

# Khazana: A Flexible Wide Area Data Store

*Sai Susarla and John Carter*  
{*sai, retrac*}@cs.utah.edu

UUCS-03-020

School of Computing  
University of Utah  
Salt Lake City, UT 84112 USA

October 13, 2003

## *Abstract*

Khazana is a peer-to-peer data service that supports efficient sharing and aggressive caching of mutable data across the wide area while giving clients significant control over replica divergence. Previous work on wide-area replicated services focussed on at most two of the following three properties: aggressive replication, customizable consistency, and generality. In contrast, Khazana provides scalable support for large numbers of replicas while giving applications considerable flexibility in trading off consistency for availability and performance. Its flexibility enables applications to effectively exploit inherent data locality while meeting consistency needs. Khazana exports a file system-like interface with a small set of consistency controls which can be combined to yield a broad spectrum of consistency flavors ranging from strong consistency to best-effort eventual consistency. Khazana servers form failure-resilient dynamic replica hierarchies to manage replicas across variable quality network links. In this report, we outline Khazana's design and show how its flexibility enables three diverse network services built on top of it to meet their individual consistency and performance needs: (i) a wide-area replicated file system that supports serializable writes as well as traditional file sharing across wide area, (ii) an enterprise data service that exploits locality by caching enterprise data closer to end-users while ensuring strong consistency for data integrity, and (iii) a replicated database that reaps order of magnitude gains in throughput by relaxing consistency.