

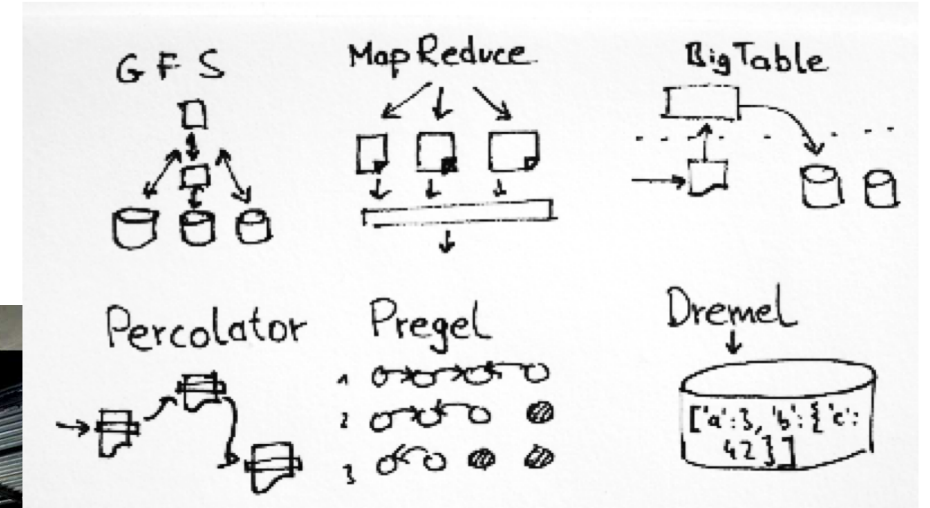
Everything breaks  
All the time

# Intro

- After IT related jobs;
- Dutch Police – enterprise IT
- OCOM group - hosting & cloud
- SDL – Software As A Service
- Spoken Communications – Call Center as Service


# History of Cloud backend

- Behind the Website and Apps;
- Running in to scale problems
  - Cost
  - Complexity
  - Speed



