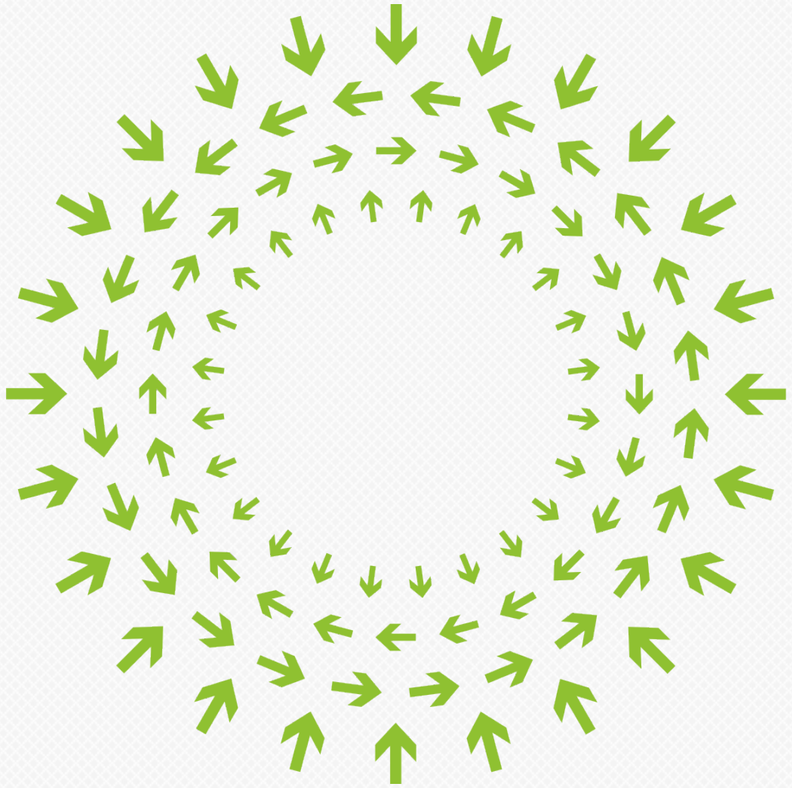

OPEN Compute Summit

January 28–29, 2014San Jose

OPEN Compute Summit

Rack Storage for Cloud Applications

Sharing experimental results

Manoj Wadekar
QLogic Corp.
Chief Technologist

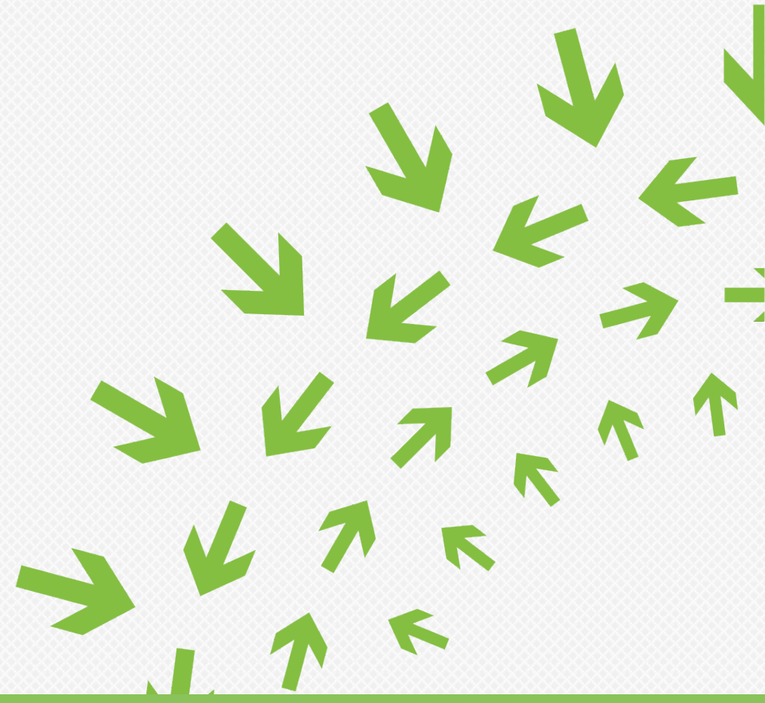

Disclaimer

- **Sharing some early results of tinkering with Hadoop workload**
  - Findings may apply to other similar workloads
- **Feedback welcome**

