Georgia Tech Research Institute
Problem. Solved.

Why Hadoop?

Jason Poovey
Branch Head of HPC and Data Analytics

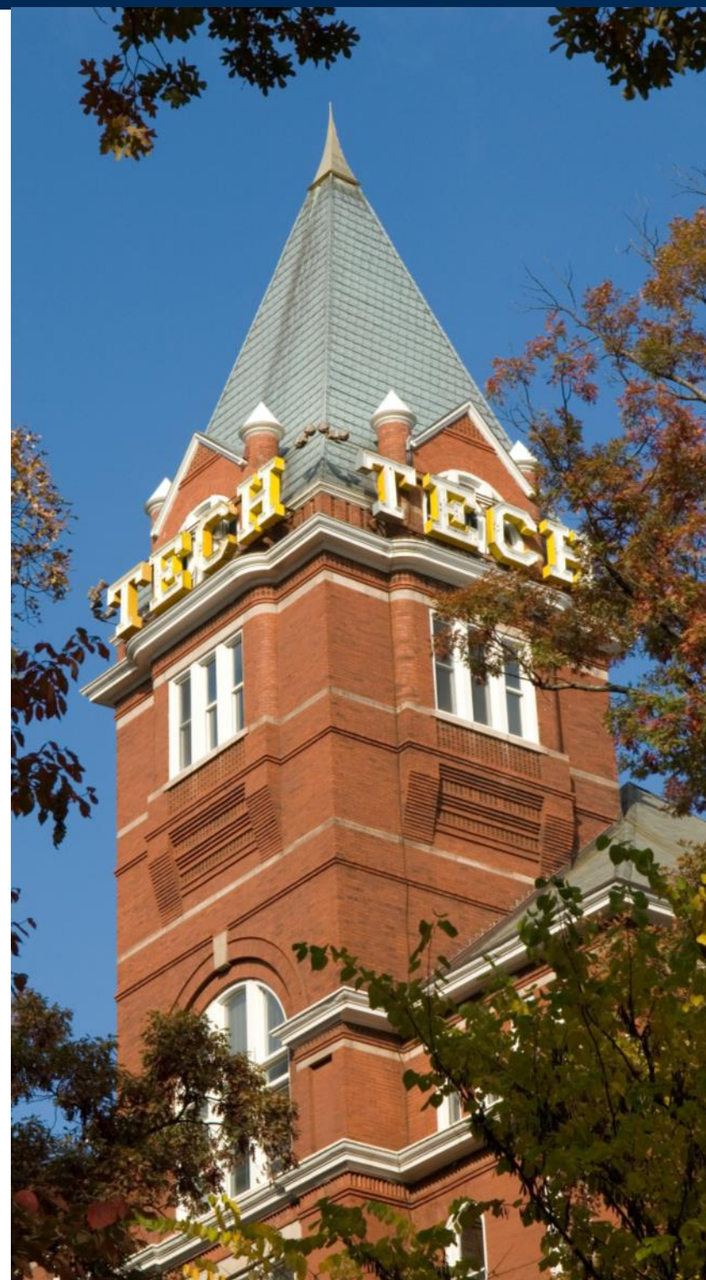

Georgia Tech + OHDSI

- Work with our partners data to
  - Merge together data from many sources: structured, unstructured, omics, sensor/IoT
  - Harmonize and Standardize across health domains and sources
  - Build computable and shareable phenotypes that describe cohorts uniformly across multi-modal data
  - Build statistical models for: precision medicine, care recommendation, clinical decision support
  - Provide interoperable reports and dashboards

