


Google Datacenter




Datacenter Organization



Single server:

- 4-8 cores
- DRAM: 4-16GB @ 100ns
- Disk: 2 TB @10ms



Rack:

- 50 machines
- DRAM: 200-800GB @ 300 μ s
- Disk: 100TB @ 10ms



Row/cluster:

- 30+ racks
- DRAM: 6-24TB @ 500 μ s
- Disk: 3 PB @ 10ms

Sun Containers



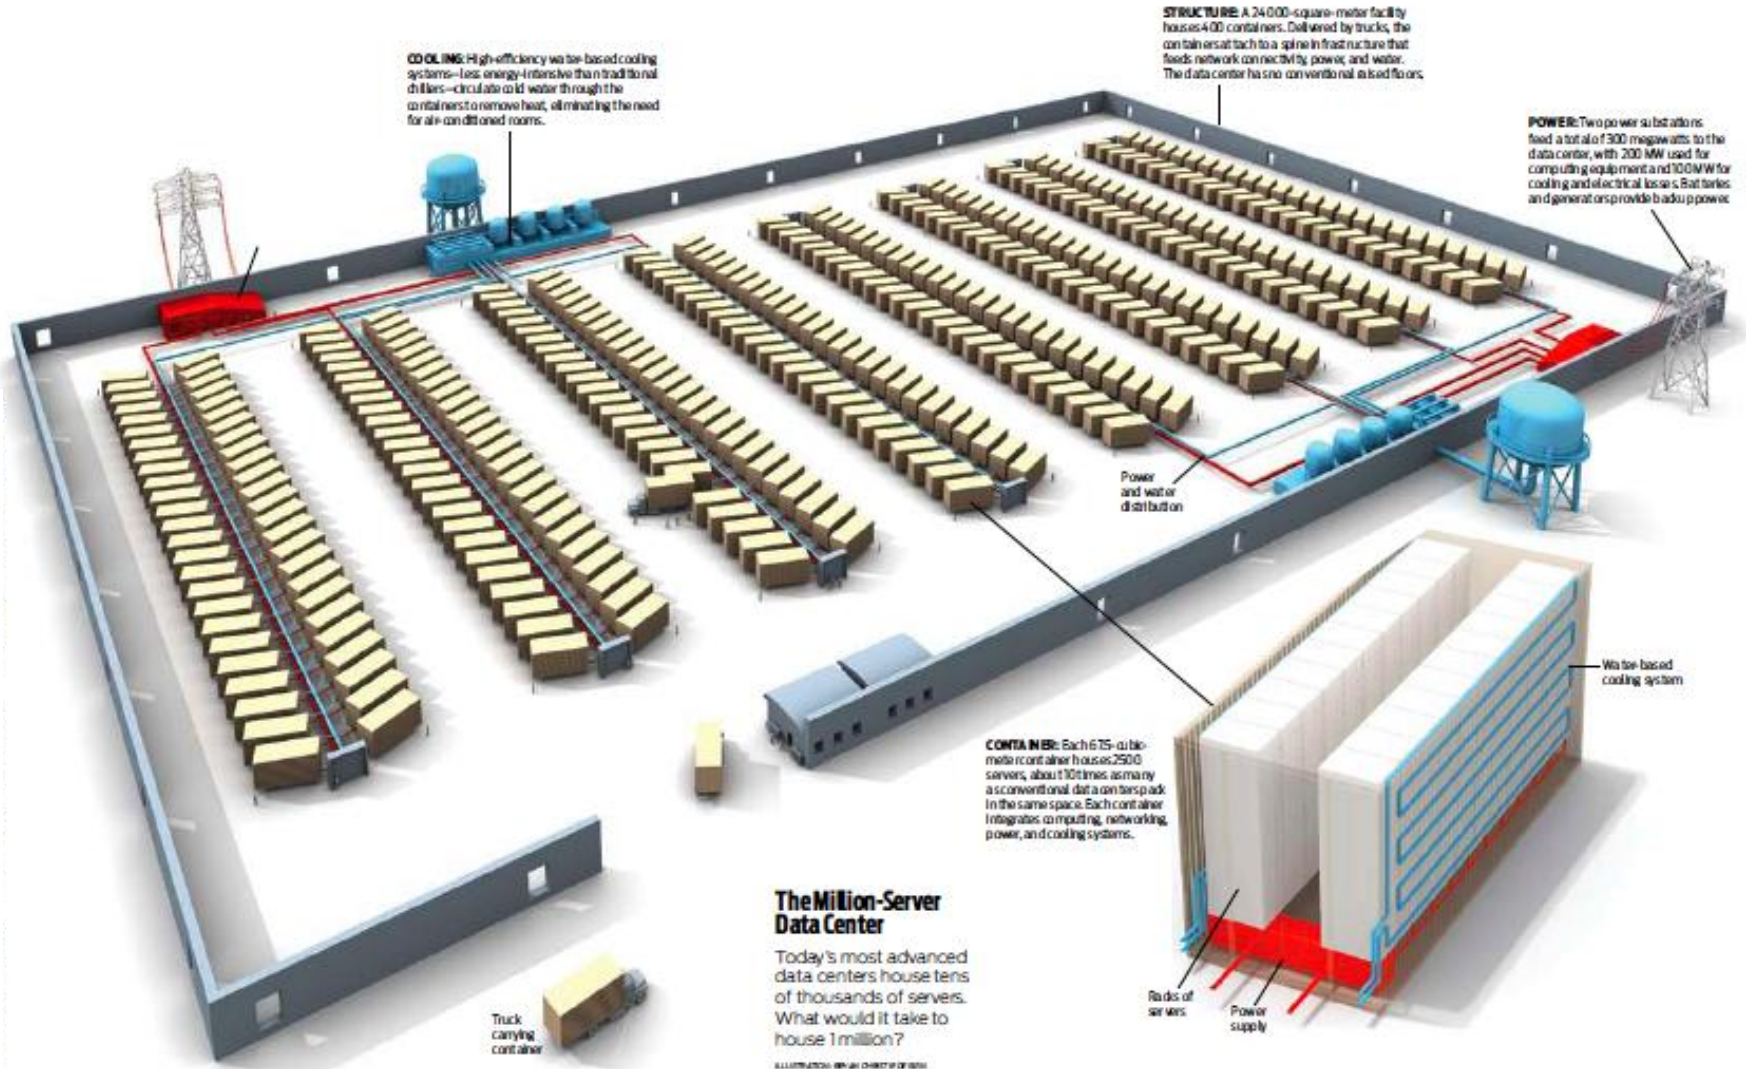
Sun Containers, cont'd



Google Containers



Microsoft Containers



Microsoft Containers, cont'd



Failures are Frequent

Typical first year for a new cluster (Jeff Dean, Google):

- ~0.5 **overheating** (power down most machines in <5 mins, ~1-2 days to recover)
- ~1 **PDU failure** (~500-1000 machines suddenly disappear, ~6 hours to come back)
- ~1 **rack-move** (plenty of warning, ~500-1000 machines powered down, ~6 hours)
- ~1 **network rewiring** (rolling ~5% of machines down over 2-day span)
- ~20 **rack failures** (40-80 machines instantly disappear, 1-6 hours to get back)
- ~5 **racks go wonky** (40-80 machines see 50% packet loss)
- ~8 **network maintenances** (4 might cause ~30-minute random connectivity losses)
- ~12 **router reloads** (takes out DNS and external vips for a couple minutes)
- ~3 **router failures** (have to immediately pull traffic for an hour)
- ~dozens of minor **30-second blips** for DNS
- ~1000 **individual machine failures**
- ~thousands of **hard drive failures**
- **Slow disks, bad memory, misconfigured machines, flaky machines**, etc.
- Long distance links: **wild dogs, sharks, dead horses, drunken hunters**, etc.

How Many Datacenters?

- **1-10 datacenter servers/human?**
- **100,000 servers/datacenter**

	U.S.	World
Servers	0.3-3B	7-70B
Datacenters	3000-30,000	70,000-700,000

- **80-90% of general-purpose computing in datacenters?**