Agenda

- **Problem**
- **Potential Solution**
- **Hadoop with Rack Storage**
- **Early results**

Exploration Area

- **Distributed Storage:**
  - Scale-out Storage: Clustered DAS
  - Distributed compute and Storage: Shared-Nothing
  - Converged infrastructure: Compute-Storage Platforms
  - Etc.
- **What are challenges?**
  - Performance
  - Deployment
  - Maintenance
  - Etc.
- **Hadoop: an example application for study**

Hadoop – Storage Dependency

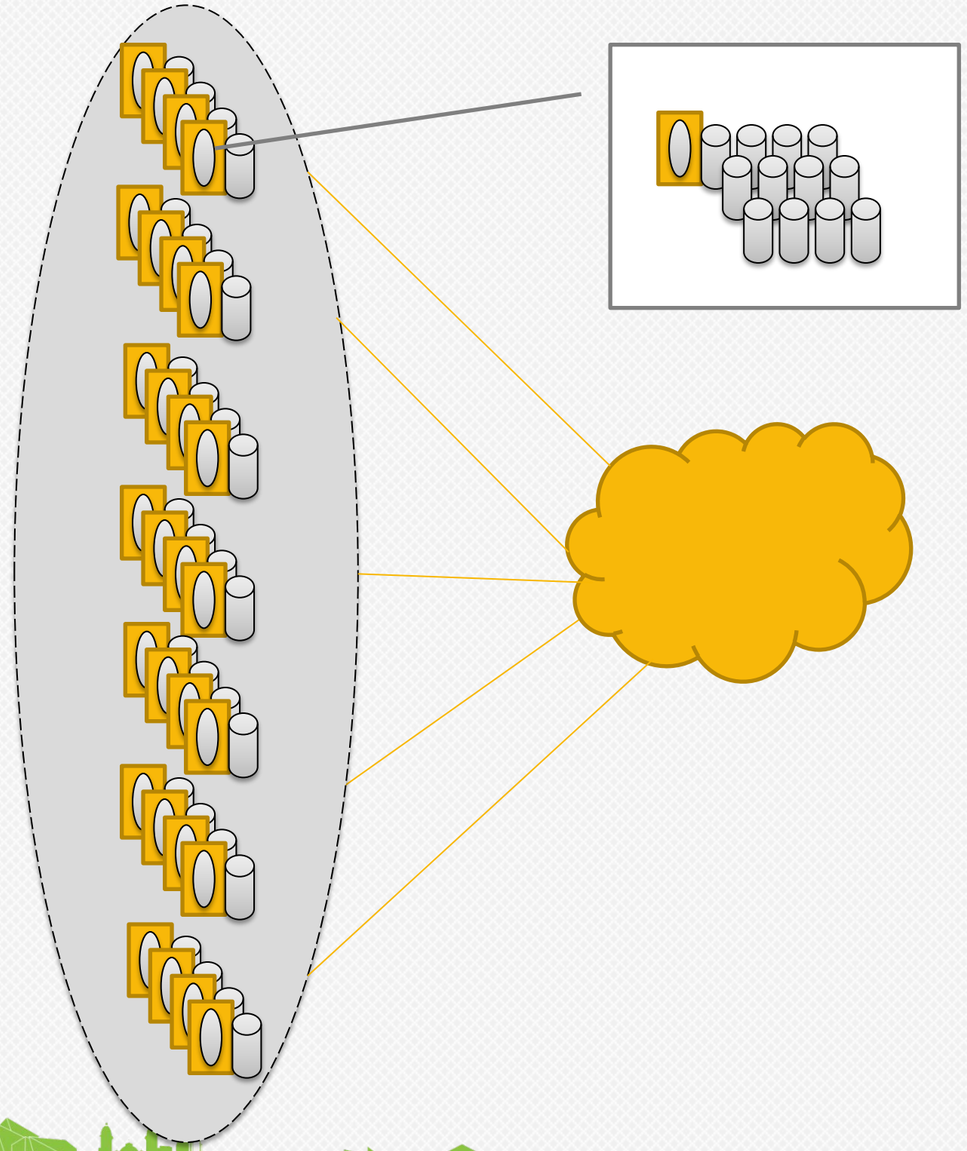

- Platform Variables:
  - CPU, Memory, Networking, Storage
- Small test bed: 5 nodes, Map-Reduce
- (No news here) Storage is most important variable
  - More disks – better performance
  - For Ingest as well as Map-Reduce
  - All other variables follow after disk bottleneck is removed

