

# Introduction

- The warehouse-scale computer (WSC)' is the foundation of Internet services many people use every day: search, social networking, online maps, video sharing, online shopping, email services, and so on. The tremendous popularity of such Internet services necessitated the creation of WSCs that could keep up with the rapid demands of the public. Although WSCs may appear to be just large datacenters, their architecture and operation are quite different.
- Today's WSCs act as one giant machine and cost on the order of \$150M for the building, the electrical and cooling infrastructure, the servers, and the networking equipment that connects and houses 50,000 to 100,000 servers. Moreover, the rapid growth of cloud computing makes WSCs available to anyone with a credit card.
- The WSC is the modern descendant of the supercomputer-making Seymour Cray the godfather of today's WSC architects. WSCs have many orders of magnitude more users than high-performance computing, and they represent a much larger share of the IT market.
- WSC architects share many goals and requirements with server architects:
  - 1. Cost-performance-Work done per dollar is critical in part because of the scale. Reducing the capital cost of a WSC by 10% could save \$15M.
  - 2. Energy efficiency-Power distribution costs are functionally related to power consumption; you need sufficient power distribution before you can consume power. You need to get out the heat that you put in. Hence, peak power and consumed power drive both the cost of power distribution and the cost of cooling systems. . Moreover, energy efficiency is an important part of environmental stewardship. Hence, work done per joule is critical for both WSCs and servers.

- 3. Dependability via redundancy-The long-running nature of Internet services means that the hardware and software in a WSC must collectively provide at least 99.99% of availability; that is, it must be down less than 1 hour per year. Redundancy is the key to dependability for both WSCs and servers. While server architects often utilize more hardware offered at higher costs to reach high availability, WSC architects rely instead on multiple cost-effective servers connected by a low-cost network and redundancy managed by software.
- 4. Network I/O-Server architects must provide a good network interface to the external world, and WSC architects must also. Networking is needed to keep data consistent between multiple WSCs as well as to interface to the public.
- 5. Both interactive and batch processing workloads-While you expect highly interactive workloads for services like search and social networking with millions of users, WSCs, like servers, also run massively parallel batch programs to calculate metadata useful to such services. For example, MapReduce jobs are run to convert the pages returned from crawling the Web into search indices.
- There are also characteristics not shared with server architecture:
  - 1. Ample parallelism-A concern for a server architect is whether the applications in the targeted marketplace have enough parallelism to justify the amount of parallel hardware and whether the cost is too high for sufficient communication hardware to exploit this parallelism. A WSC architect has no such concern. First, batch applications benefit from the large number of independent datasets that require independent processing, such as billions of Web pages from a Web crawl. This processing is data-level parallelism applied to data in storage instead of data in memory. Second, interactive Internet service applications, also known as software as a service (SaaS), can benefit from millions of

independent users of interactive Internet services. Reads and writes are rarely dependent in SaaS, so SaaS rarely needs to synchronize. For example, search uses a read-only index and email is normally reading- and writing-independent information. We call this type of easy parallelism request-level parallelism, as many independent efforts can proceed in parallel naturally with little need for communication or synchronization; for example, journal-based updating can reduce throughput demands.

- 2. Operational costs count-Server architects usually design their systems for peak performance within a cost budget and worry about power only to make sure they don't exceed the cooling capacity of their enclosure. They usually ignore operational costs of a server, assuming that they pale in comparison to purchase costs. WSCs have longer lifetimes-the building and electrical and cooling infrastructure are often amortized over 10 or more years-so the operational costs add up: Energy, power distribution, and cooling represent more than 30% of the costs of a WSC in 10 years.
- 3. Scale and the opportunities/problems associated with scale-Often extreme computers are extremely expensive because they require custom hardware, and yet the cost of customization cannot be effectively amortized since few extreme computers are made. However, when you purchase 50,000 servers and the infrastructure that goes with it to construct a single WSC, you do get volume discounts. WSCs are so massive internally that you get economy of scale even if there are not many WSCs.