

The Datacenter as a Computer: An Introduction to the Design of Warehouse-Scale Machines

Chapters 1-2

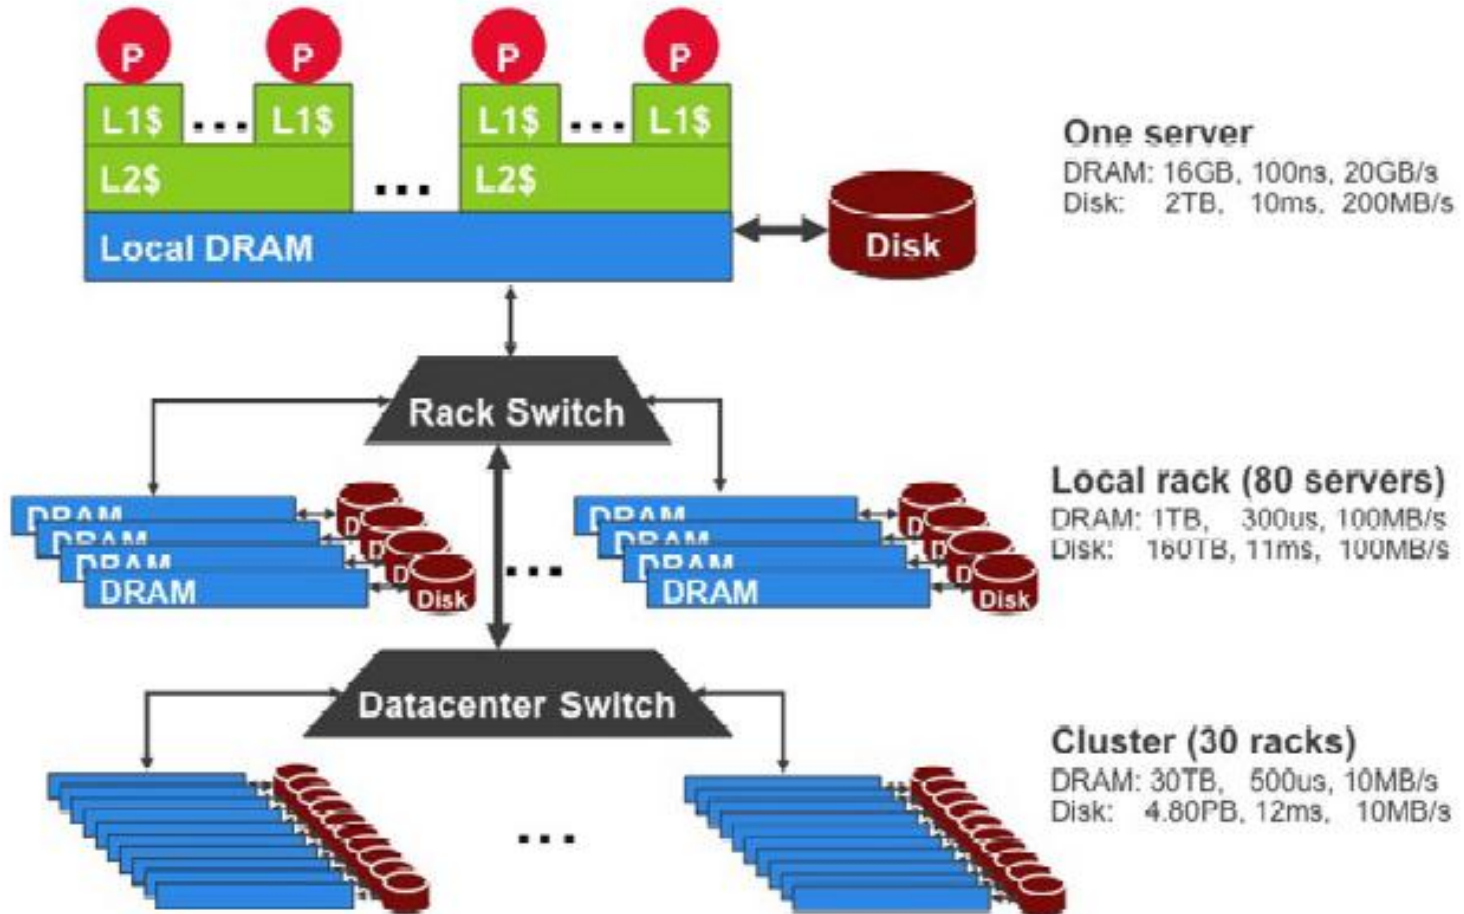
