What might a Cloud Computing approach mean for CDR?

John D. Evans

Global Science & Technology, Inc. (GST)
Greenbelt, MD

Elasticity & CDR needs

- Primary driver: data volumes
- Secondary driver: data processing loads
- Real driver: Uncertainty on processing requirements
  - System Architecture focused on flexibility and scalability
- Scheduler/Resource Manager
  - Allow system to self-manage as a quasi-batch system
    - Variable jobs loads over time
    - Optimize resource utilization balanced against throughput requirements

Excerpts from: NPP/NPOESS CDR/CIR PRODUCTION SYSTEM ARCHITECTURE OVERVIEW (Derek Yu, Scientific Research Corporation, June 2010)

CDR, CLASS, and bandwidth

- Goal of reprocessing one year of RDRs/Climate RDRs in one day means
  - 600GiB daily RDR volume total is ~214TiB/Year
  - Transferring this volume from CLASS in one day requires ~20Gbps of network bandwidth
    - 1 Gbit Ethernet (~125MB/sec) - ~499 hrs (~21 days)
    - 8G Fiber Channel (~850MB/sec) - ~73hrs (~3 days)
    - 10 Gbit Ethernet (~1.2GB/sec) - ~51 hrs (~2 days)
    - 12X QDR Infiniband (~11GB/sec) - ~6 hrs

Excerpts from: NPP/NPOESS CDR/CIR PRODUCTION SYSTEM ARCHITECTURE OVERVIEW (Derek Yu, Scientific Research Corporation, June 2010)

CDR, CLASS, and bandwidth



Excerpts from: NPP/NPOESS CDR/CIR PRODUCTION SYSTEM ARCHITECTURE OVERVIEW (Derek Yu, Scientific Research Corporation, June 2010)

Scaling out: Data-Intensive Science

1. Scientific computing is becoming increasingly data intensive
   - Storage bandwidth hasn't kept pace with capacity
2. The solution is in a “scale-out” architecture
   - Map / Reduce computing patterns
   - Distributed, non-relational databases (“NoSQL”)
   - “Storage Bricks” / “Amdahl blades” *(Alex Szalay)* (with balanced IO vs. CPU)
3. Bring computations to the data
   - “Eliminate gratuitous data movement”
   - cf. Apache Hadoop: “rack-aware” computing
4. Start with “20 queries”
   - As a basis for indexing, system design & configuration
5. Go from “working to working”
   - Build & test early & often

Photo by flickr.com/dullhunk via Creative Commons Attribution 2.0 Generic (CC BY 2.0)



Jim Gray

cf. “The Fourth Paradigm: Data-Intensive Scientific Discovery”

Scaling out: MapReduce pattern



amazon web services™

Learning from others

- Learning from other petabyte-scale science projects is important
  - CERN LHC Compact Muon Solenoid (1 MB @ 40MHz)
  - DARPA ARGUS-IS *(1.8 gigapixel @ 15 frames/sec)*
  - Sloan Digital Sky Survey, …
- But "Science is not the big kid anymore" *(Jim Hendler)*
  - Facebook: logs 25 TB/day
  - Google: processes 24PB/day
- Both offer lessons, resources
  - Computation patterns
  - Datacenter & server design
  - Distribution and interoperability
  - Services, tools, APIs
  - Data linking and semantics

Photo by Robert Scoble / http://scobleizer.com

Facebook's new data center – Prineville, OR

Share Tweet

Crowd-sourced science via Web 2.0

Galaxy Zoo and Hanny's Voorwerp



Hanny Van Arkel

In summary

- Elastic computing may suit CDR well
- Bandwidth may be a showstopper for traditional system design (archive-network-processing)
- Data-intensive science: “Scaling out”
  - Massive parallelization of data and processing
  - Eliminate gratuitous data movement
  - Look at other petabyte-scale projects (science, Web)
- Getting started on the cloud is easy; but making the most of it may require rethinking things
  - Approaches to storage and computation
  - System and network architectures