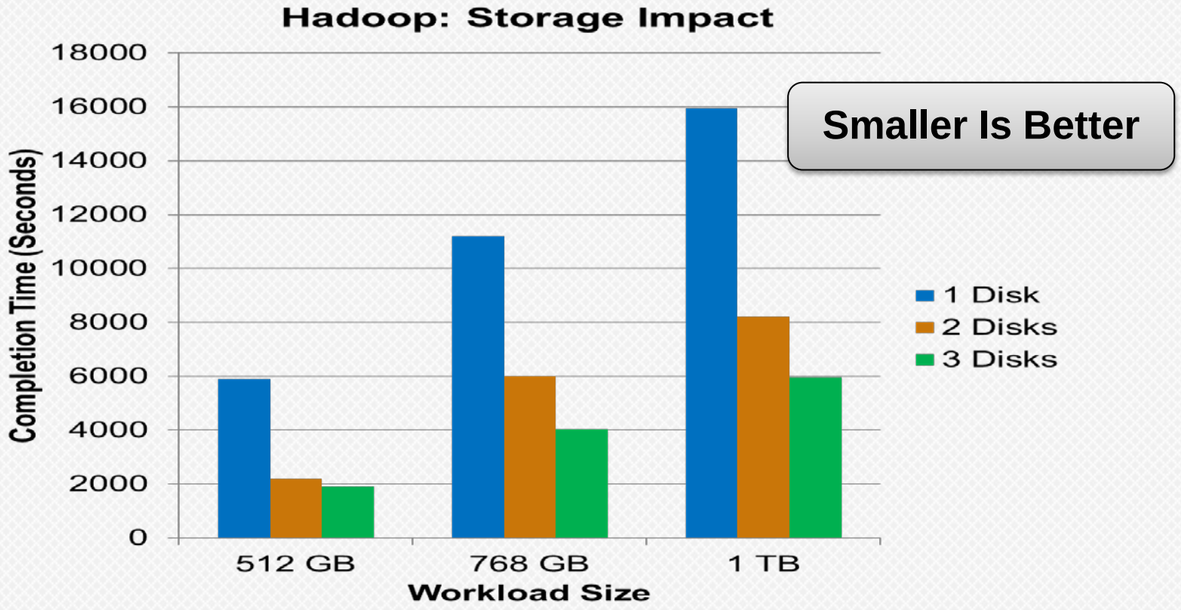

Storage Server: Challenges

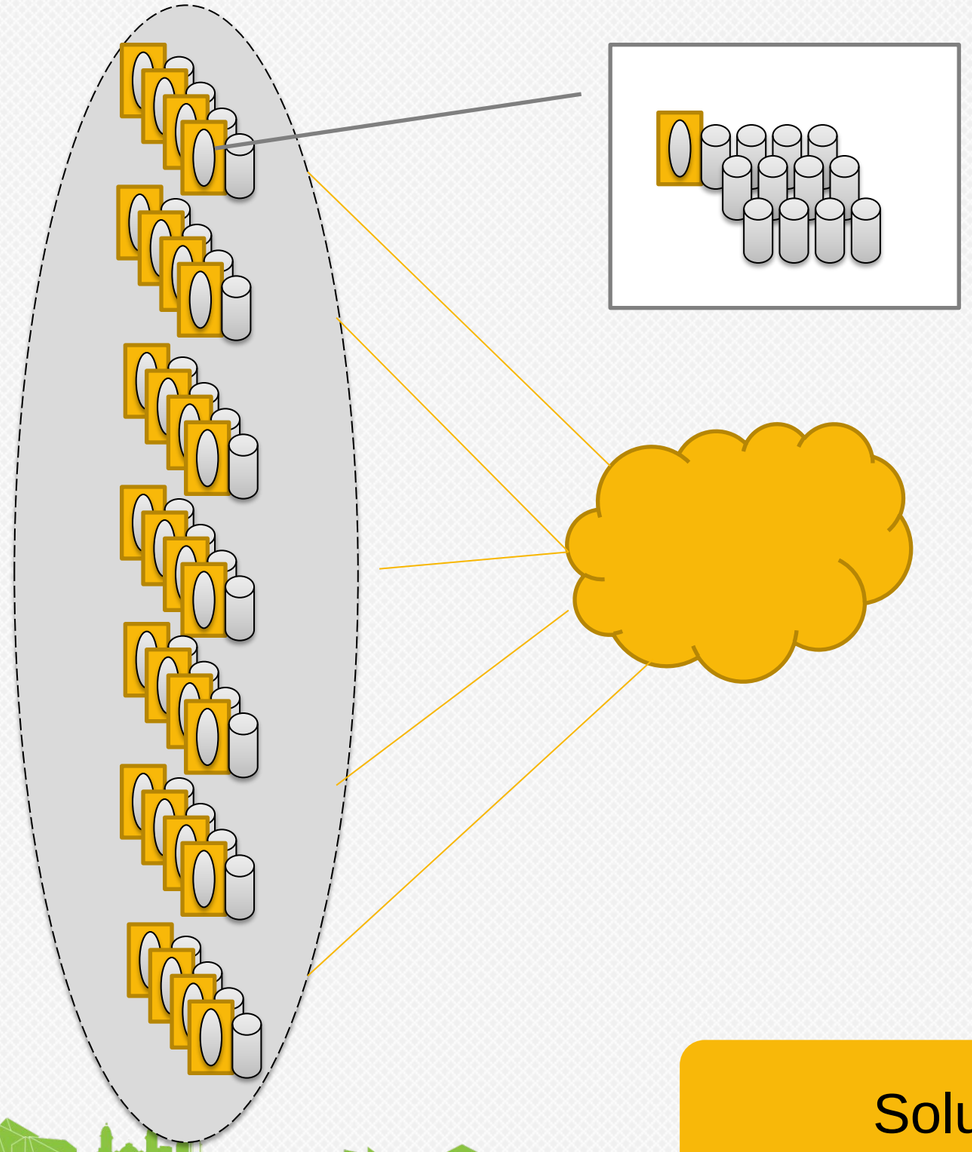

- Capacity / Performance bottleneck
  - How many disks in a server?
  - How to add capacity in a server?
  - Disk Failure maintenance?
- Separate refresh cycle storage and compute
- SKU reduction
  - Servers with SSDs, Fast disks, slow disks?
- Silos

Solution: Disaggregated Compute/Storage

Rack Storage

- Rack Storage: Local Pool
  - Without changing Hadoop-view of the world: shared-nothing
- Disaggregated Compute and Storage
  - Allow independent refresh
- Improve TCO
  - Optimal compute and storage solution
- Other potential opportunities
  - Storage specific enhancements
  - Tiered storage

