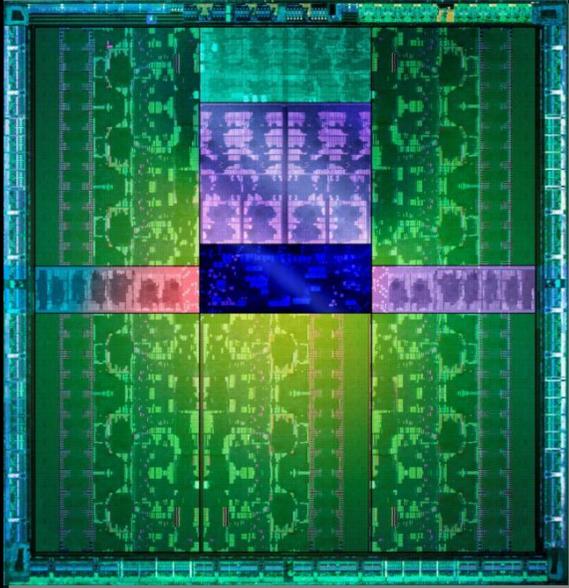


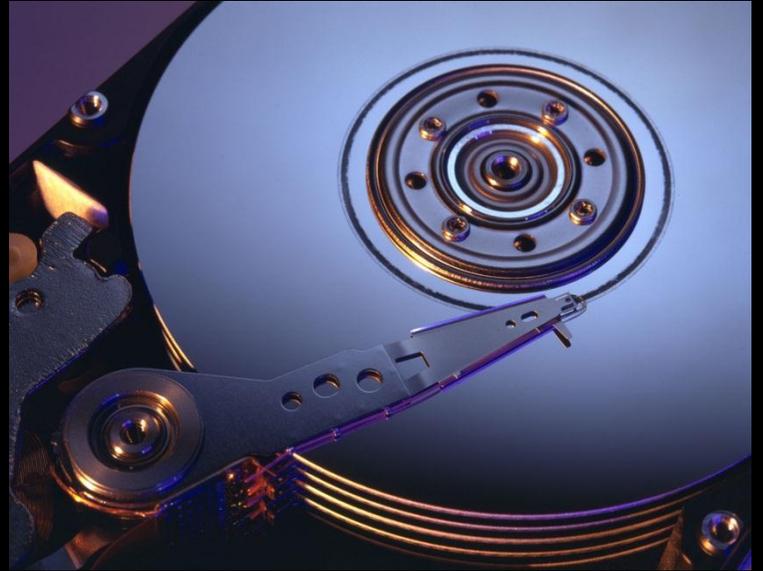
NVIDIA and Big Data

Steve Scott
CTO Tesla, NVIDIA
SOS17
March 27, 2013





& &



GPUs and Big Data Today

- Accelerating the Cloud + Mobile transformation
- Computational acceleration for Big Data
- Visualization

