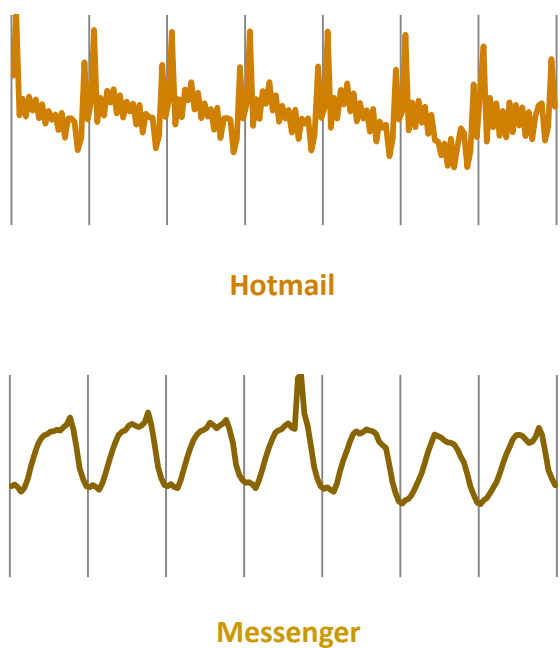


Sierra: practical power-proportionality for data center storage

Eno Thereska, Austin Donnelly, Dushyanth Narayanan

1. motivation




We observe that ...
Many workloads have peak and trough periods
Idle servers consume significant power

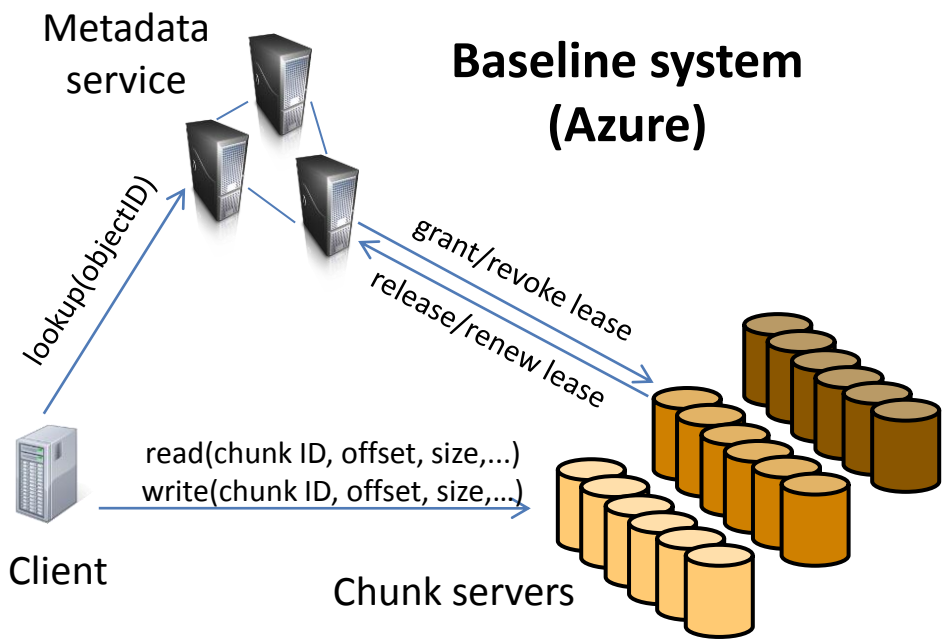
We would like to. ...
Turn some servers off during the troughs
Consolidate load on remaining servers
Maintain availability, consistency, performance, fault-tolerance

Storage is the challenge: cannot migrate

What Sierra does ...
Turn some data replicas off at low load
Keep 1 replica available for reads
Send writes to *distributed virtual log*



