scc: Cluster Storage Provisioning Informed by Application Characteristics and SLAs

UC RIVERSIDE UNIVERSITY OF CALIFORNIA Bourns College of Engineering

Harsha V. Madhyastha*, John C. McCullough, George Porter, Rishi Kapoor, Stefan Savage, Alex C. Snoeren, and Amin Vahdat

UC Riverside and UC San Diego*

UCSDCSE Computer Science and Engineering

Cluster Storage Provisioning

- Diverse requirements in cluster applications:
  - Sequential throughput for bulk data analysis
  - IOPS for key-value storage
  - Mix of the two in other scenarios
- Currently, requirements roughly characterized
- Configurations typically selected by rule of thumb, often homogeneous

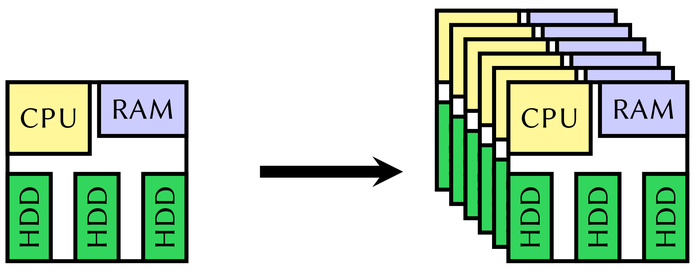

- Homogeneous configurations result in hard-to-utilize resources and unnecessary increase in cost

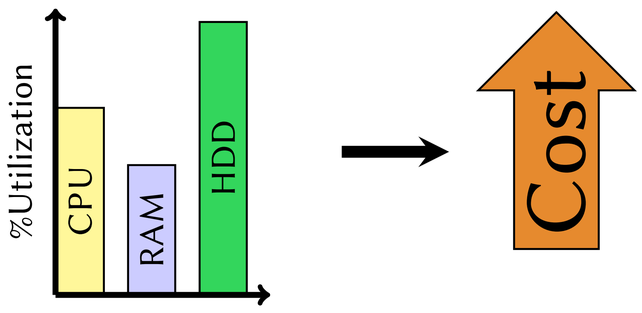

- We provide a storage configuration compiler to find low-cost, performant configurations for a given application

Challenge

- Diverse storage options:

	size	MB/s	IOPS	Cost
7.2k-rpm	500GB	90/90	125/125	$213
15k-rpm	146GB	150/150	285/285	$296
SSD	32GB	250/80	2500/1000	$296
DRAM	1GB	13K/13K	1.6B/1.6B	$35

- Best storage choice:
  - Coupled to workload
  - May change with growth
- Storage choice constrained to chassis capacity and bus capacity:

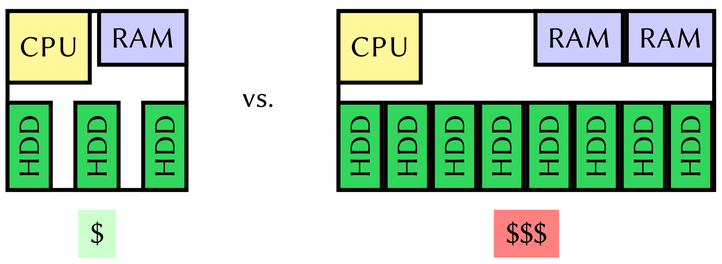

- Need CPU to drive storage and application
- Must understand relationship between application, workload, and storage

Overview of scc

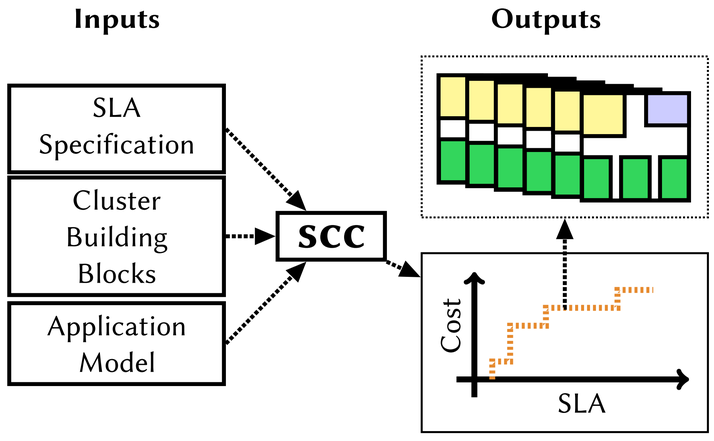

- Storage Configuration Compiler
- Automates cluster configuration
- Combine:
  - Formal application specification
  - Hardware properties
  - Workload specification