World Moving Increasingly to Mobile + Cloud



NVIDIA Tegra™



NVIDIA GRID™

GPU Accelerators For Big Data Analytics

Analyzing Twitter



Audio Search



Shazam

Visual Search

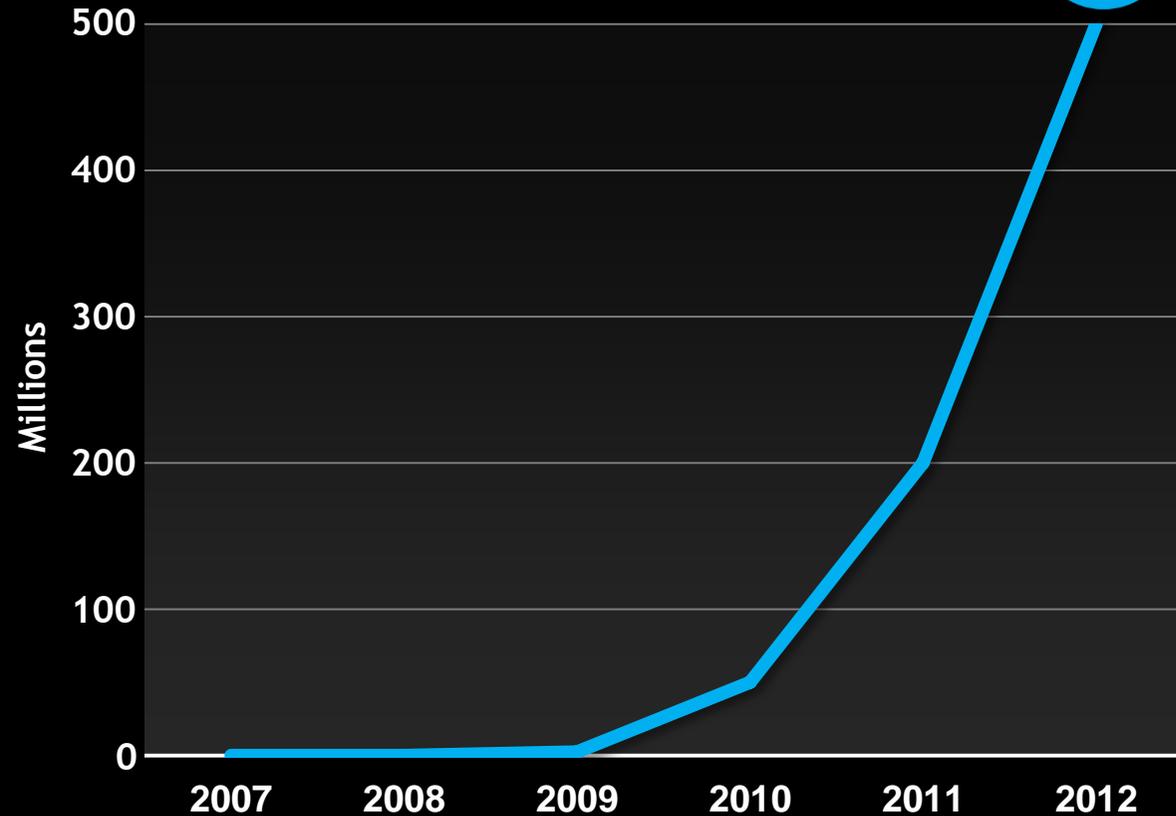


Real-time
Video Delivery



SalesForce.com: Analyzing Twitter Real-Time

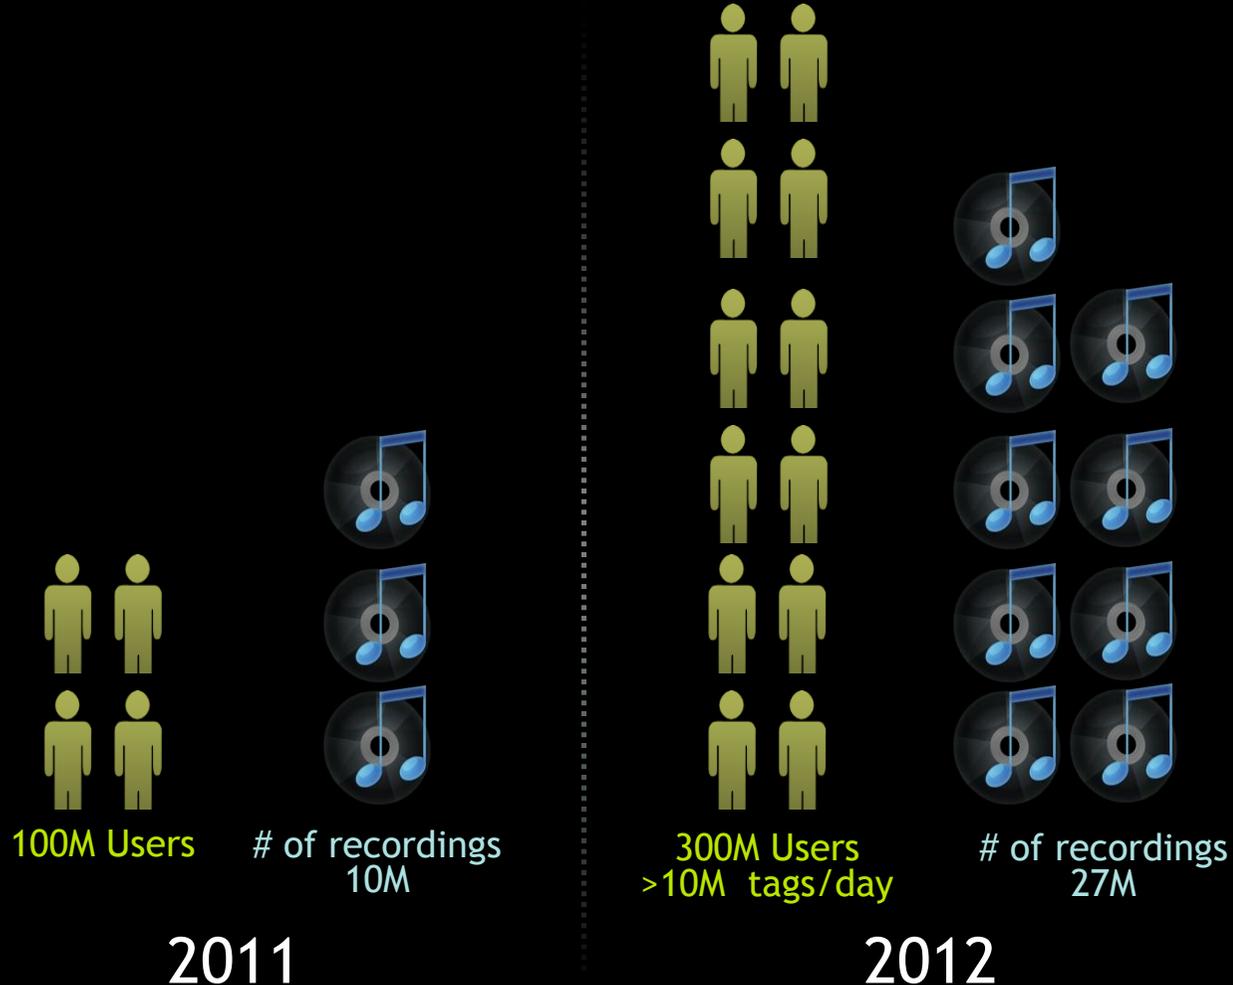
500 Million Tweets per Day



- SFDC one of 6 companies to get full Twitter firehose
- Realtime monitoring using keyword boolean expressions
- Customers may track many thousands of conversations in multiple languages



Shazam: Massive Audio Searches



- Acoustic fingerprint generated on phone, 2KB file sent to Shazam
- Matching done on hundreds of GPUs