2. overview



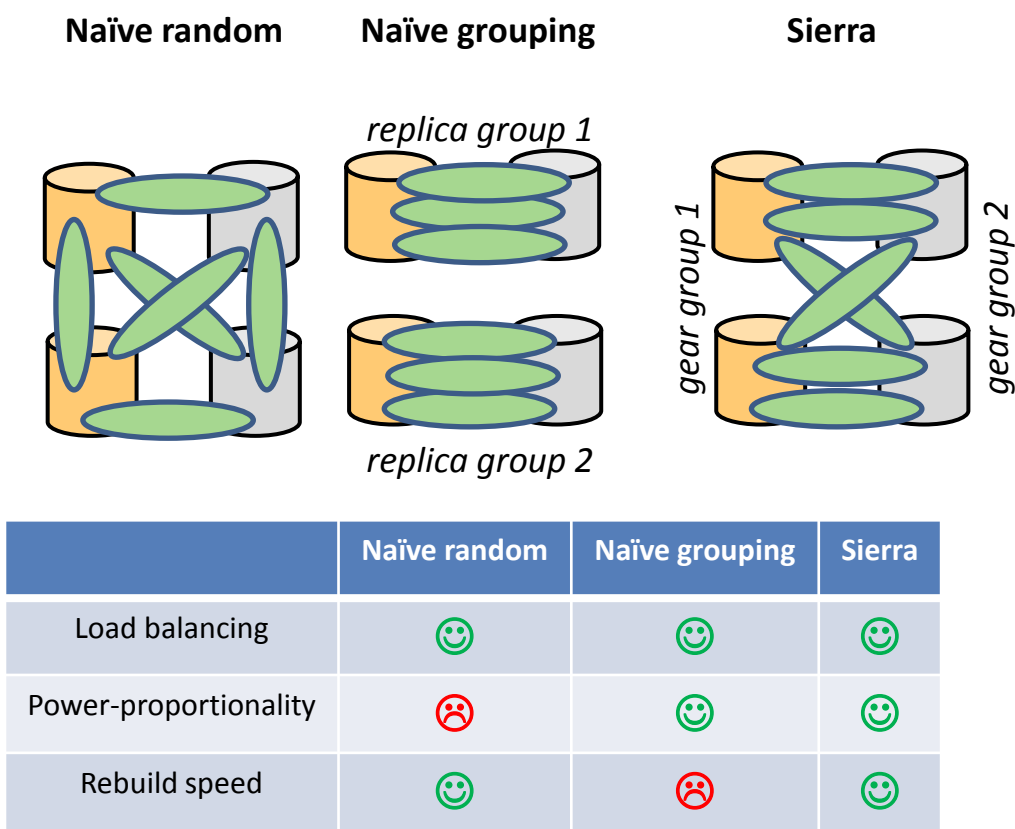
Challenge: add power-proportionality without losing

Read availability
Write availability
Consistency
Fault tolerance
Rebuild speed

Sierra design features

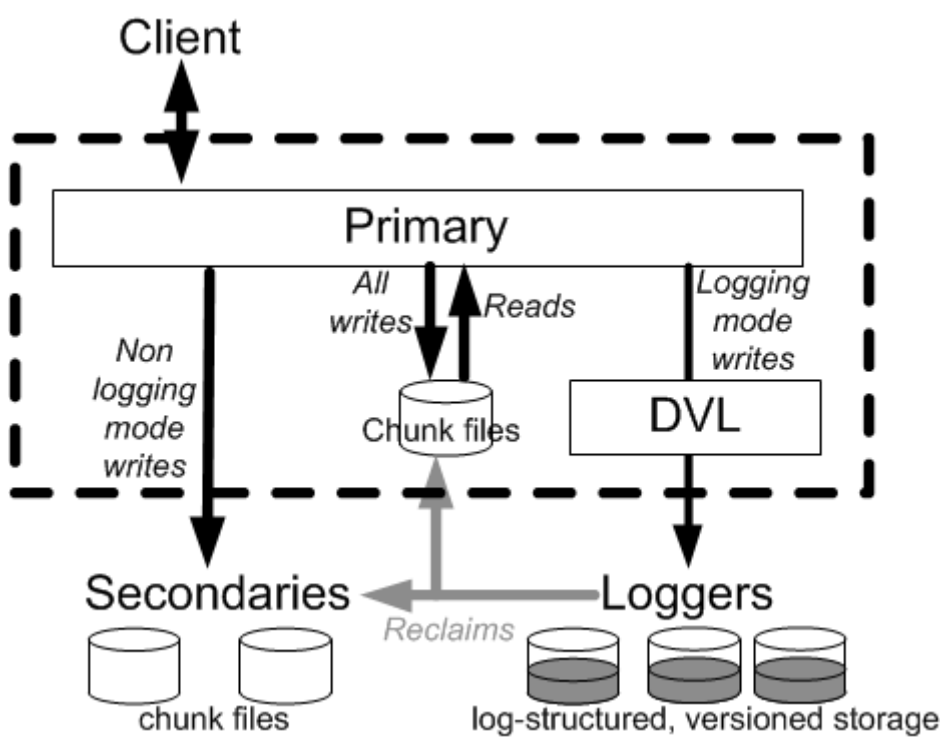
Replica placement
Distributed virtual log
Lease migration protocols
Diurnal gear scheduler

