


# Large-Scale Architecture

## The Unreasonable Effectiveness of Simplicity

Randy Shoup  
@randyshoup

# Background



STITCH FIX™

**we work**

# eBay Architecture



- 1995 Monolithic Perl
- 1996-2002 v2
  - Monolithic C++ ISAPI DLL
  - 3.4M lines of code
  - Compiler limits on number of methods per class
- 2002-2006 v3 Migration
  - Java “mini-applications”
  - Shared databases
- 2012 v4 Microservices
- 2017 v5 Microservices

# Amazon Architecture



- 1995-2001 Obidos
  - Monolithic Perl / Mason frontend over C backend
  - ~4GB application in a 4GB address space
  - Regularly breaking Gnu linker
  - Restarting every 100-200 requests for memory leaks
  - Releasing once per quarter
- 2001-2005 Service Migration
  - Services in C++, Java, etc.
  - No shared databases
- 2006 AWS launches

No one starts with microservices

...

Past a certain scale, everyone ends  
up with microservices

# Large-Scale Architecture



- Simple Components



- Simple Interactions



- Simple Changes



- Putting It All Together

# Large-Scale Architecture



- Simple Components



- Simple Interactions



- Simple Changes



- Putting It All Together

# Modular Services


- Service boundaries match the problem domain
- Service boundaries encapsulate business logic and data
  - All interactions through published service interface
  - Interface hides internal implementation details
  - No back doors
- Service boundaries encapsulate architectural -ilities
  - Fault isolation
  - Performance optimization
  - Security boundary






# Orthogonal Domain Logic

- Stateless domain logic
  - Ideally stateless pure function
  - Matches domain problem as directly as possible
  - Deterministic and testable in isolation
  - Robust to change over time
- “Straight-line processing”
  - Straightforward, synchronous, minimal branching
- Separate domain logic from I/O
  - Hexagonal architecture, Ports and Adapters
  - Functional core, imperative shell




# Sharding

- Shards partition the service's “data space”
  - Units for distribution, replication, processing, storage
  - Hidden as internal implementation detail
- Shards encapsulate architectural -ilities
  - Resource isolation
  - Fault isolation
  - Availability
  - Performance
- Shards are autoscaled
  - Divide or scale out as processing or data needs increase
  - E.g., DynamoDB partitions, Aurora segments, Bigtable tablets




# Service Layering

- Common services provide and abstract widely-used capabilities
- Service ecosystem
  - Services call others, which call others, etc.
  - Graph, not a strict layering
- Services grow and evolve over time
  - Factor out common libraries and services as needed
  - Teams and services split like “cellular mitosis”



# Common Platform


- “Paved Road”
  - Shared infrastructure
  - Standard frameworks
  - Developer experience
  - E.g., Netflix, Google
- Separation of Concerns
  - Reduce cognitive load on stream-aligned teams
  - Bound decisions through enabling constraints




Large-scale organizations often invest more than 50% of engineering effort in platform capabilities

@randyshoup

# Evolving Services



- Variation and Natural Selection
  - Create / extract new services when needed to solve a problem
  - Services justify their continued existence through usage
  - Deprecate services when they are no longer used
- Domains grow and divide over time



“Every service at Google is either deprecated or not ready yet.”

# Large-Scale Architecture



- Simple Components



- Simple Interactions



- Simple Changes




- Putting It All Together



# Event-Driven


- Communicate state changes as stream of events
  - Statement that some interesting thing occurred
  - Ideally represents a semantic domain event
- Decouples domains and teams
  - Abstracted through a well-defined interface
  - Asynchronous from one another
- Simplifies component implementation



# Services and Events

- Events are a first-class part of a service interface
- A service interface includes
  - Synchronous request-response (REST, gRPC, etc)
  - Events the service produces
  - Events the service consumes
  - Bulk reads and writes (ETL)
- The interface includes *any mechanism for getting data in or out of the service (!)*

# Immutable Log



- Store state as immutable log of events
  - Event Sourcing
- Often matches domain
  - E.g., Stitch Fix order processing / delivery state
- Log encapsulates architectural –ilities
  - Durable
  - Traceable and auditable
  - Replayable
  - Explicit and comprehensible
- Compact snapshots for efficiency


# Immutable Log

- Example: Stitch Fix order states
  - History of the states of a order

Request a fix	<i>fix_scheduled</i>
Assign fix to warehouse	<i>fix_hizzy_assigned</i>
Assign fix to a stylist	<i>fix_stylist_assigned</i>
Style the fix	<i>fix_styled</i>
Pick the items for the fix	<i>fix_picked</i>
Pack the items into a box	<i>fix_packed</i>
Ship the fix via a carrier	<i>fix_shipped</i>
Fix travels to customer	<i>fix_delivered</i>
Customer decides, pays	<i>fix_checked_out</i>


# Embrace Asynchrony

- Decouples operations in time
  - Decoupled availability
  - Independent scalability
  - Allows more complex processing, more processing in parallel
  - Safer to make independent changes
- Simplifies component implementation



