


# Cloud Storage FUD

Alyssa Henry  
General Manager  
Amazon S3



# Amazon S3: Storage for the Internet

Billions of Objects Stored



# Design Goals

“In life, as in football, you won’t go far unless you know where the goalposts are.”

Arnold H. Glasgow

# Durable

Won't lose or corrupt objects



# Available

Always on

No planned downtime

Engineer for 99.99%



# Scalable

Virtually infinite

Support an unlimited number of web-scale apps

Use scale as an advantage



# Secure



Secure protocols

Authentication mechanisms

Access controllable, log-able

# Fast

Support high performance apps

S3 latency insignificant relative to Internet latency

Reduce Internet latency by adding new locations





# Simple



Self-service

Straightforward API

Few concepts to learn

# Cost Effective



Pay as you go

Pay only for what is used

No long-term contracts or commitments

Use software and scale to reduce costs

# Uncertainty

“Everything is vague to a degree  
you do not realize till you have  
tried to make it precise.”

Bertrand Russell

# What Don't We Know?

Customer usage consistent or changing over time

Predominant workload type

Object access frequency

Object access volume

Object access locality

Object lifetime

Object size



# Uncertainty Is Certain

Inherent in general purpose systems

Use cases varied

May change over time

May change suddenly

Have to make assumptions

# Failure

“Try again. Fail Again. Fail better”

Samuel Beckett

# What Are The Odds?

Many failures happen frequently

Even low probability events happen at high scale



# Failure Happens

Natural disasters destroy data centers

Load balancers corrupt packets

Technicians pull live fiber

Routers black hole traffic

Power and cooling fails

NICs corrupt packets

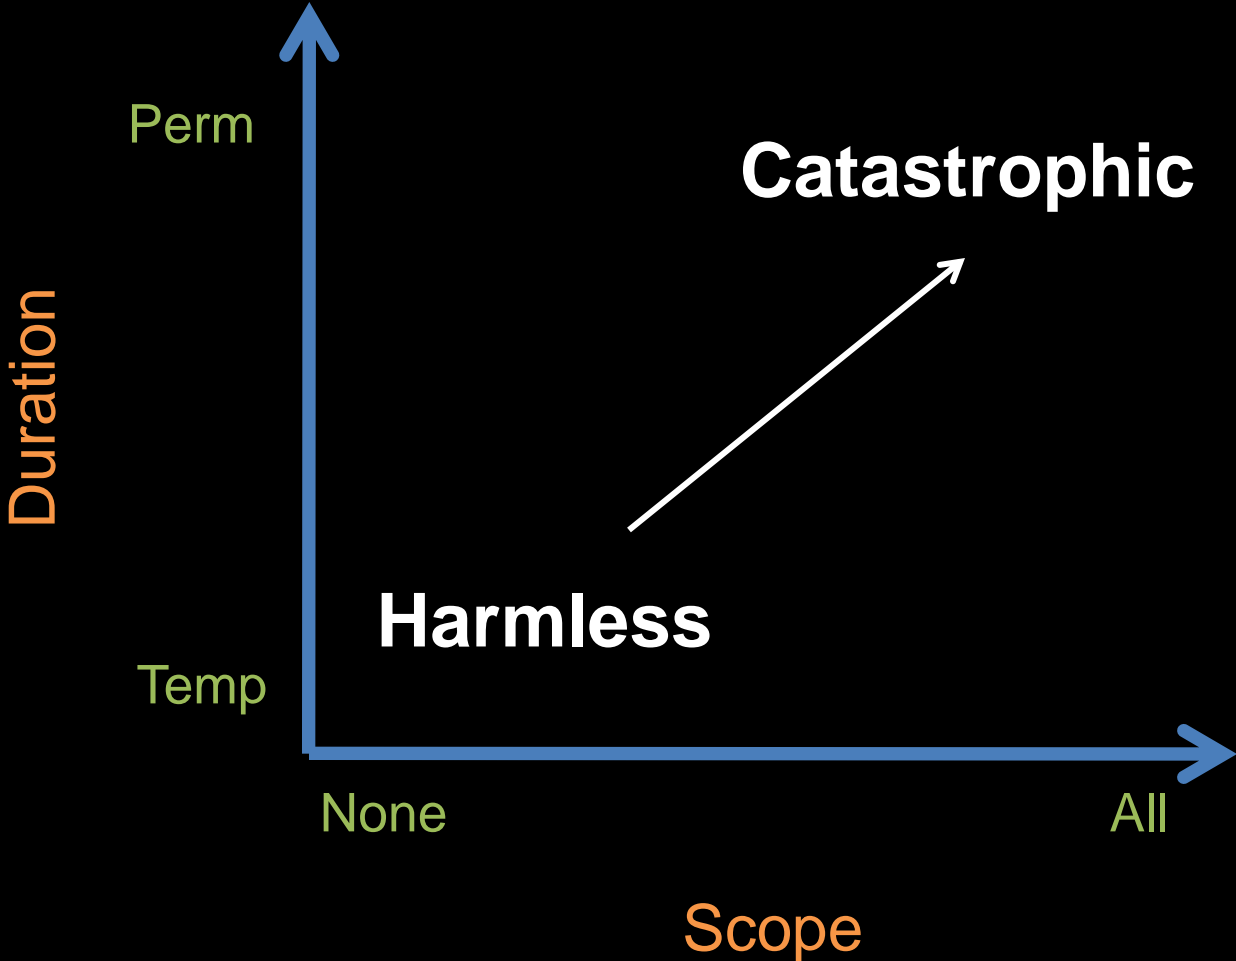
Disk drives fail

Bits rot





# Failure Types



# Techniques

“Do not let what you cannot do interfere with what you can do.”

John Wooden

# Redundancy



Broadly applicable technique

Increases durability, availability, cost, complexity

Seat belt & air bag vs. belt & suspenders

Plan for catastrophic loss of entire data center

# Retry

Resolves temporal failures

Real-time or later date

Leverage redundancy

Idempotency



**LATHER, RINSE, REPEAT**

# Surge Protection

Rate limiting

Exponential back off

Cache TTL extension



# Eventual Consistency



Spectrum of choices

Time lapse typically result of node failure

Sacrifice some consistency for availability

Sacrifice some availability for durability

# Routine Failure

Failure of components is normal

Routinely fail disks, servers, data centers



# Diversity

Software

Hardware

Workloads





# Integrity Checking



Identifies corruption inbound, outbound, at rest  
Increases cost, complexity for the customer  
Increases durability, availability

# Telemetry

Internal, external

Real time, historical

Per host, aggregate



# Autopilot

Human processes fail

Human reaction time is slow



# Summary

# Design Goals

Durable

Available

Scalable

Secure

Fast

Simple

Cost Effective



# Techniques

Redundancy

Retry

Surge Protection

Eventual Consistency

Routine Failure

Diversity

Integrity Checking

Telemetry

Autopilot



# Final Thoughts

Storage is a lasting relationship

Requires trust

Reliability at low cost achieved through  
engineering, experience, and scale



# More Information

Amazon S3

<http://aws.amazon.com/s3>

Amazon Web Services blog

<http://aws.typepad.com>

Werner Vogel's blog

<http://www.allthingsdistributed.com>

Email me directly

[ahenry@amazon.com](mailto:ahenry@amazon.com)



**Thank You!**