Data Integrity in Stateful Services

Velocity, China, 2016

Data Integrity

Bringing Sexy Back

“Protect the Data.”

-Every DBA who doesn’t want to be fired

Breaking Integrity Down

- Physical Integrity - Help, my data files are gone!
- Logical Integrity - Help, my emails disappeared!

"Data Integrity is the mission of the entire organization."

When to plan for integrity?

If you are planning for recovery after your application and infrastructure is built…

You're guarding a henhouse that already has a fox in it.

Goals for Integrity

Elimination

Empowerment

Detection

Flexibility

Elimination

Where possible, eliminate the potential for corruption and data loss.

Optimize for durability based on your user needs

ACID vs BASE

Consistency vs Availability

Velocity Levers

Empowerment

Help people and systems to recover rapidly from their own mistakes.

don’t trust destructive requests

soft deletes w/recovery API

data versioning

Detection

Early detection of corruption is as important as the ability to recover from it.

unit and regression testing

data validation pipelines

tools for investigation

Flexibility

You cannot predict all of the ways you can lose data. Focus on flexibility in your toolbox.

Tiered Storage

Replication and Data Portability

What could go wrong?

Famous Last Words

Planned Recovery

- Production deployments
- Environment duplication
- Downstream services
  - (analytics, compliance)
- Operational tests

Unplanned Recovery

"Google estimates 24 combinations of data integrity failures possible"

Category

Scope

Impact

Scenario Scope

- Small:
  - Localized or single instance in redundant scenarios.
  - Small subset of data (1000 customers)
- Medium:
  - Cluster-Wide or a full Zone
  - A full dataset (all customers in a shard)
- Large:
  - Multiple clusters, or a full DC
  - Multiple datasets (full data loss, all customers across shards)

Scenario Impact

- Small:
  - Some features impacted, non-SLO threatening.
  - Small subset of users impacted.
- Medium:
  - SLO threatening.
  - Moderate subset of users impacted.
- Large:
  - SLO impacting, application down.
  - Majority of users impacted.

Operator Error

- Data Deletion
- Data Corruption
- Relaxed Constraints
- Storage removal

Application Errors

- Removing pointers to assets in external storage
- Character set mutilation
- Duplication of data

Infrastructure Services

- Orchestration got frisky?
- Configuration management change some durability parameters?
- Proxies or DNS points to the wrong node?

OS and Hardware Errors

- Silent corruption due to failed ECC error checks?
- Filesystem corruption
- Data loss during a power down

Hardware Failures

- Disk Failures
- Memory Failures
- Controller Failures

The Iconic image of buildings at Pruitt-Igoe being imploded, as seen in THE PRUITT-IGOE MYTH, a film by Chad Freidrichs. A First Run Feature release.
Photo Credit: State Historical Society of Missouri.

Datacenter Failures

- Catastrophic power loss
- Wiped out storage
- Fires, catastrophic events

Anatomy of a Recovery Strategy

Building Block 1

"Early Detection, Bad Data Propagates"

- A culture of unit and regression testing
- Data validation test suite
  - Example: Storing external media
- Tools and analytics to investigate errors

Building Block 2

"Tiered Storage"

- Fast, expensive storage for dataset portability
- Slow, inexpensive storage for long-term backups
- Long-term storage (tape, offsite)
- Object storage for versioning
- Distributed logs (i.e. Kafka) for versioning

Building Block 3

"Toolbox"

- Full and incremental online backups
- Full and incremental offline/long-term backups
- APIs for soft deletion/undeletion
- APIs for version rollback/play forward
- Producers for event streams to recreate objects

Building Block 4

"Testing"

- Daily use as testing - incorporate recovery into daily work
- Continuous testing of less-used recovery methods
- Regular game days - team scenario testing

Final Considerations

Data Integrity is cultural

- Design, Build, Test, Deploy - each stage is an opportunity to think about data integrity
- Checks and balances between teams keeps us honest and focused on the goal
- Data becomes too complex for one person or team to understand it. We must help each other.

