

# Slides for Netflix Visitor Sandia Cloud Effort

---

- **Hadoop, Cassandra, other “big data” frameworks**
  - Graph algorithms
  - Machine learning
- **Commodity cluster (“Nebula”)**
  - 280 cores + 280 disk drives
  - 560 TB raw disk capacity
  - 876 GB RAM
  - 2 racks






# Web Refer Graphs

---

- Directed graph based on HTTP Referer header and URL destination
- Count degree, triangles, SSSP, connected component, page rank, graph spectra



# Power Law Distribution of In-Degree

