GISFederal

GPUdb

A distributed database for many-core devices

801 N Quincy St
Suite 601
Arlington, VA 22203

GIS Federal Proprietary Information

GPUdb

- A distributed database designed around taking advantage of many-core devices such as GPUs, Intel Xeon Phi, and traditional x86 devices
  - Capabilities of GPUdb stem from second generation big data needs of the DoD and Intelligence community
  - GPUdb re-introduces SQL style functions on the big data scale with no pre-indexing or costly custom index development

GIS Federal Proprietary Information

GPU Supercomputing Cloud Overview

- Graphics Processing Units (GPU) based supercomputing power in the cloud
  - General Purpose Graphic Units (GP-GPUs) used for problems requiring floating point computations
  - Tremendous gains in computing power
  - Great reductions in hardware and power
    - 1/10th cost of comparable server/cluster/supercomputer
    - 1/20th the power consumption
- GIGA-threading per node vs. multi-threading per node
- Use CPUs + GPUs
- NVidia Kepler based hardware 20X the double precision horsepower of a quad-core x86 CPU

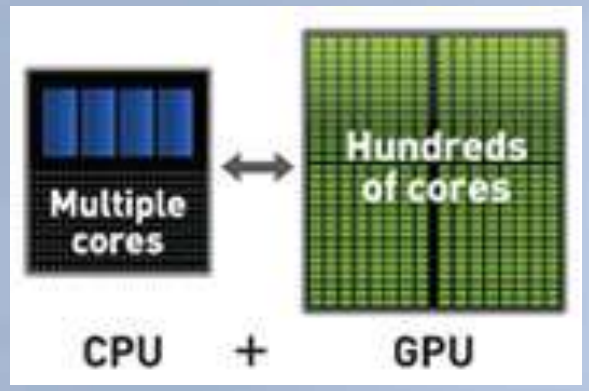

GPUdb in a box

- A big data object store and calculation engine that is accelerated with GPUs and other HPC many core hardware devices that enables big data analytics on the fly with streaming near real time data
- Calculate multi-dimensional algorithms with big data in sub-second time
- Native geospatial object support for big data visualization as a raster or video result
- High Performance Computing with commodity hardware costs
  - Order of magnitude performance gain compared to CPUs based clouds
  - Order of magnitude power reduction savings
  - Order of magnitude cost savings when compared to similar type of FPGA based big data stores while providing greater performance

GIS Federal Proprietary Information

GPUdb advantages in NoSQL space

GISFederal

- Orders of magnitude faster than relational and ‘NoSQL’ competitors
- Reduced development costs for data scaling and data analytics
  - GPUdb does not require complicated key sharding techniques that several of the major NoSQL players require such as MongoDB, Hbase, and Cassandra
- Vastly smaller power and space footprint for greater computational capability