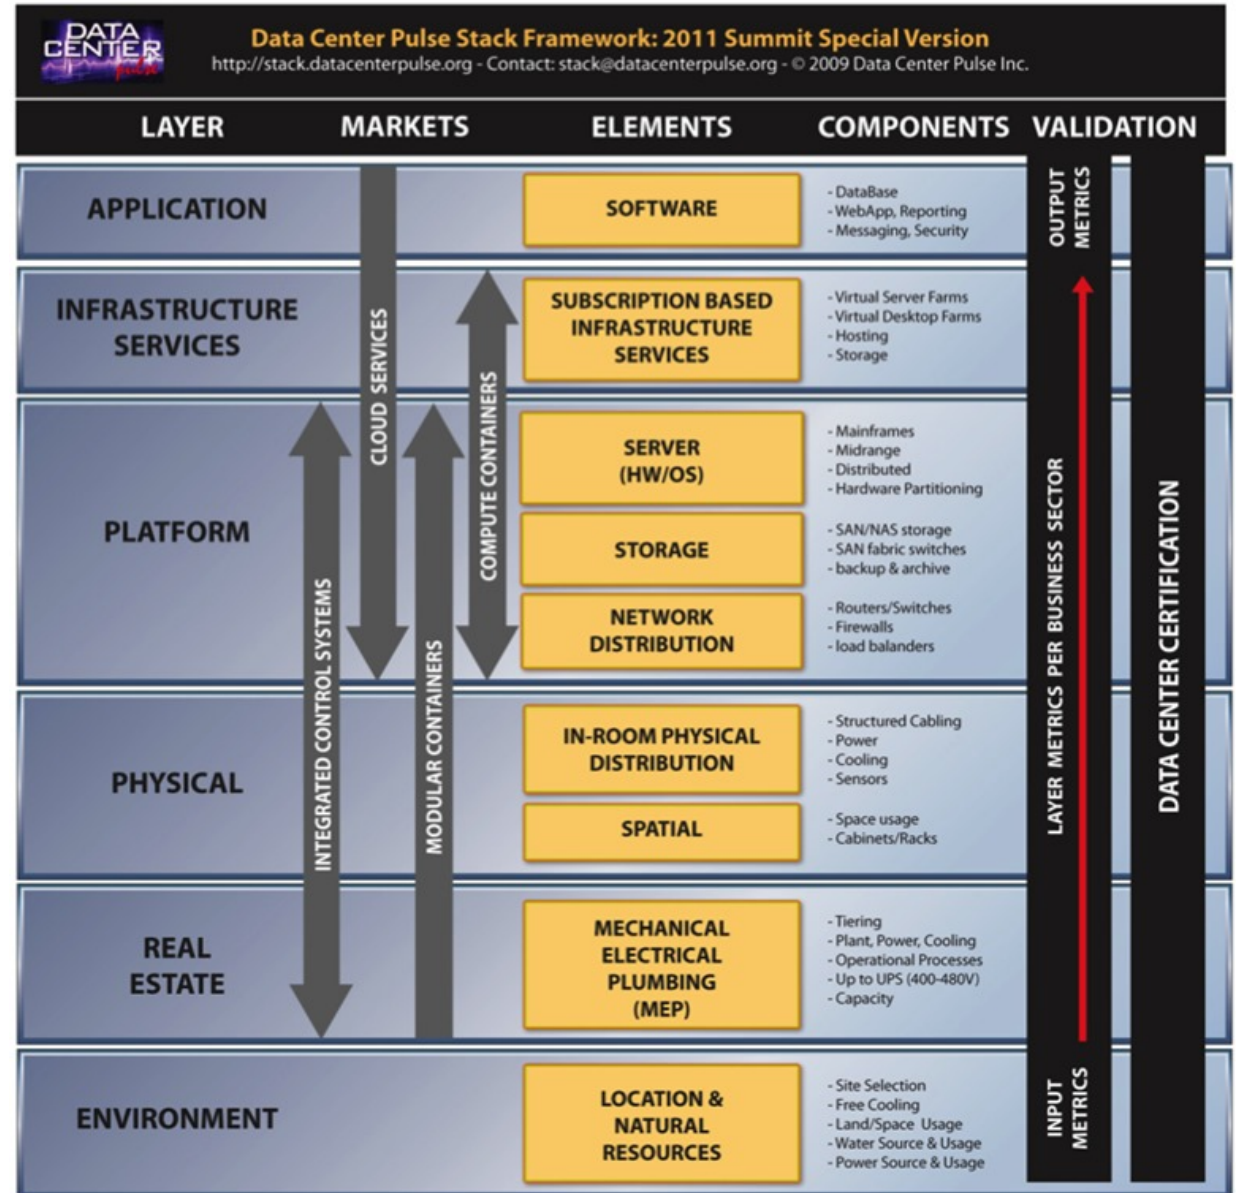
# Running at scale

- The 'normal' things didn't work any more.
- New server hardware, networking, software
- New datacenters
- New economics



# Running at scale



- Many layers, abstracted away



# Large scale – everybody's problem now

- For cloud providers – millions of servers
- For cloud costumers – millions of instances