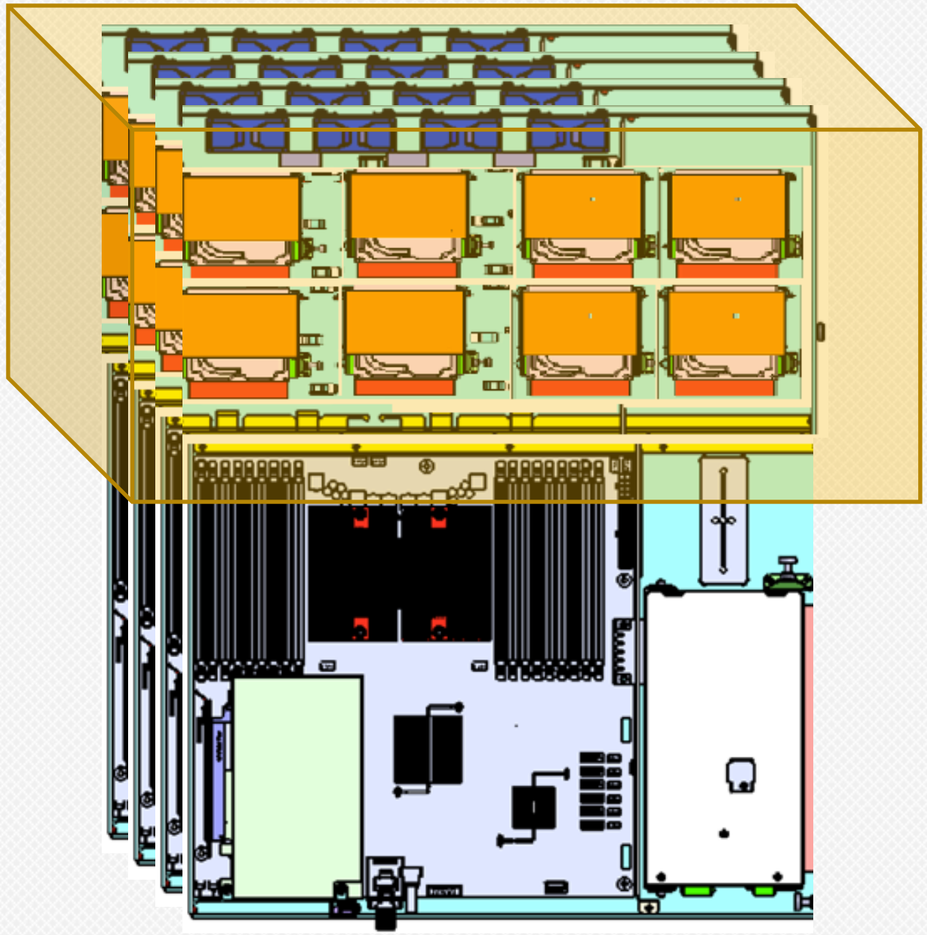

Question

- **Does moving local storage into the network reduce Hadoop performance?**
- **Assumption:**
  - Keep “Shared Nothing” concept from Hadoop perspective
  - Pool storage for manageability, and perhaps reliability and performance

Hadoop Experiment –Setup

Local Storage

- 5 Data Nodes:
  - Xeon Servers
    - 1 Boot Disk
    - 1 Local Disk for Hadoop
- 1 Name Node

iSCSI Storage

- 5 Data Nodes:
  - Xeon Servers
    - 1 Boot Disk
    - 1 Network SAN LUN for Hadoop
      - (Could be FC or FCoE as well)
- 1 Storage Server: iSCSI
  - 5 Disks
- 1 Name Node