GIS Federal Proprietary Information

Case Study: Procter & Gamble

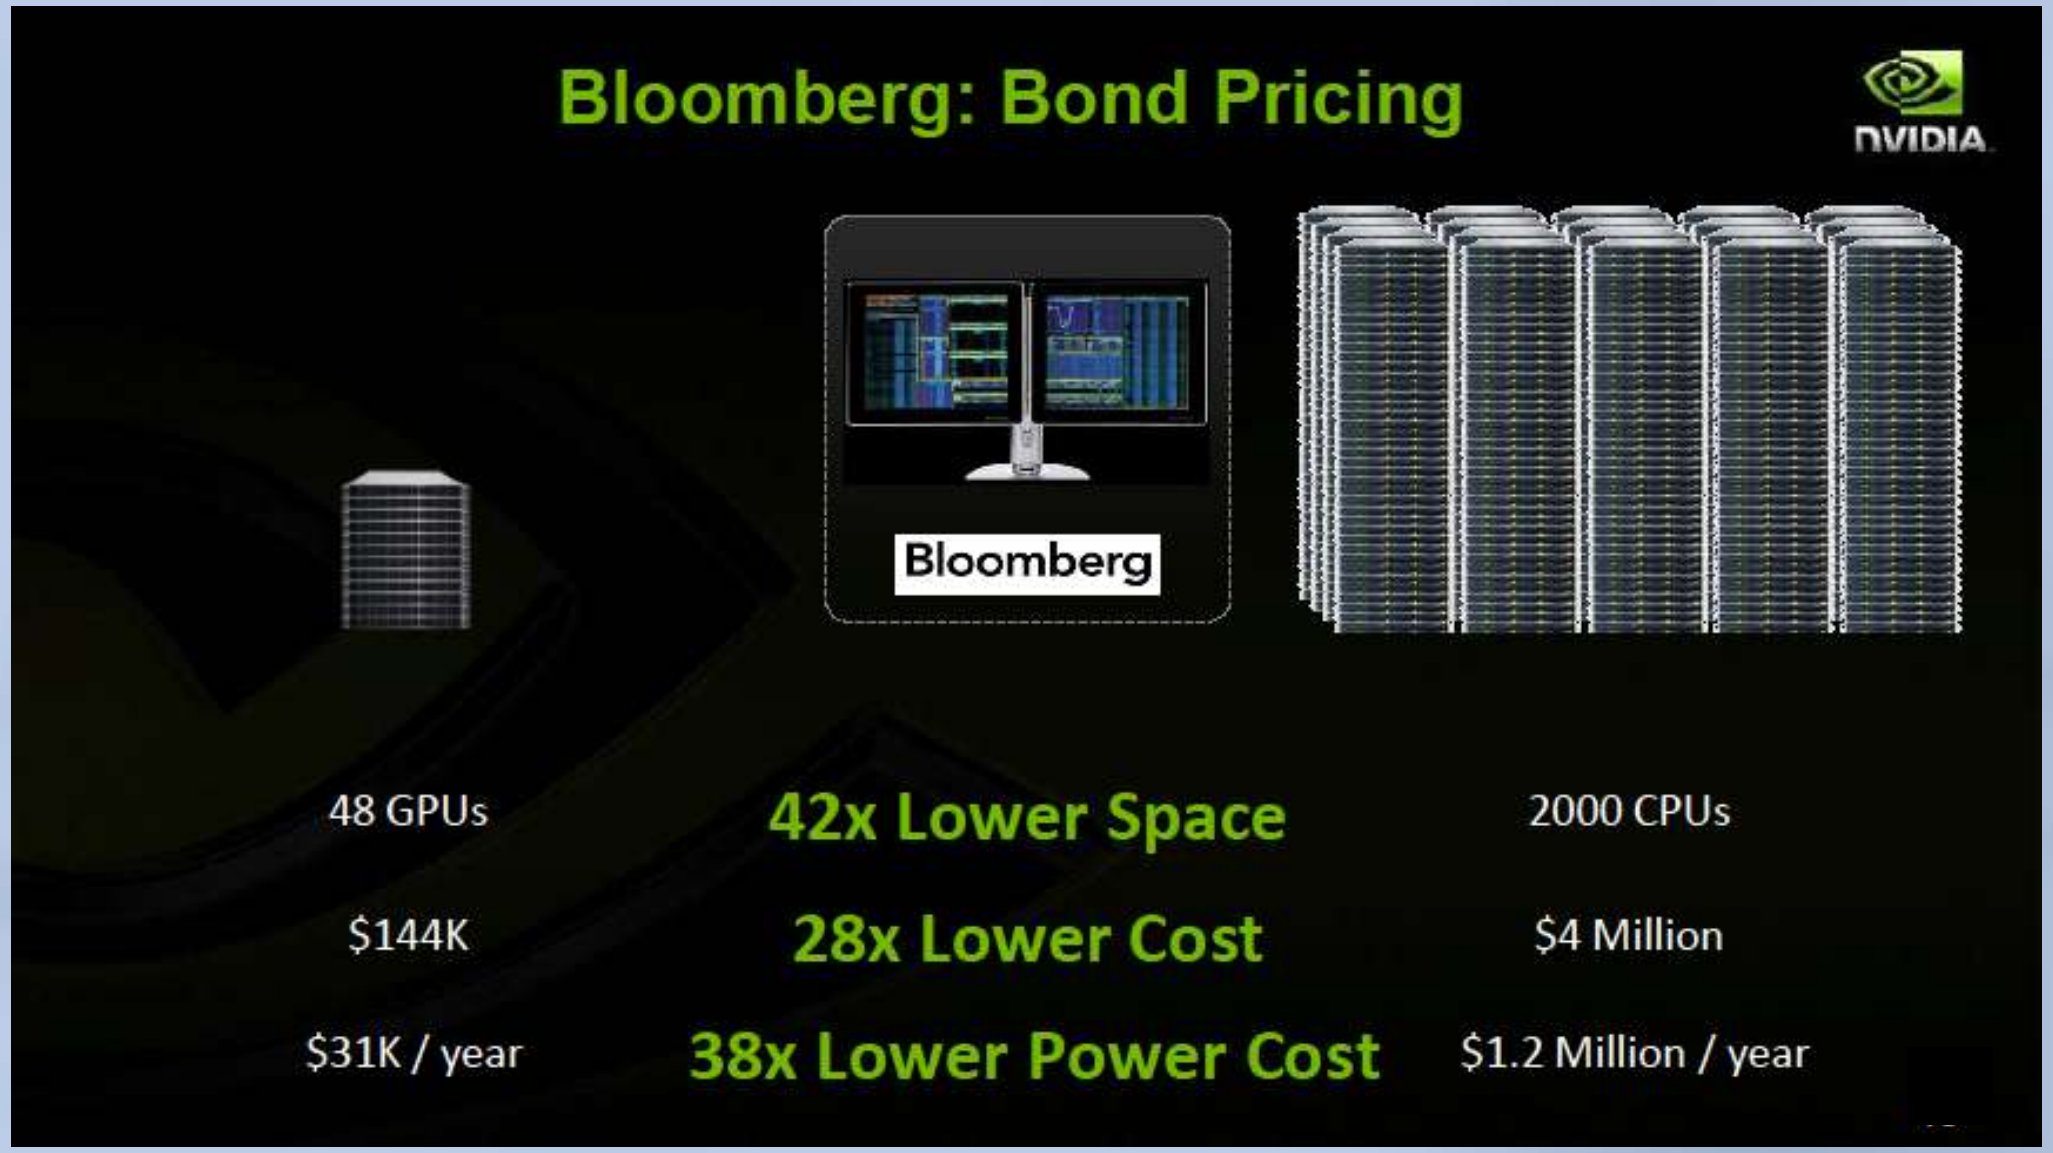

GIS Federal Proprietary Information

GPUdb Gandhi Cluster

GISFederal

- 5 node Gandhi Cluster — Total $5,040.00
  - Single Gandhi Node — $1,008.00
    - 1U used server from EBay – $869.00
      - 2x XEON HEX-CORE
      - 72GB RAM
      - 3TB HDD
    - 1x New NVidia 750Ti - $140.00
      - 640 cores
      - 2GB VRAM

Able to filter and render over 5 Billion Tweets in seconds

GIS Federal Proprietary Information

On-the-fly Time Referenced Heatmap Video Generation

(Click Play)