Patrick Wendell

QUICK SUMMARY / KEY POINTS

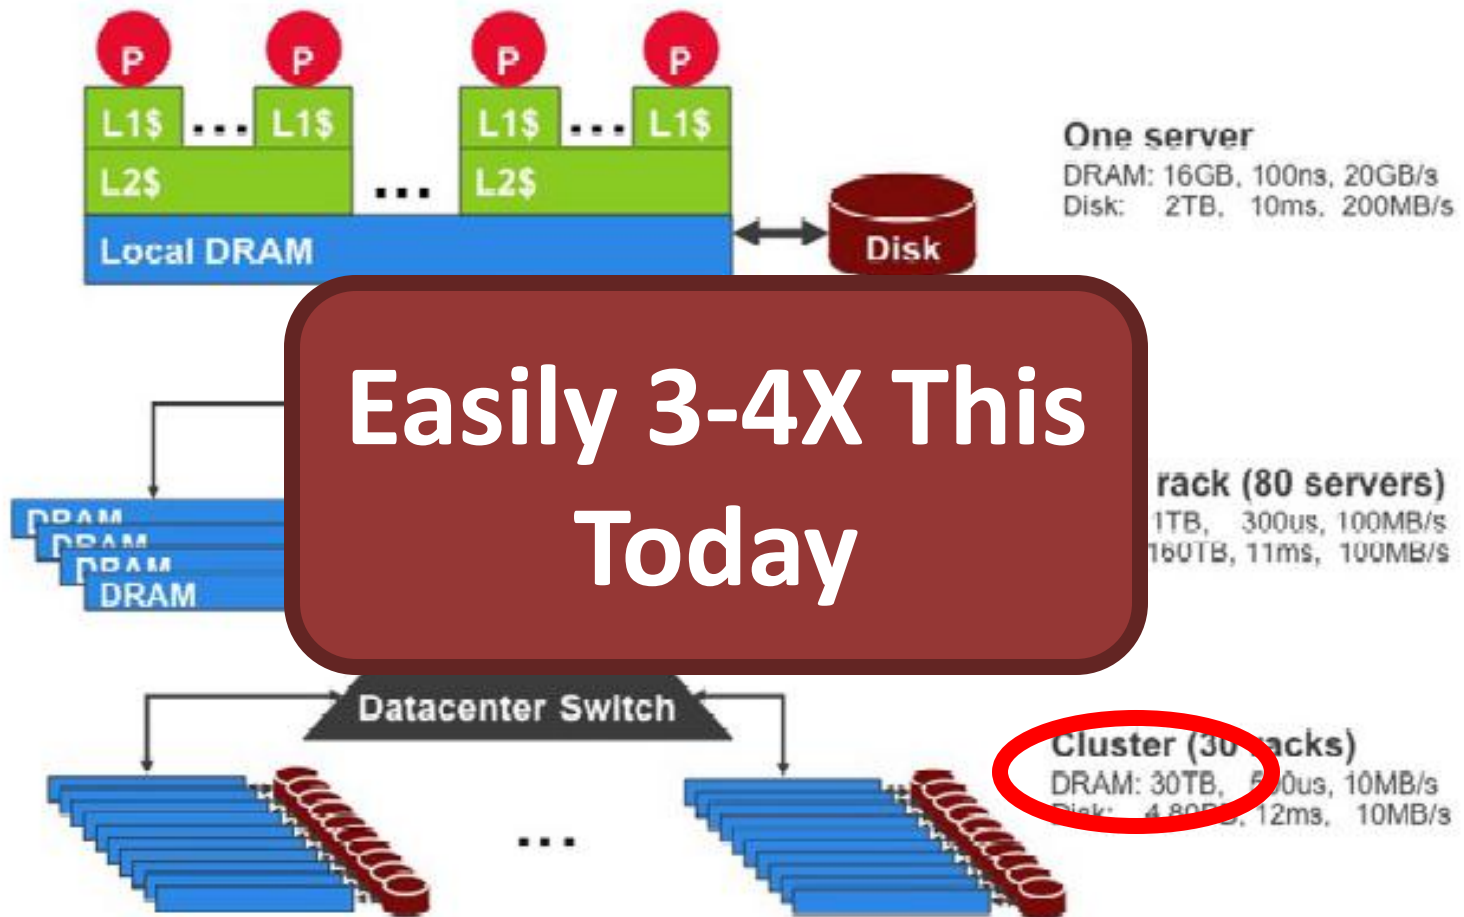
Why do we need cluster-scale computing?