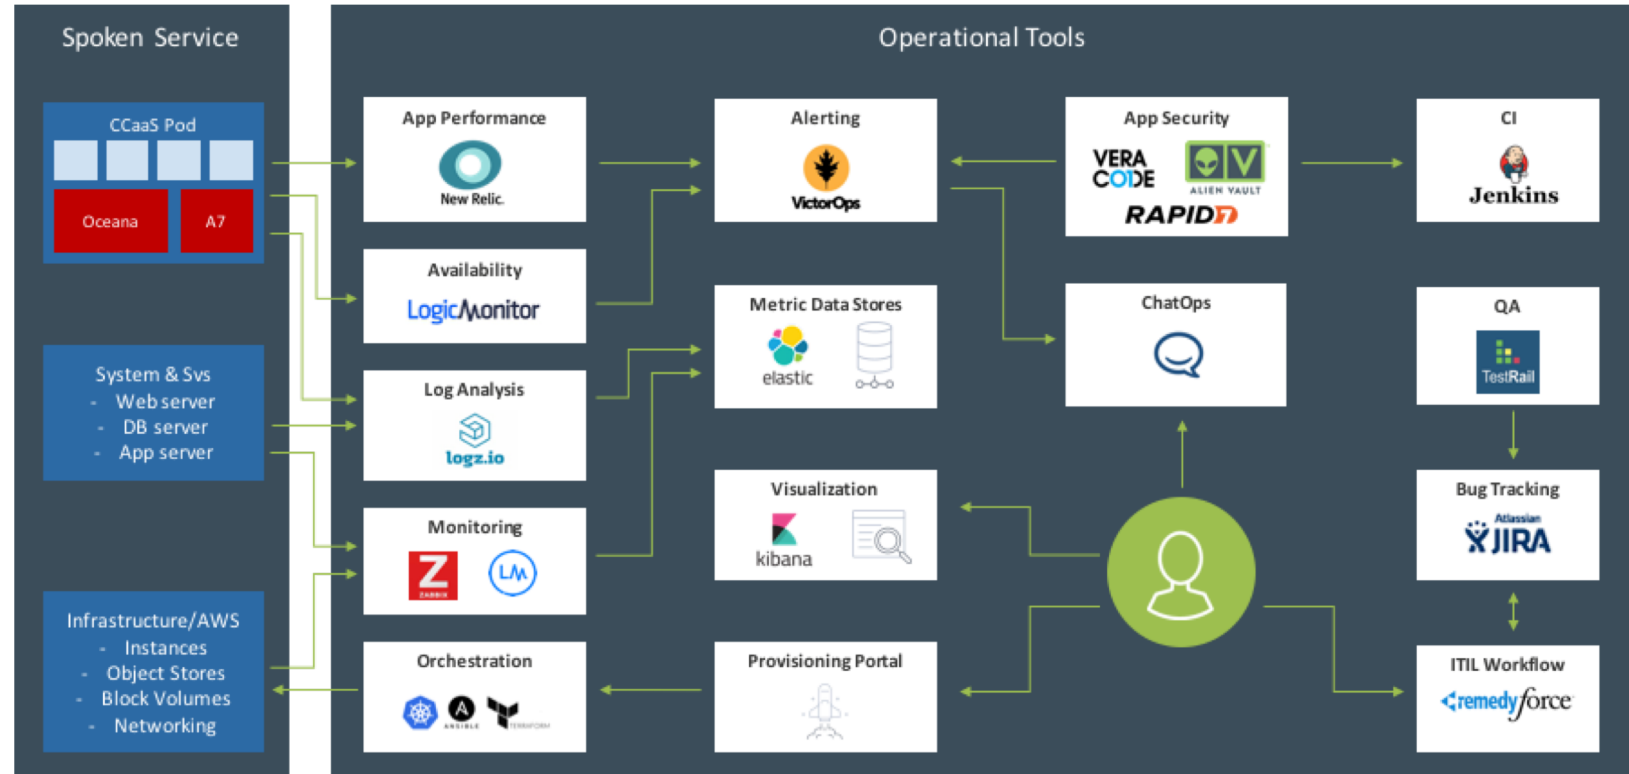
Jevons Paradox will apply to 'Cloud Computing': Consumption of the processing power will increase (rather than decrease)



*Technological progress that increases the efficiency with which a resource is used tends to increase (rather than decrease) the rate of consumption of that resource*