- Benchmark: Terasort
- Workload Size: 512G, 768G, 1TB

Hadoop Disk Usage Scenario – Rack Storage

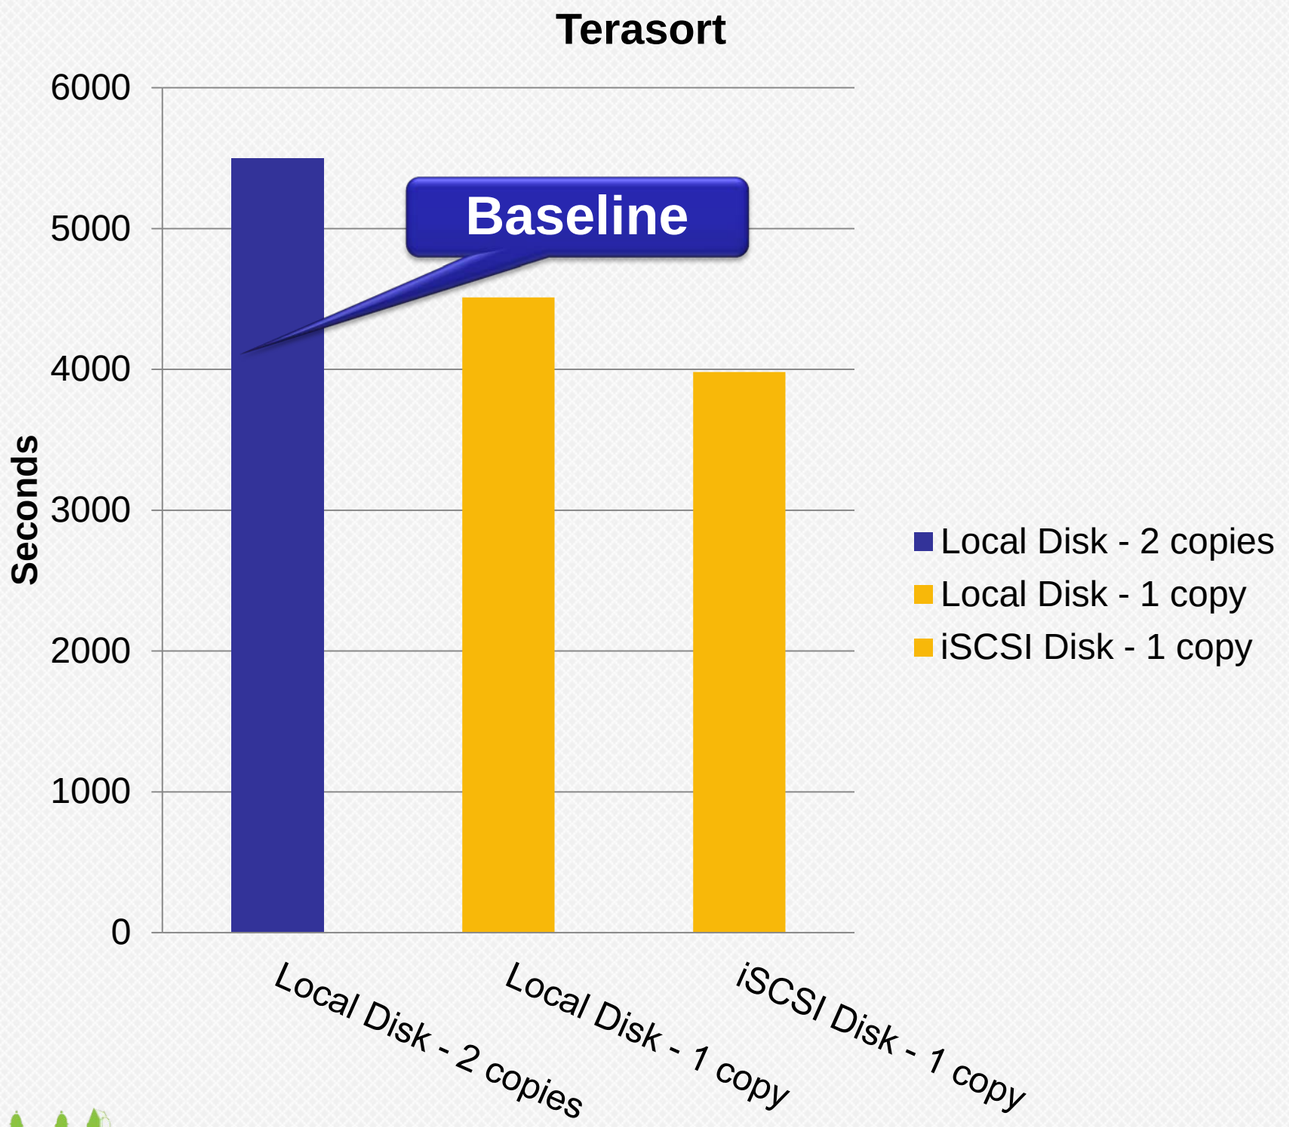

- Benchmark: Terasort
  - 256GB workload
- Hadoop data local or remote
- Remote storage: iSCSI
  - Server with 5 disks
- Test Cases:
  - Local Disk: 2 Copies
  - Local Disk: 1 Copy
  - iSCSI Disk: 1 copy

Summary

- **Rack Storage: Enables Disaggregated Compute-Storage**