3. replica placement



4.

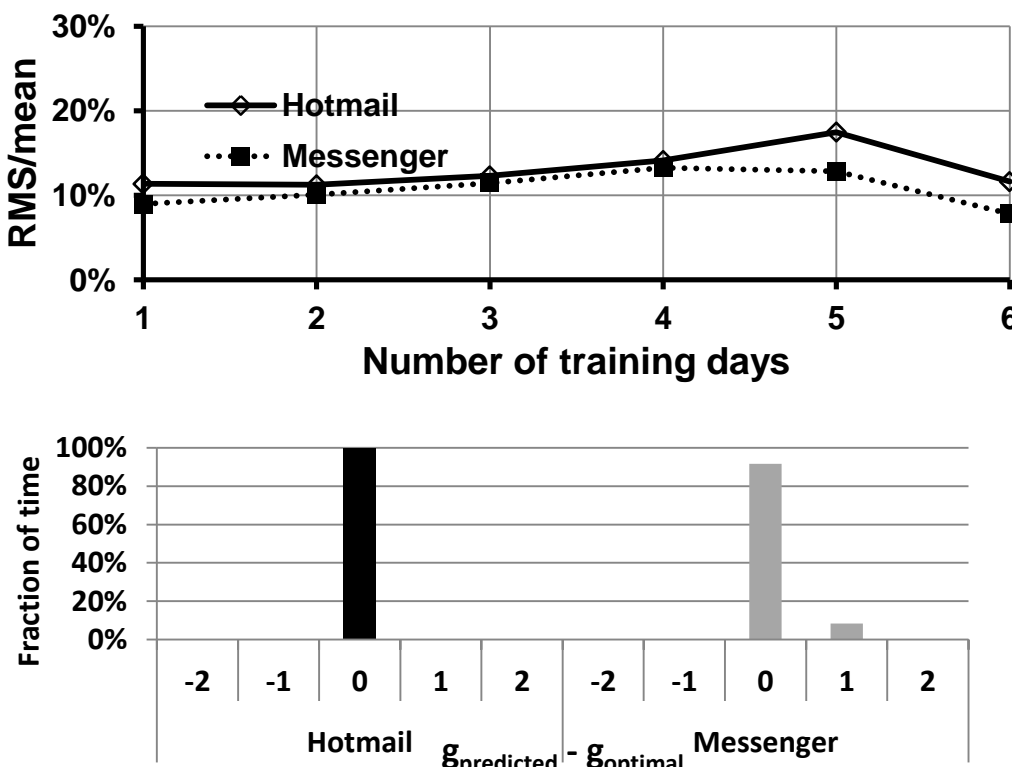
distributed
Virtual log




5.

evaluation

Load patterns are predictable




Power savings are significant



based on 1-week load traces

Performance impact is modest



based on 48-hour Hotmail disk I/O traces