- OHDSI provides the common model for the predictive pipelines

Center for Health Analytics and Informatics

Using Hadoop in GT

- Machine Learning Pipelines

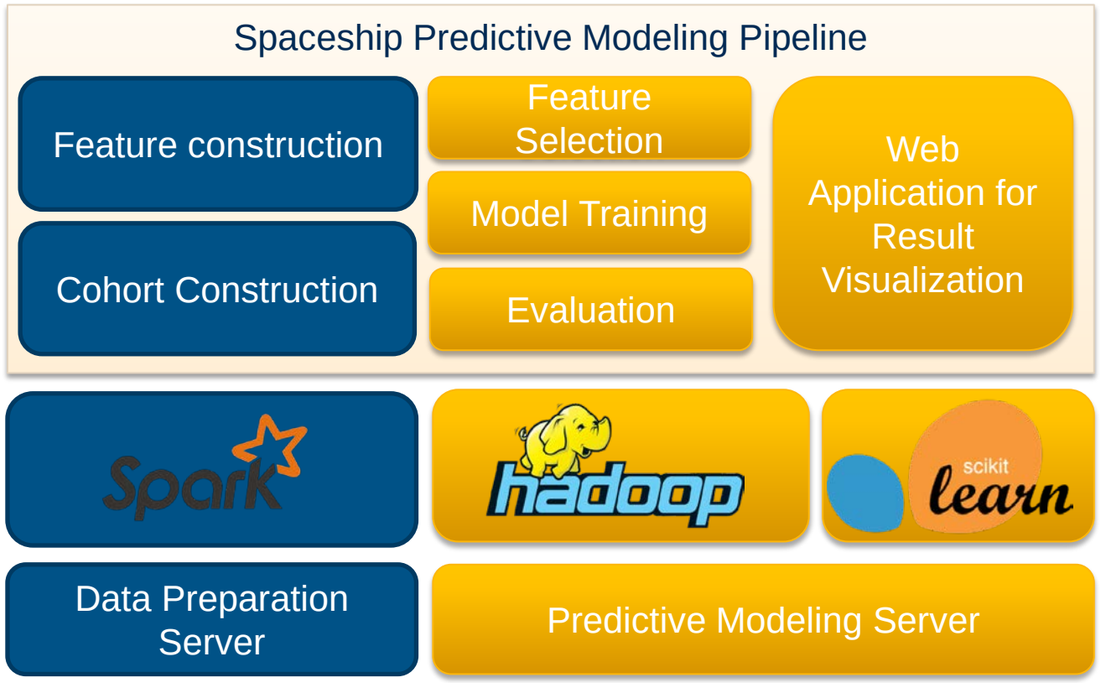

256GB memory
64 processing units

- Storing Sensor Data (H5 files + HIVE)

- NLP / Clinical Notes

cTAKES

Solr

OHDSI
Observational Health Data Sciences and Informatics

Infrastructure Needs

- **Data Storage and Access**
  - GT has invested in a secure, PHI environment to hold several large datasets
  - Current capacity
    - 96TB for ***PHI data***
    - 256TB for ***Hadoop***
    - 20TB for ***Solr Cluster***
  - Expect need for an environment that holds a mixture of PHI and non-PHI
    - ***Clinical Data*** – Standardized Data Format
    - ***Sensor Data*** – Patient Generated Data, IoT/SmartCities Data
    - ***Unstructured Text*** – Clinician Notes, Drug Labels, Publications, etc.
    - ***Omics Data***

- **Computational Resources**
  - ***GPGPU arrays*** – Deep Learning, Large scaled simulation and modeling
  - ***Traditional Compute Arrays / VM clusters*** – Machine Learning, Simulation and Modeling, Databases, etc.
  - ***Hadoop-like architectures*** – Machine Learning, Analysis/Storage of Sensor Data
  - ***Solr/Lucene architectures*** – Text Mining / NLP

GT Hadoop V1 Design Parameters

- Started with 8 nodes to achieve minimum parallelism goals
- Preferred frequency over core count for NameNode and Solr
- Balanced frequency and cour count for DataNodes
  - Considered single socket E3 v5 series for price (could have bought a much larger cluster)
  - All public clusters that we found used dual-socket servers
  - There is a widely-held belief that Hadoop makes good use of multiple cores/sockets
- 4TB SATA drives represent "good value" at today's prices
- Solr performs best on SSDs
  - Can query multi-terabyte indexes in 100s of milliseconds

Hardware Specifications

- (1) Login/Management/NameNode
  - Dual Intel Xeon E5-2623 v4 2.6GHz, 10M Cache, 4C/8T
  - 64GB DDR4 2400MHz DRAM
  - 4x 1TB 7200RPM Hot-plug SATA Hard Drives (RAID-5)
  - Intel X550 Dual Port 10GBASE-T

- (8) DataNodes
  - Dual Intel Xeon E5-2650 v4 2.2GHz, 30M Cache, 12C/24T
  - 64GB DDR4 2400MHz DRAM
  - 8x 4TB 7200RPM Hot-plug SATA Hard Drives (JBOD)
  - Intel X550 Dual Port 10GBASE-T (bonded)
  - Dual Hot-plug Redundant Power Supplies (1+1), 750W

Hardware Specifications

- (5) Solr Nodes
  - Intel Xeon E3-1270 v5 3.6GHz, 8M Cache, 4C/8T
  - 64GB DDR4 2400MHz DRAM
  - 4x 1TB Samsung 850 Pro SSD drives (RAID-5)

- Network: Dell Force10 S4820T
  - 48 x 10GBASE-T
  - 4 x 40GbE QSFP+
  - 23 production ports + 14 OOB ports + KVM/PDUs/UPS = 41 ports in use

Networking

- Network is often the bottleneck
- 10GbE over copper saves a little on cabling ($8 Cat6 vs. $70 DAC)
- Copper vs SFP+ doesn't seem to impact switch pricing
- We bond both 10GbE interfaces for more bandwidth
- 48-port Top of Rack switch can accommodate 24 hosts
  - With 2U hosts in a 42/48U rack, unlikely to actually fill this switch
- Upgrading to 40GbE would a) double the NIC cost, b) add $100/cable, c) add 20% network switch

Software Infrastructure

- CentOS 7.2 with Cloudera Manager
- PXE Bare Metal Provisioning with Cobbler + Kickstart
- Puppet for Configuration Management