- Problems are data-intensive, e.g. Web Search
- Client-server model software decreases costs
(Network Computer redux???)
- Competition for functionality spurs increased hardware investment


Available Resources and Their Properties



Available Resources and Their Properties



Available Resources and Their Properties



Parallel Parallel Parallel

- **Great news!** Data is highly parallelizable
 - Workloads easily partitioned (by website, user, time period, etc)
 - Replicate data, distribute computation, and we're finished
- **Terrible news!** Data is highly parallelizable
 - Latency becomes an issue: e.g. Hang on slowest task, gets worse as # tasks increases
 - Debugging & monitoring in parallel environments much more complicated

OPINIONS AND RESPONSE

Overall Impressions

- Important publication from cluster computing thought leaders (Google) but likely obsolete now
- Good wrap-up to formative years (2000 – 2006) of datacenter applications growth. Primitives here likely to persist for a while.

Overall Impressions

- Seems to ignore multi-tenancy environments
 - Maybe because App-Engine is shitty?
 - Wonder how much datacenter computation is going to be in multi-tenant settings looking forward
- No mention of the “cloud” except intro. What is this “cloud” thing I keep hearing about?

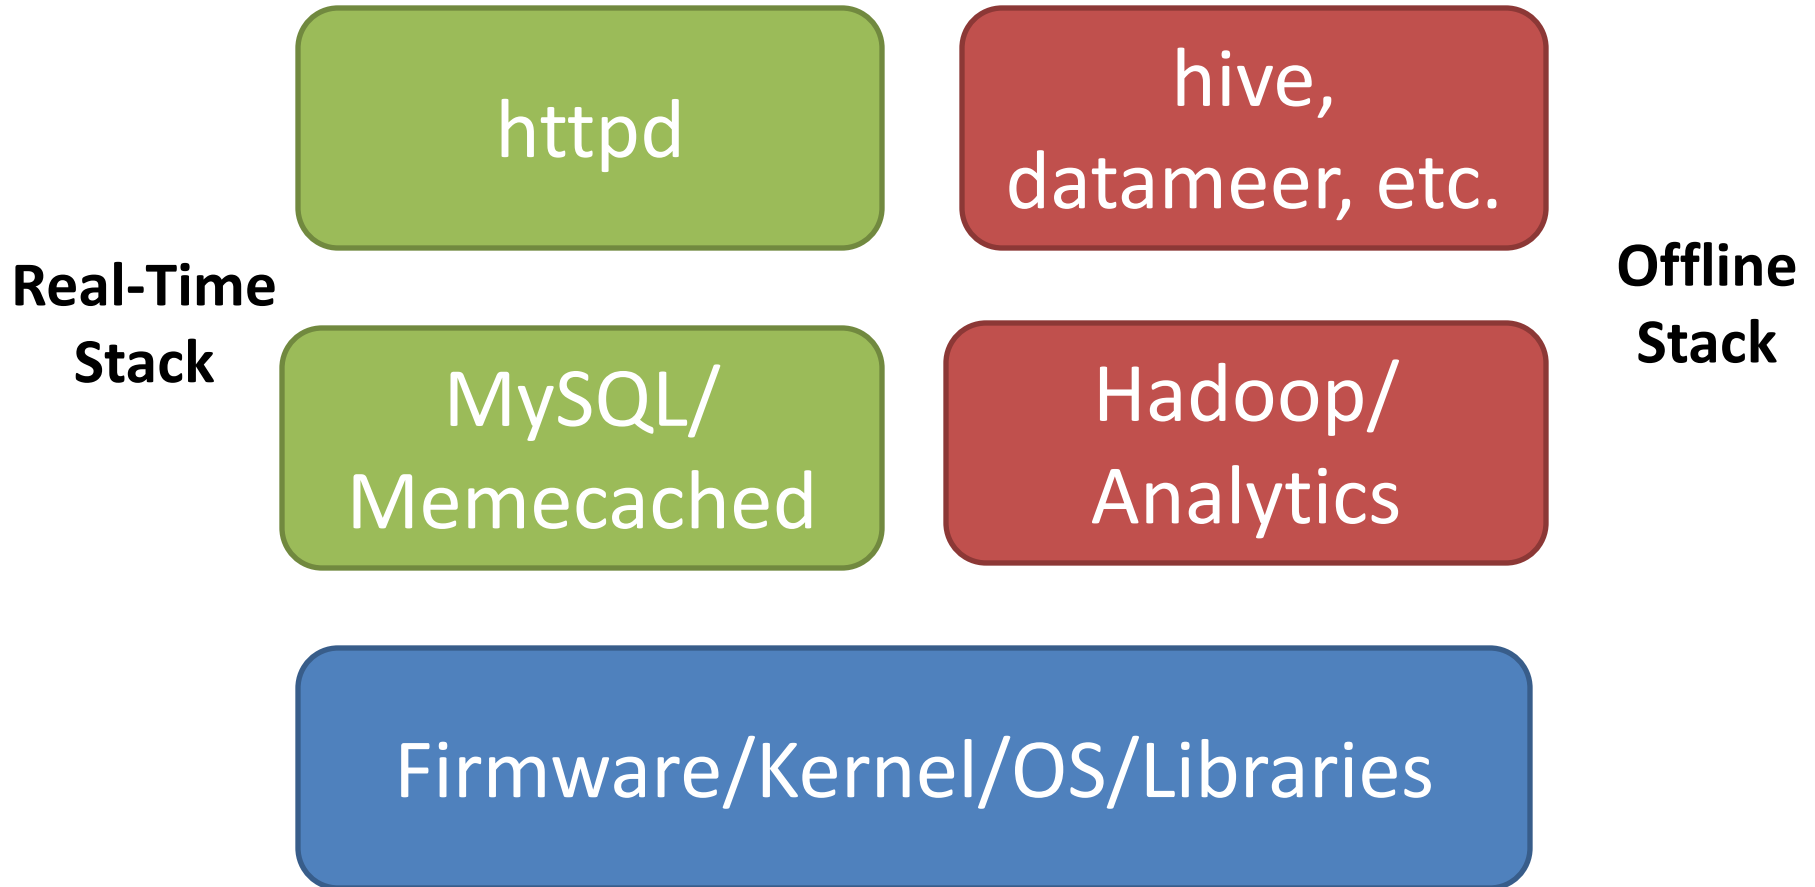
The Datacenter Stack (Proposed)

Application Software
(Google, Gmail, Google Maps)

Cluster Software
(MapReduces, GFS, BigTable)

Firmware/Kernel/OS/Libraries

The Datacenter Stack (Today)



The Datacenter Stack

- Google had trouble introducing Gmail (real-time) on existing infrastructure (GFS)
- How do we design the “middle layer” to provide a useful substrate for many services with different requirements?
- How much application-specific optimization is appropriate?

Claim: Massive Cluster Computing to Become Pervasive (?)

- From lecture
 - Google: = 1 mil servers
 - Microsoft, Yahoo!, IBM, HP, Amazon: 100,000(s)
 - Ebay, GoDaddy, Facebook, Akamai: > 50,000
- My experience
 - Very long tail of companies well served by clusters of 100 or fewer nodes (4 OOM < Google)
 - Large corporations also partition clusters
 - Multi-tenancy environments have large number of small users, not inverse
 - Individual clusters likely to stay same size (??)

Inter-Datacenter Application Logic

- Punts on inter-datacenter design
- I think this is an interesting area
 - Availability during datacenter failures
 - Understanding consistency between datacenters (mostly master-slave replication today)
 - Workload migration in response to diurnal patterns