Data Integrity Must be Continuous

- These processes are crucial and cannot be allowed to gather dust.
- Humans will not do this on their own, integration and automation is required.
- This must be put into project functional requirements.

You cannot plan for everything

- New and interesting things will occur that will challenge your plans.
- Flexibility and multiple options must be made available.
- Early detection is crucial for ensuring problems do not propagate out of control.

Good luck!

“@lainevcampbell, laine@opsartisan.com”

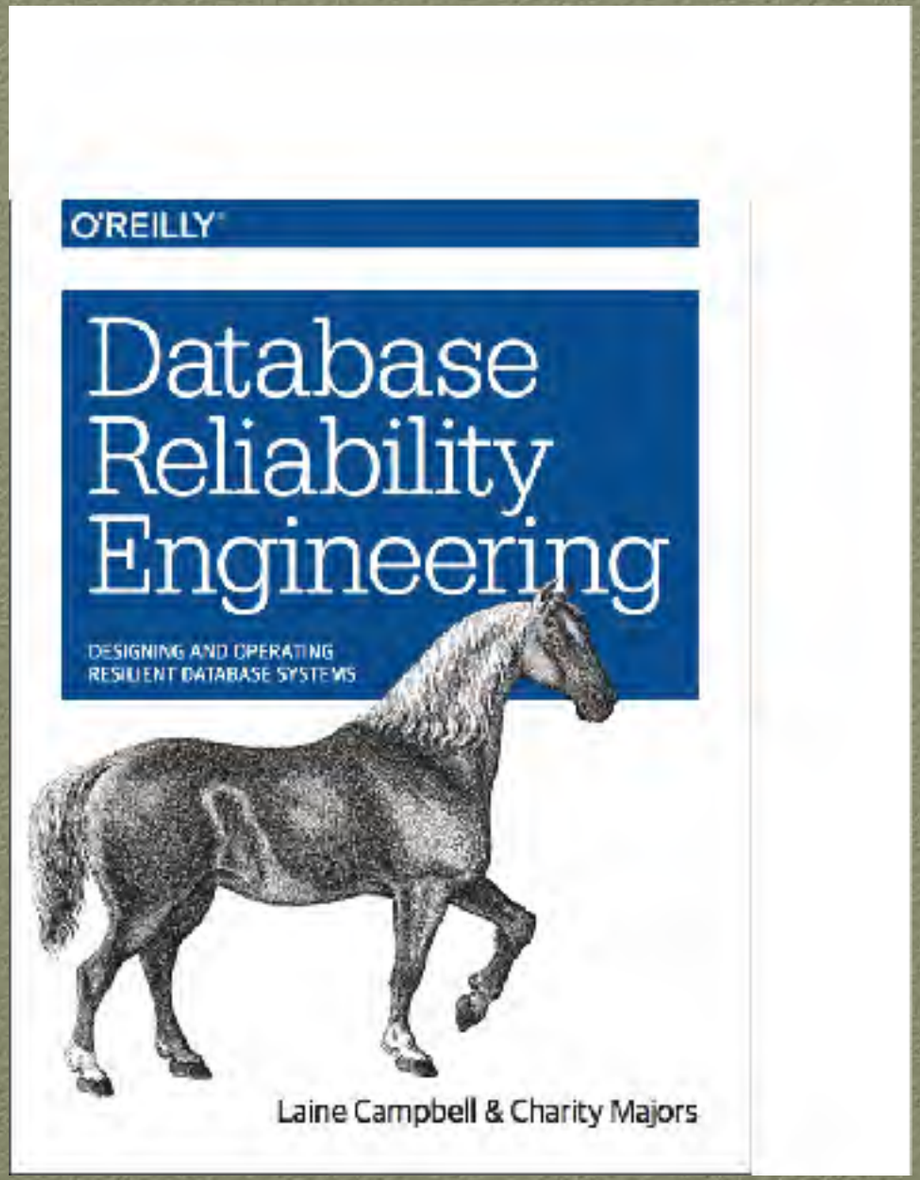

Check out our book

"Laine Campbell and Charity Majors"