Specifying Apps

- Application broken into tasks and datasets

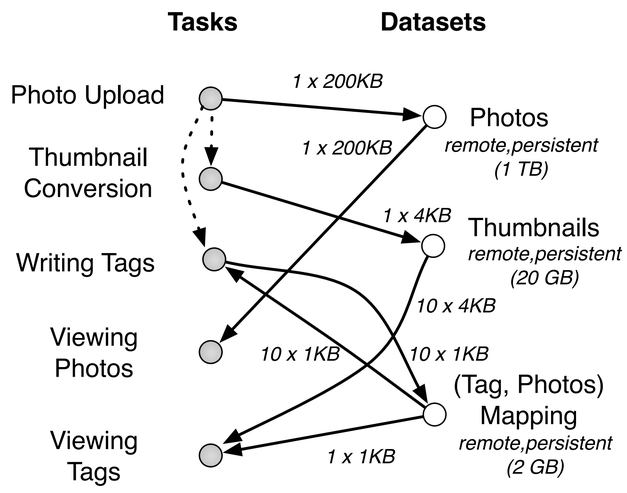

Model for photo sharing application

- Tasks parametrized by cpu-time, write count and size for each dataset, and dependencies
- Datasets parametrized with overall size
- SLA is throughput of tasks per unit-time

Algorithm

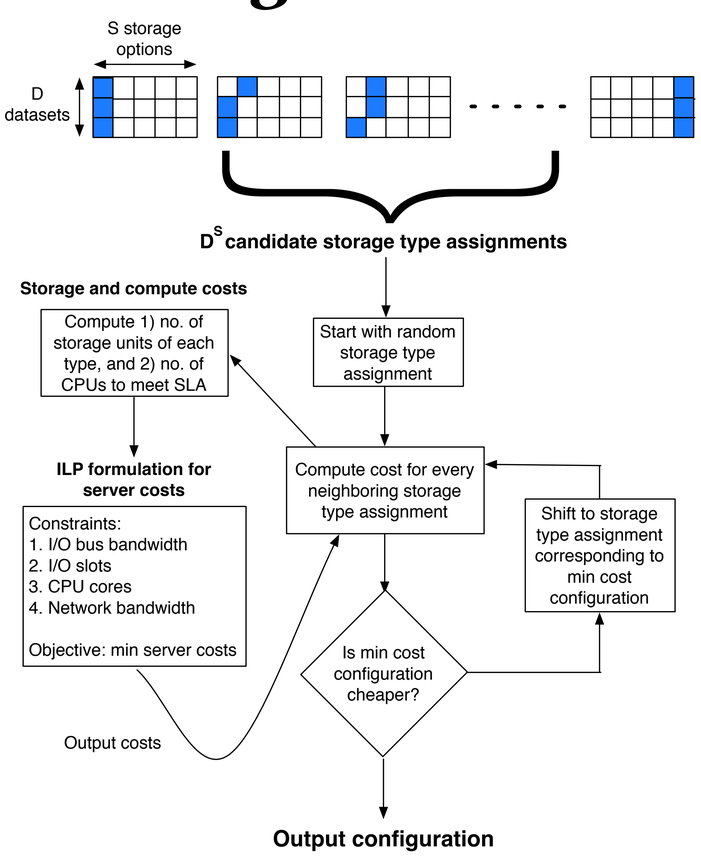

Heterogeneous beats scale-out

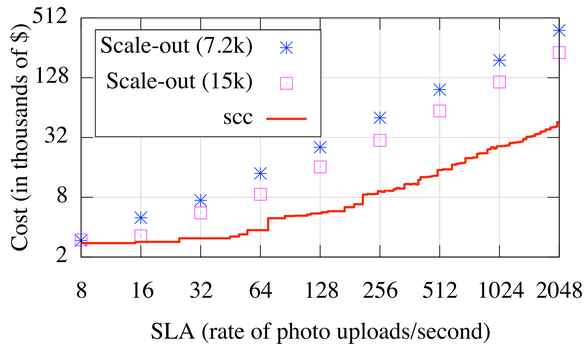

scc achieves $2 - 4.5\times$ cost average savings vs. scaling fixed server configurations

Storage regimes

Uploads/s	Photos	Thumbnails	Tags
≤ 5	Disk	Disk	Disk
$5 - 25$	Disk	Disk	Disk+ DRAM
$25 - 330$	Disk	SSD	Disk+ DRAM
$330 - 930$	SSD	Disk+ DRAM	Disk+ DRAM
$930 - 10k$	Disk+ DRAM	Disk+ DRAM	Disk+ DRAM

Storage unit type

Storage-types transitions with workload

Modeling Storage

- Parametrized by:
  - Capacity
  - Sequential throughput
  - Operation *Gap* (e.g. seeks, erasures)
- *Gap* captures more than IOPS
- Operation latency $= \frac{\text{size}}{\text{rate}} + \text{gap}$

Validation

Cheaper configurations miss the SLA, and simple scale-out solution requires much higher cost

Robustness

Attribute	Range with same architecture
Photo size	50 KB ← 200 KB → 850 KB
Thumbnail size	1 KB ← 4 KB → 30 KB
SSD unit price	$200 ← $450 → $900

scc can estimate the robustness of its output