Improve the Availability of Cloud Datacenter

Group 3

FENG Jinghao, HU Lin, KANG Wangcheng, WEN Zechu, ZHENG Jie

Background and Problem

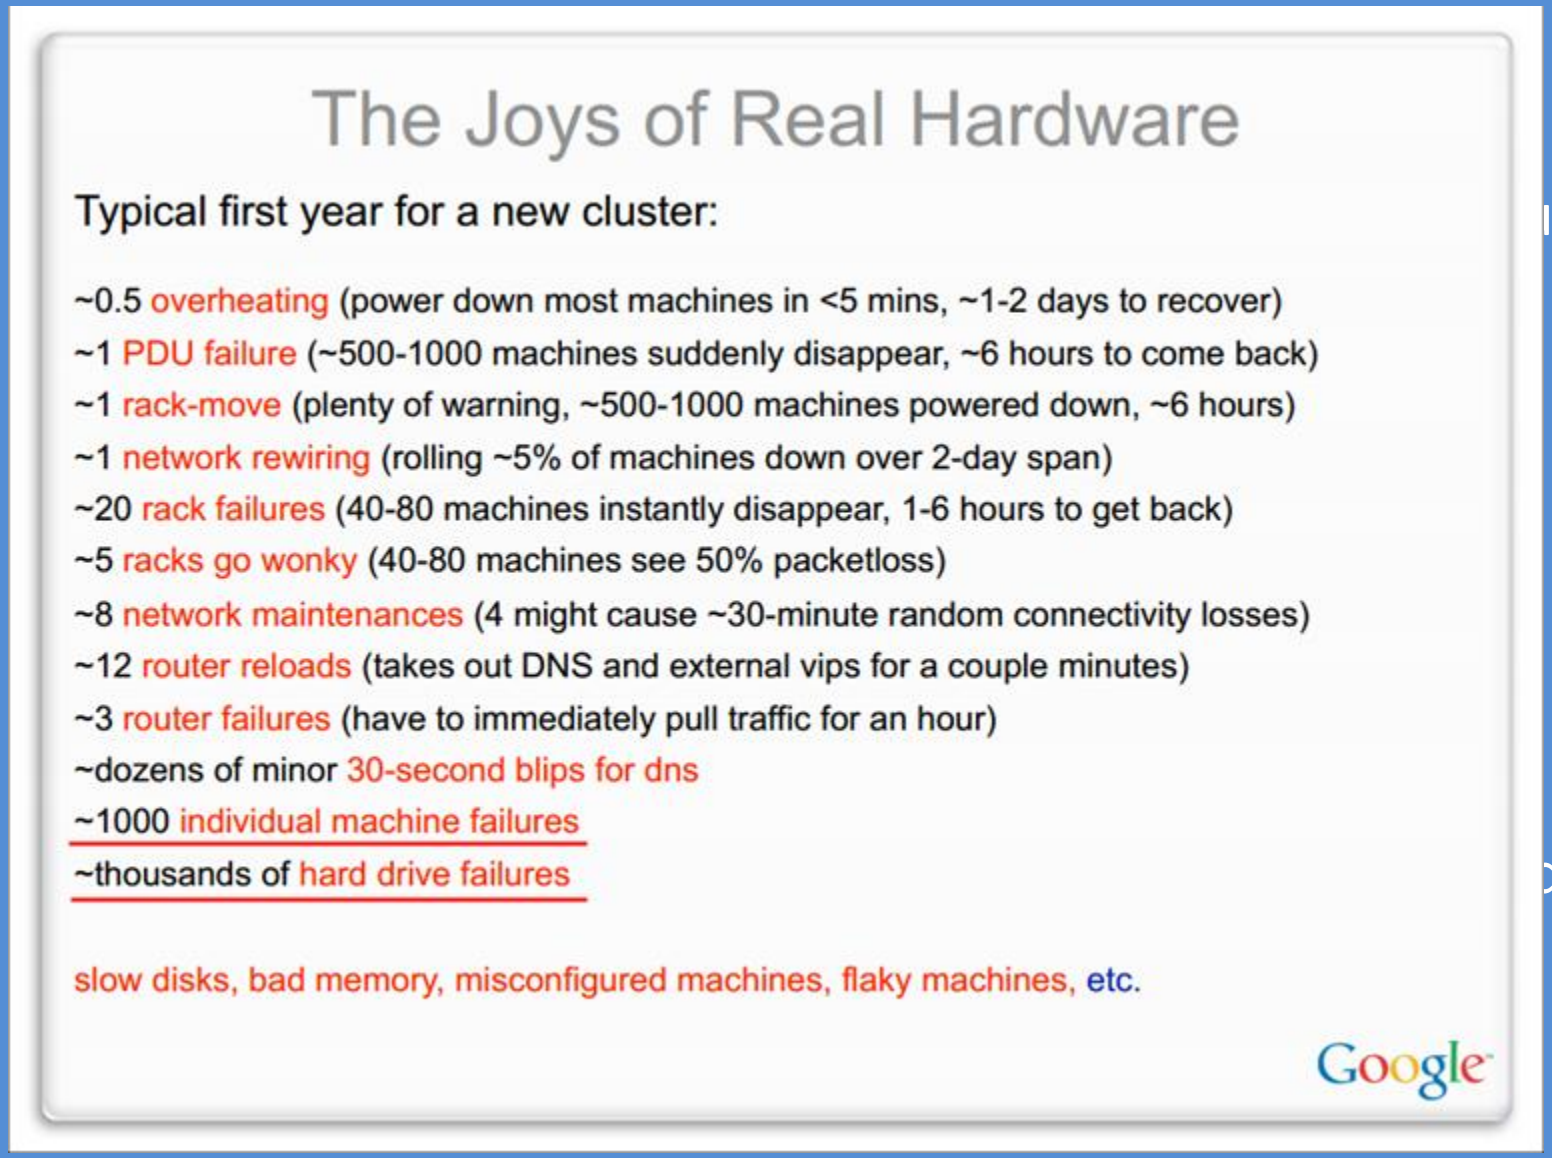

The Joys of Real Hardware

Typical first year for a new cluster:

~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% packetloss)
~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.