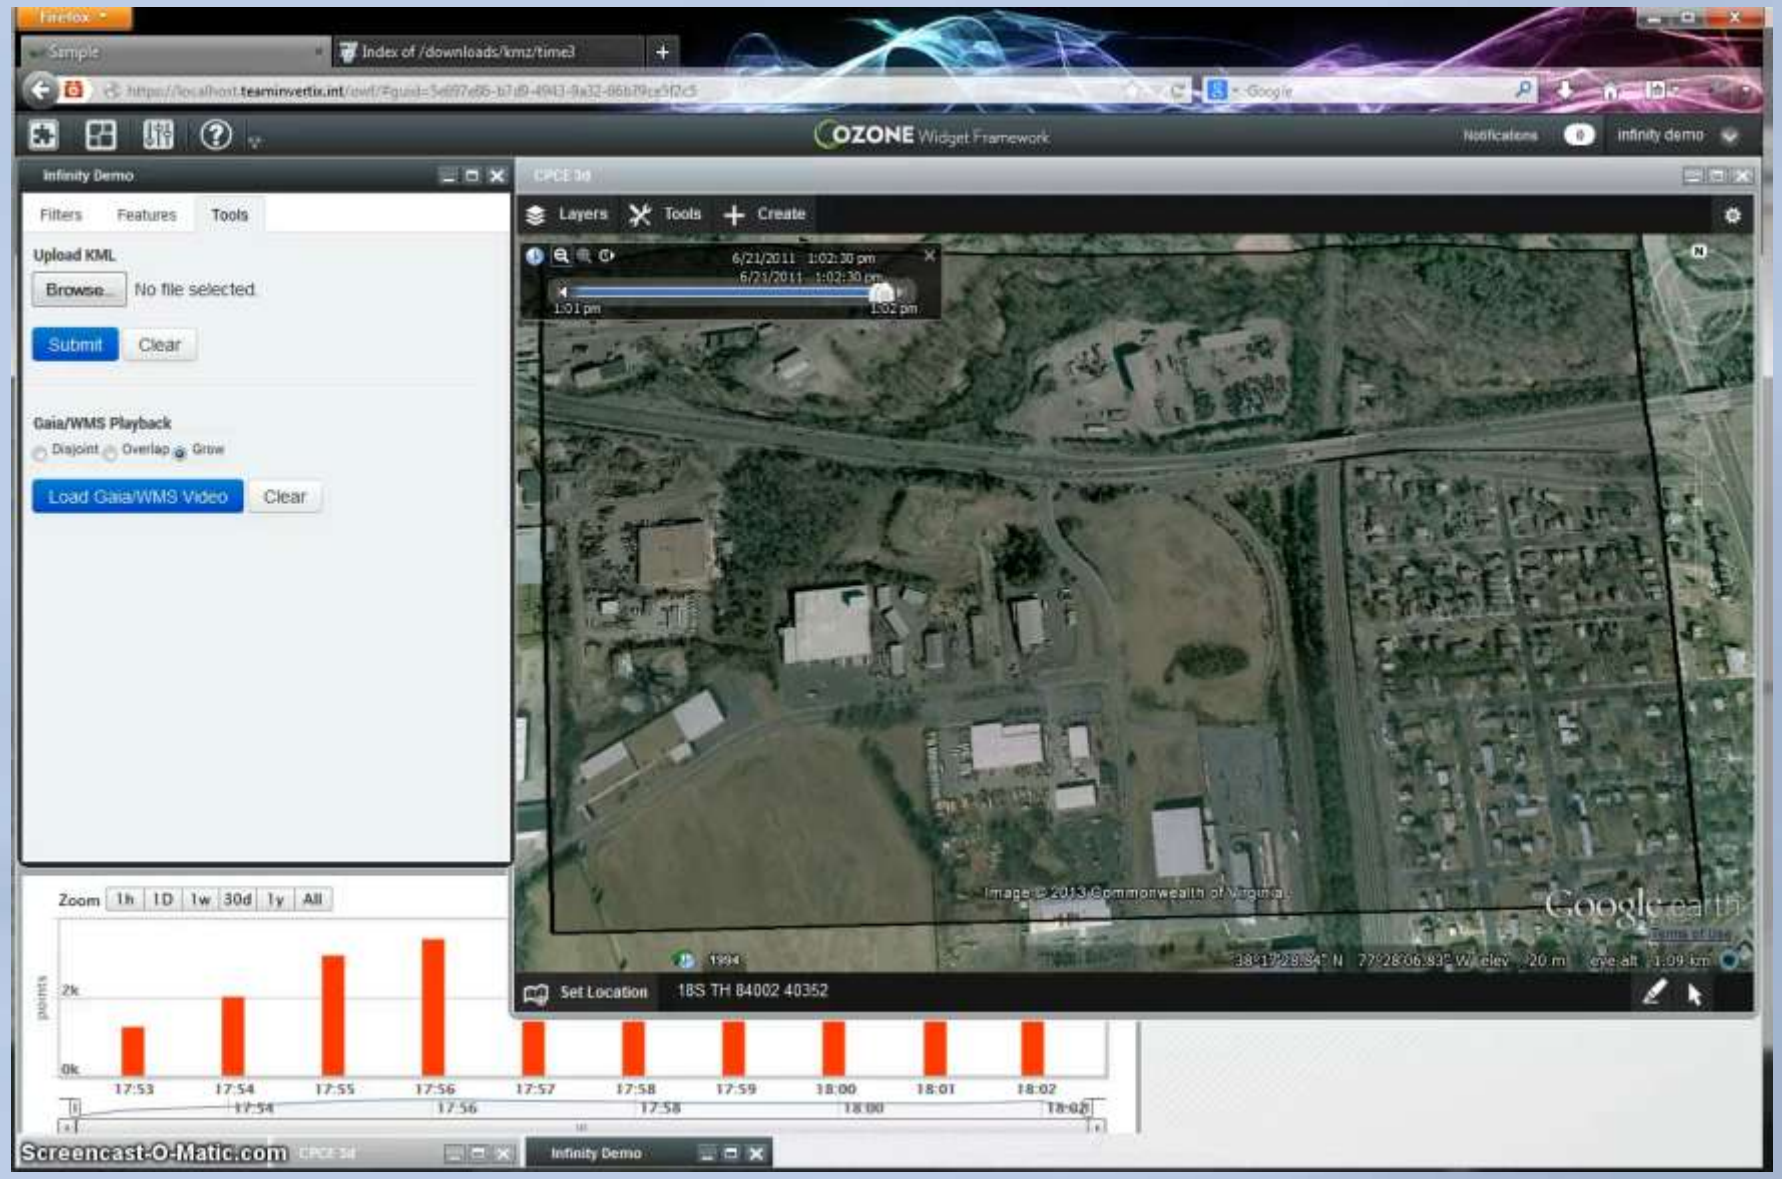

GIS Federal Proprietary Information