Cortexica: Enabling Visual Shopping

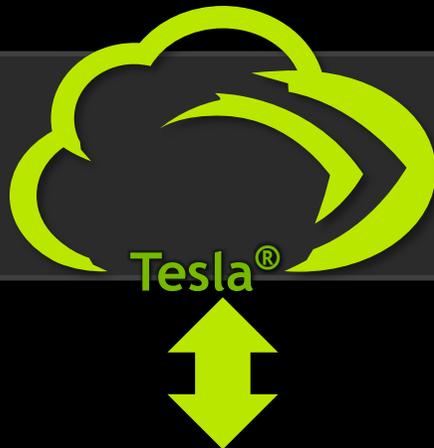


Take a Photo



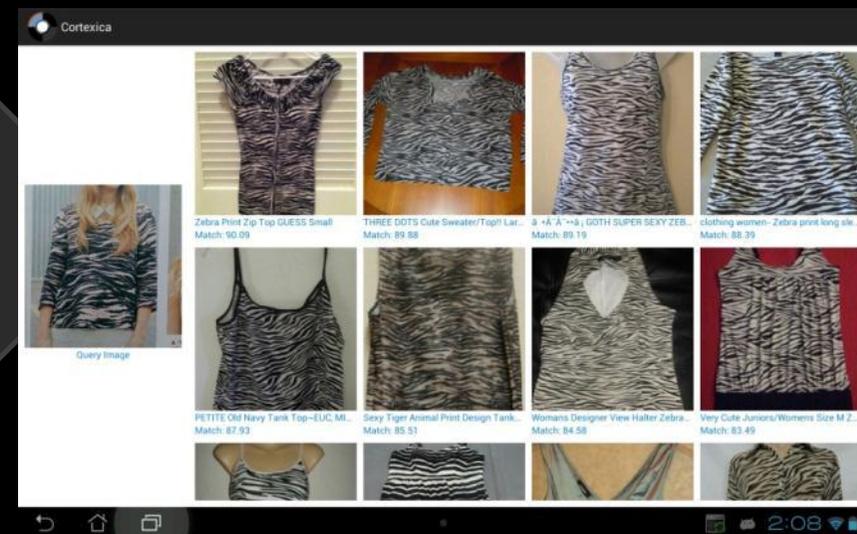
1000 key points per photo

Real-time Visual Search



Database: 1M+ Apparel

Results in seconds



Now think video...



Elemental: Managing Live Video Streams

12M+ live video streams during Hurricane Sandy



The Weather Channel

Rescale to every screen size & resolution in real-time



ELEMENTAL 

38M mobile consumers monthly



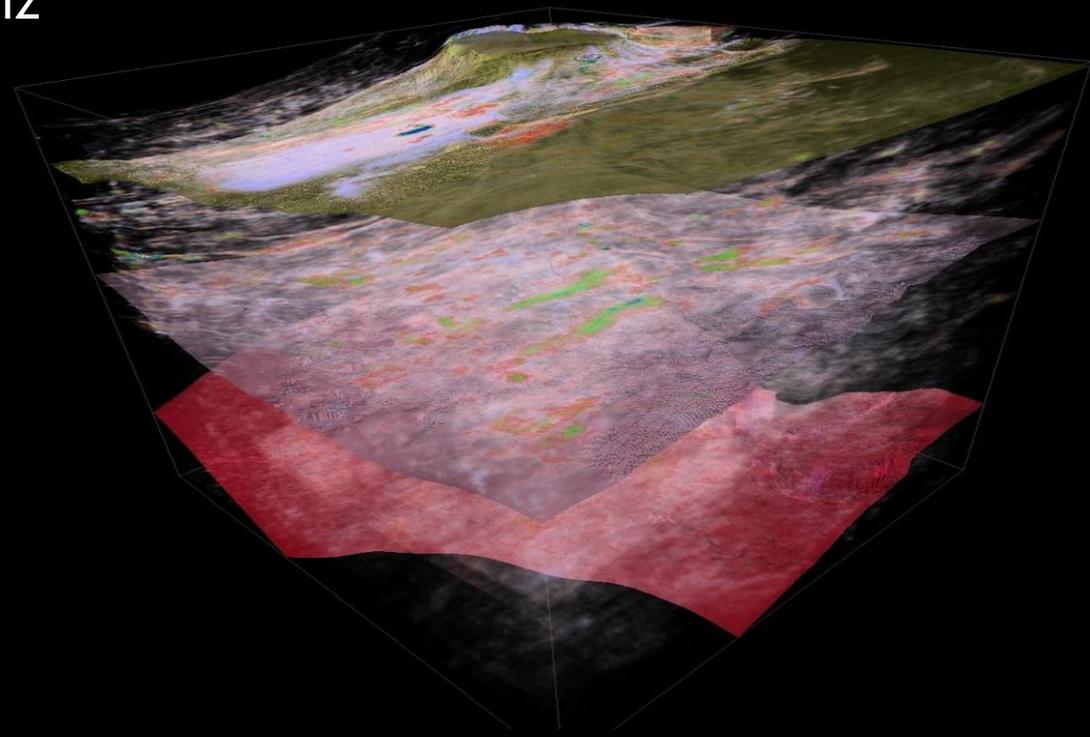
Also used for other live events (e.g.: Olympics)