- **Early results: Rack Storage (over the network) Hadoop performance at par with local disk**
- **Areas for further investigation**
  - Shared storage with more reliability, performance, features
  - More (Cost/power/performance) efficient processors in storage system
- **Optimal system for shared storage: HoneyBadger?**

Questions?

QLogic is at Yttibrium and Quanta booths.

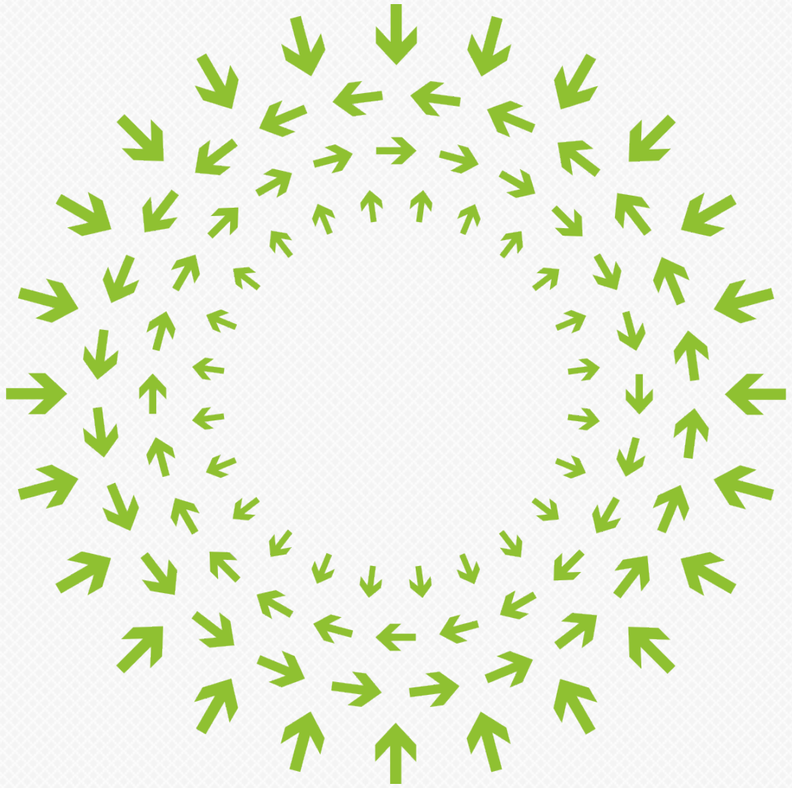

OPEN Compute Summit

January 16-17, 2013Santa Clara