Google

- ıite & as
- des

Dynamic Anomaly Detection

- Based on the dataset collected from node and sensor in every minutes
  - CPU load/memory user/disk accesses/network traffic ,etc.
  - fan speed/temperature/moisture ,etc.
- Use **Dynamic Anomaly Detection** to warn possible malfunction in advanced
- Use **map-reduce** because of massive dataset

Methodology

- D
- A
- N

An

1.

2.

3.

```
r          #the nearness between the data values in cluster
total      #the total number of reported processes
data[]     # the reported and sorted values
index=1    #iterator for each process
while index <=total:
    new cluster   #a distribution zone of the data values
    add data[index] to cluster
    while index <=total-1:
        if data[index]-data[index +1]<=r:
            index +=1
            add data[index+1] to cluster
        else:
            index +=1
            break
mark the first max cluster  as "normal"
mark clusters with higher  value as "high"
mark clusters with lower values as "low"
compute the deviation of each abnormal cluster
```

Figure 2. Pseudo code

Andrew Ng

- Notify staff to check suspicious nodes

Preliminary result

- We
- An d ove we

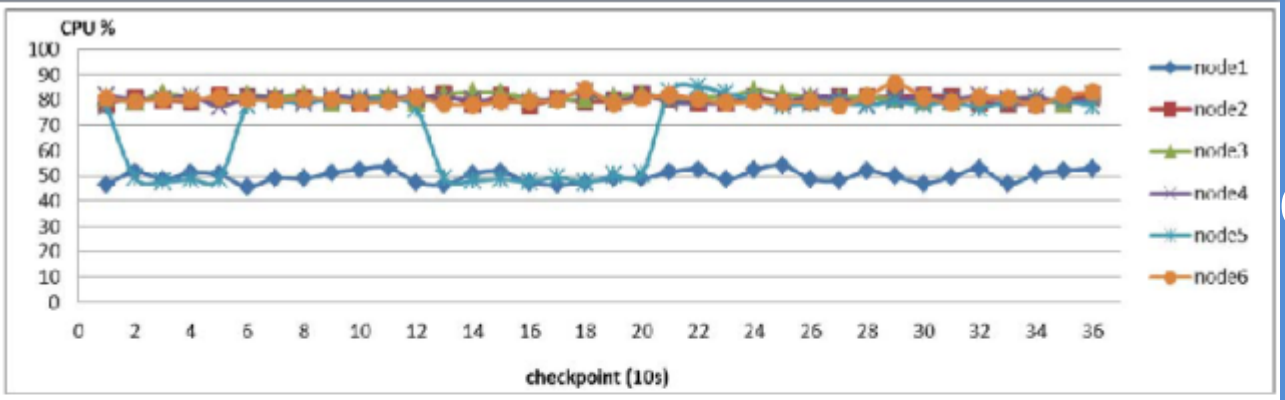

Figure 4. CPU utilizations

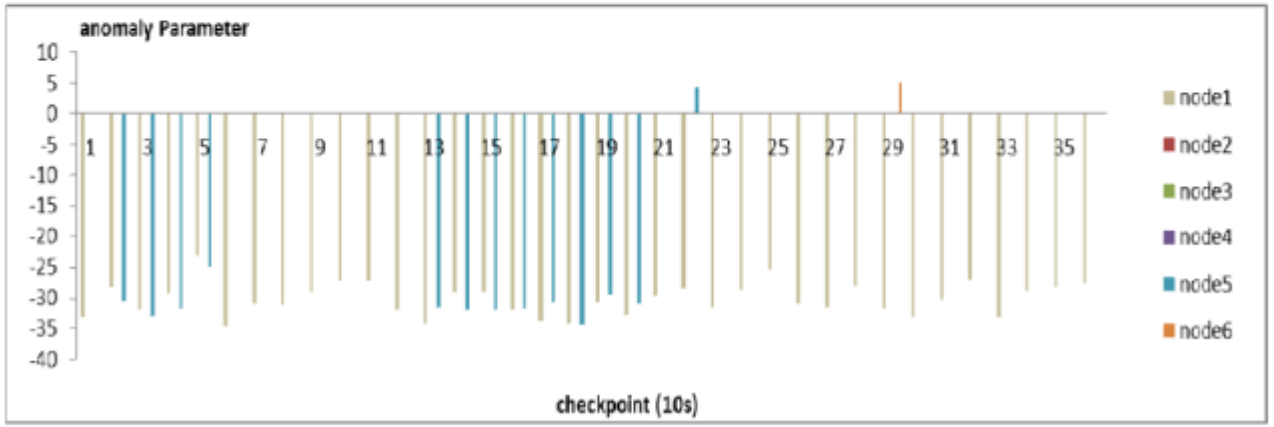

Figure 5. Anomaly detection of CPU utilizations

Q & A

References

- Liao S, Hung T H, Nguyen D, et al. Machine learning-based prefetch optimization for data center applications[C]//Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis. ACM, 2009: 56.
- Bodık P, Griffith R, Sutton C, et al. Statistical machine learning makes automatic control practical for internet datacenters[C]//Proceedings of the 2009 conference on Hot topics in cloud computing. 2009: 12-12.