NVIDIA indeX - Scalable Big Data Visualization



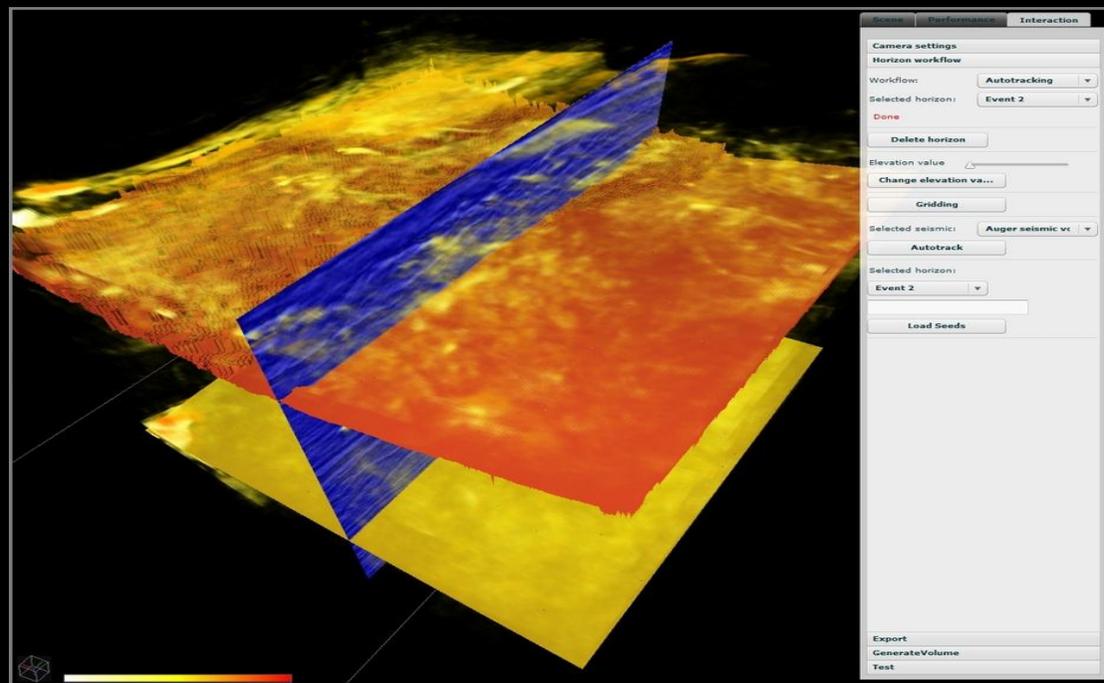
NVIDIA indeX

- Graphics infrastructure which can be integrated with ISV applications or other viz tools to deliver a scalable and interactive visualization solution
 - C++ library
 - Extensible via user-defined functions
- Volumetric ray tracing, iso-surfaces
- Large display device rendering (Quad HD)
- Visualizes combinations of volumetric and surface data accurately and at high visual quality



NVIDIA indexX

- Enables comprehensive visual insight into complex data models by providing feature extraction, data analysis, and attribute generation
- Interactive 3D navigation on O(TB) datasets
- Cloud-based GPU rendering for data sharing to support collaborative analysis
 - H.264 streaming output
 - Multiple independent cameras for remote teams on shared data



GPUs Today

- Computational acceleration for Big Data
- Visualization
- Accelerating the Cloud + Mobile transformation