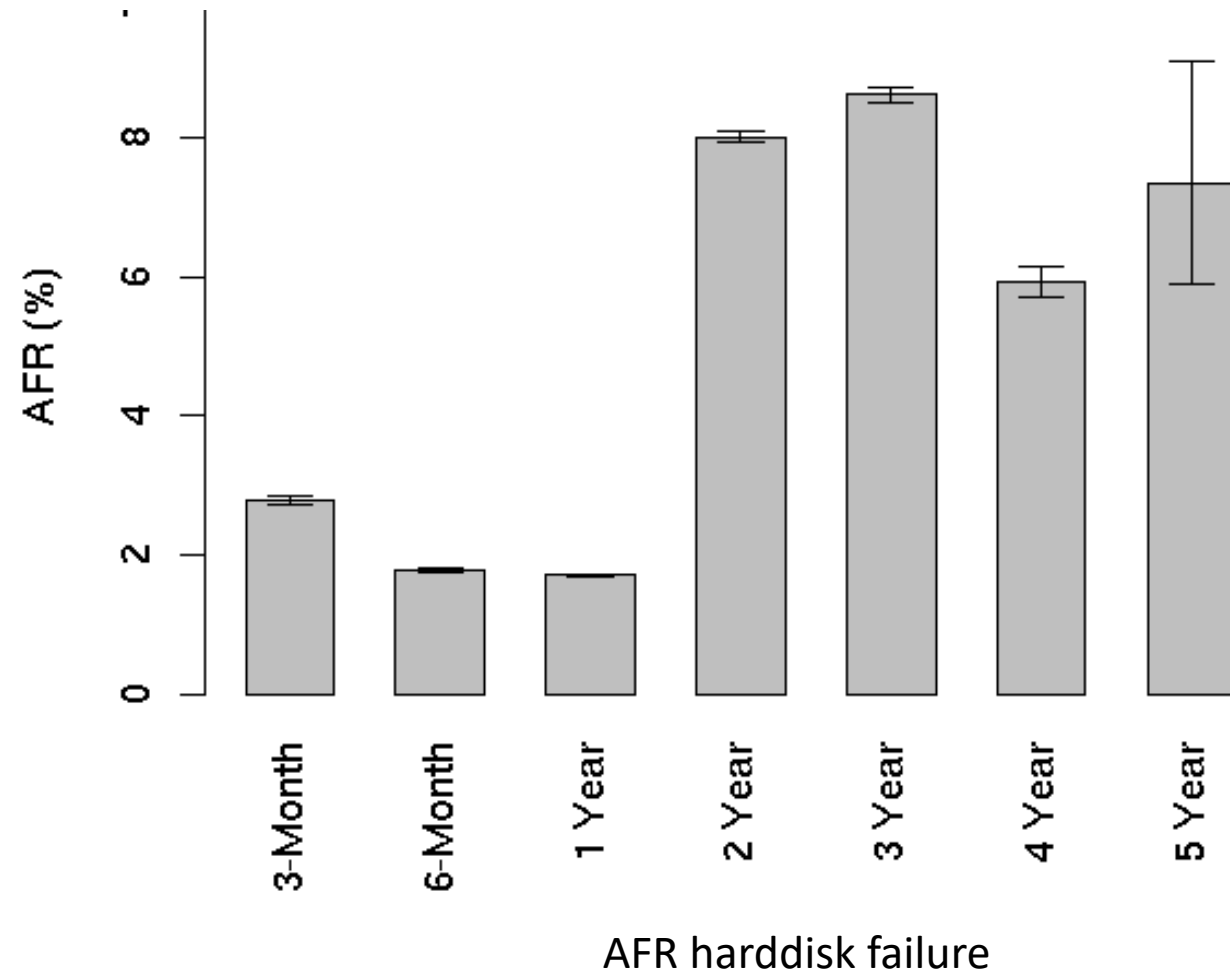
# Automation

- Large scale;
  - Deployment
  - Config
  - CI/CD
  - Capacity mgmt
  - Monitoring



# Everything fails... all the time...


- Hardware will fail.
- Software will have bugs
- Force Majeure will hit
  
- **Focus on restore not on preventing**





# Everything fails... all the time...

- Force Majeure will hit; Amazon datacenter failure 2010



# Everything fails... all the time...

- Human error; Amazon S3 failure 2017



[Amazon Web Services](#) » Service Health Dashboard

## **[RESOLVED] Increased Error Rates for Amazon S3**

**Update at 2:08 PM PST:** As of 1:49 PM PST, we are fully recovered for operations for adding new objects in S3, which was our last operation showing a high error rate. The Amazon S3 service is operating normally.


**Update at 1:12 PM PST:** S3 object retrieval, listing and deletion are fully recovered now. We are still working to recover normal operations for adding new objects to S3.

**Update at 12:52 PM PST:** We are seeing recovery for S3 object retrievals, listing and deletions. We continue to work on recovery for adding new objects to S3 and expect to start seeing improved error rates within the hour.

**Update at 11:35 AM PST:** We have now repaired the ability to update the service health dashboard. The service updates are below. We continue to experience high error rates with S3 in US-EAST-1, which is impacting various AWS services. We are working hard at repairing S3, believe we understand root cause, and are working on implementing what we believe will remediate the issue.

# Balancing cost and risk

Level of Availability	Percent of Uptime	Downtime per Year	Downtime per Day
1 Nine	90%	36.5 days	2.4 hrs.
2 Nines	99%	3.65 days	14 min.
3 Nines	99.9%	8.76 hrs.	86 sec.
4 Nines	99.99%	52.6 min.	8.6 sec.
5 Nines	99.999%	5.25 min.	.86 sec.
6 Nines	99.9999%	31.5 sec.	8.6 msec



- Cost goes up
- Complexity goes up

# Outage !

- What happens ?
  - Typical Process (ITIL or similar)
  - Typical Human behaviour
- What happens after ?
  - Review process, retro or post mortem
  - Typical Human behaviour

# Retro's

- When ?
- Who ?
- Structure:
  - Detailed Timeline
  - Root cause
  - What went well
  - What could we have done better
  - Action items
- Automate the input (ChatOps)
- Importance of Blameless Retro's.
- Share retro output internal and external

# Materials / references

- Barraso, Clidas, Holzle: The Datacenter as a Computer. 2009/2013
- Mark Burgess: In Search of Certainty. 2013
- Woods DD. STELLA: Report from the SNAFUcatchers Workshop on Coping With Complexity. 2017.