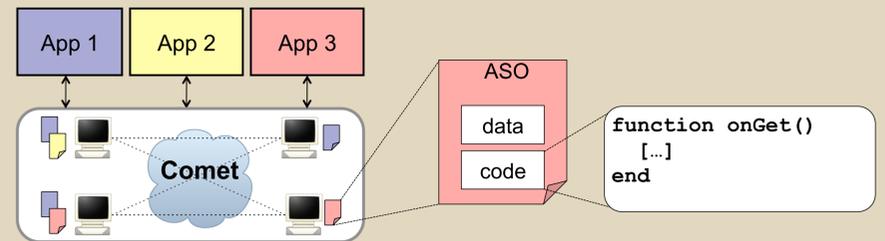
1 Today's Distributed Key-Value Storage Services

- Increasingly, key/value stores are shared by many apps
 - Avoids per-app storage system deployment
- Applications have different (even conflicting) needs:
 - Availability, security, performance, functionality
- But today's key/value stores are *one-size-fits-all*



3 Active Storage Objects (ASOs)

- The ASO encapsulates data + code
 - The data is the actual value
 - Code is a set of handlers that are called on DHT actions:
 - get/put/delete/timer



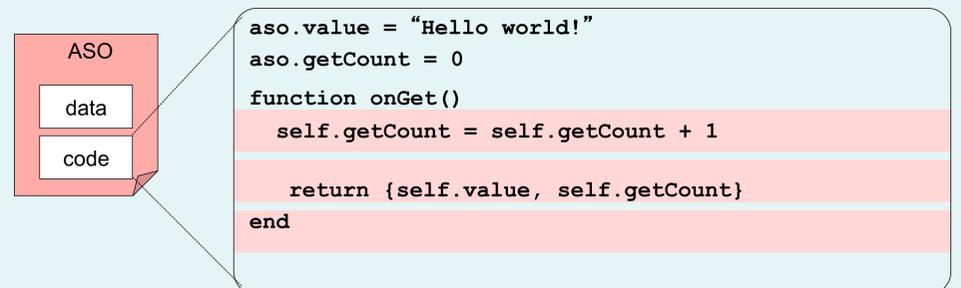
2 The ASO Extension API

- Small yet powerful API for a wide variety of applications
- We have explicitly chosen not to support:
 - I/O, arbitrary network, customized routing, etc...

Intercept accesses	Periodic Tasks	Host Interaction	DHT Interaction
onPut(caller)	onTimer()	getSystemTime()	get(key, nodes)
onGet(caller)		getNodeIP()	put(key, data, nodes)
onUpdate(caller)		getNodeID()	lookup(key)
		getASOKey()	
		deleteSelf()	

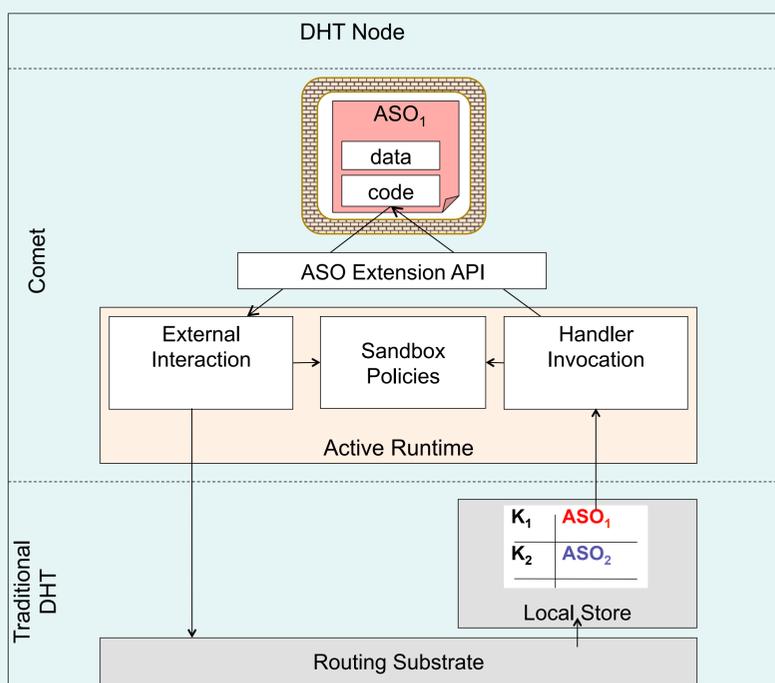
4 A Simple ASO Example

- Each replica keeps track of a gets on an object:



- Simple but powerful:
 - Difficult to track object popularity in current DHTs
 - Trivial to do so in Comet

5 Comet Architecture



- Extend basic DHT architecture with active layers
- Active Runtime handles isolation and security properties
- ASOs run within a language sandbox (Lua based)

6 Comet Applications

Applications	Customization	Lines of Code
Vanish	Security-enhanced replication	41
	Flexible timeout	15
	One-time values	15
Adeona	Password-based access	11
	Access logging	22
P2P File Sharing	Smart Bittorrent tracker	43
P2P Twitter	Recursive gets*	9
	Publish/subscribe	14
Measurement	Hierarchical pub/sub*	20
	DHT-internal node lifetimes	41
	Replica monitoring	21

7 Comet Prototype

- We built Comet on top of the popular Vuze DHT and Lua
 - Experimental deployment on PlanetLab
- In the future we hope to deploy Comet at a large scale
 - Vuze engineer is interested in Comet for debugging and experimentation purposes