GPUs Tomorrow

- **Converged architecture for Big Data and Compute**

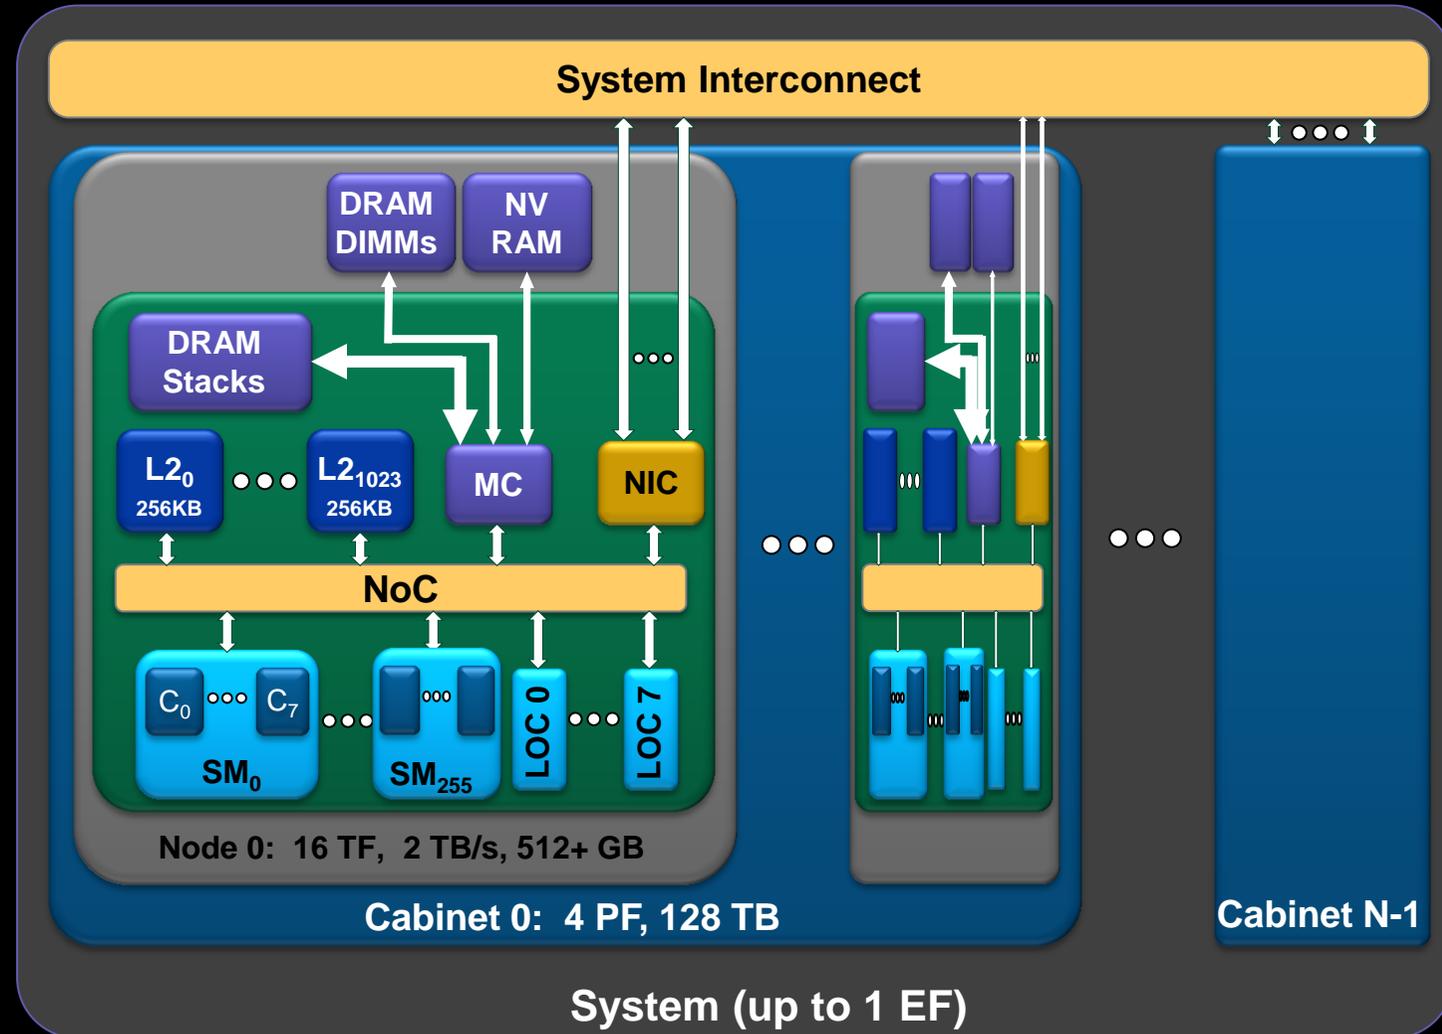
Long Term Vision: Converged Compute/Data Architecture

GPUs Today:

- Lots of on-node bandwidth and latency tolerance
- Excellent at single-node Big Graph problems today (“Medium Graph”?)
- E.g.: Breadth First Search (BFS) and Maximal Independent Set (MIS)
- Excellent parallel speedup using asymptotically efficient algorithms

GPUs Tomorrow:

- Extend single node characteristics to whole machine
- Global address space
- Aggressive and configurable memory hierarchy & interconnect
- Sophisticated/lightweight synchronization



Thank You