# Embrace Asynchrony

- Invert from synchronous call graph to async dataflow
  - Exploit asymmetry between writes and reads
  - Can be orders of magnitude less resource intensive



# Large-Scale Architecture



- Simple Components



- Simple Interactions



- Simple Changes



- Putting It All Together

# Incremental Change

- Decompose every large change into small incremental steps
- Each step maintains backward / forward compatibility of data and interfaces
- Multiple service versions commonly coexist
  - Every change is a rolling upgrade
  - *Transitional states are normal, not exceptional*



# Continuous Testing


- Tests help us go faster
  - Tests are “solid ground”
  - Tests are the safety net
- Tests make better code
  - Confidence to break things
  - Courage to refactor mercilessly
- Tests make better systems
  - Catch bugs earlier, fail faster

@randyshoup



# Developer Productivity

- 75% reading existing code
- 20% modifying existing code
- 5% writing new code




<https://blogs.msdn.microsoft.com/peterhal/2006/01/04/what-do-programmers-really-do-anyway-aka-part-2-of-the-yardstick-saga/>

@randyshoup

# Developer Productivity

- 75% reading existing code
- 20% modifying existing code
- 5% writing new code



<https://blogs.msdn.microsoft.com/peterhal/2006/01/04/what-do-programmers-really-do-anyway-aka-part-2-of-the-yardstick-saga/>

@randyshoup

# Continuous Testing

- Tests make better designs
  - Modularity
  - Separation of Concerns
  - Encapsulation

@randyshoup



“There’s a deep synergy between testability and good design. All of the pain that we feel when writing unit tests points at underlying design problems.”

-- Michael Feathers

# Test-Driven Development



STITCH FIX™

- ➔ Basically no bug tracking system (!)
  - “Inbox Zero” for bugs
  - Bugs are fixed as they come up
  - Backlog contains features we want to build
  - Backlog contains technical debt we want to repay

@randyshoup

# Canary Deployments

- Staged rollout
  - Go slowly at first; go faster when you gain confidence
- Automated rollout / rollback
  - Automatically monitor changes to metrics
  - If metrics look good, keep going; if metrics look bad, roll back
- Make deployments routine and boring

# Feature Flags

- Configuration “flag” to enable / disable a feature for a particular set of users
  - Independently discovered at eBay, Facebook, Google, etc.
- More solid systems
  - Decouple feature delivery from code delivery
  - Rapid on and off
  - Develop / test / verify in production



# Continuous Delivery


- Deploy services multiple times per day
  - Robust build, test, deploy pipeline
  - SLO monitoring
  - Synthetic monitoring
- More solid systems
  - Release smaller, simpler units of work
  - Smaller changes to roll back or roll forward
  - Faster to repair, easier to understand, simpler to diagnose
  - Increase rate of change and reduce risk of change

# Continuous Delivery

- Cross-company *Velocity Initiative* to improve software delivery
  - Think Big, Start Small, Learn Fast
  - Iteratively identify and remove bottlenecks for teams
  - “What would it take to deploy your application every day?”
- Doubled engineering productivity
  - 5x faster deployment frequency
  - 5x faster lead time
  - 3x lower change failure rate
  - 3x lower mean-time-to-restore
- Prerequisite for large-scale architecture changes




# Continuous Delivery



@randyshoup

➔ 44% more time  
on features vs.  
maintenance

# Continuous Delivery



@randyshoup

➔ **2.5x more likely  
to exceed goals**

- Profitability
- Market share
- Productivity

# Large-Scale Architecture



- Simple Components



- Simple Interactions




- Simple Changes



- Putting It All Together

# System of Record

- Single System of Record
  - Every piece of data is owned by a single service
  - That service is the **canonical system of record** for that data

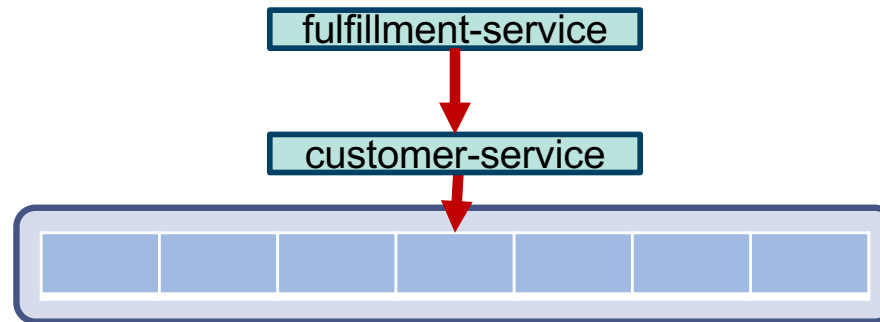


- Every other copy is a **read-only, non-authoritative cache**

# Shared Data

## *Option 1: Synchronous Lookup*


- Customer service owns customer data
- Fulfillment service calls customer service in real time



# Shared Data

## Option 2: Async event + local cache

- Customer service owns customer data
- Customer service sends `address-updated` event when customer address changes
- Fulfillment service caches current customer address



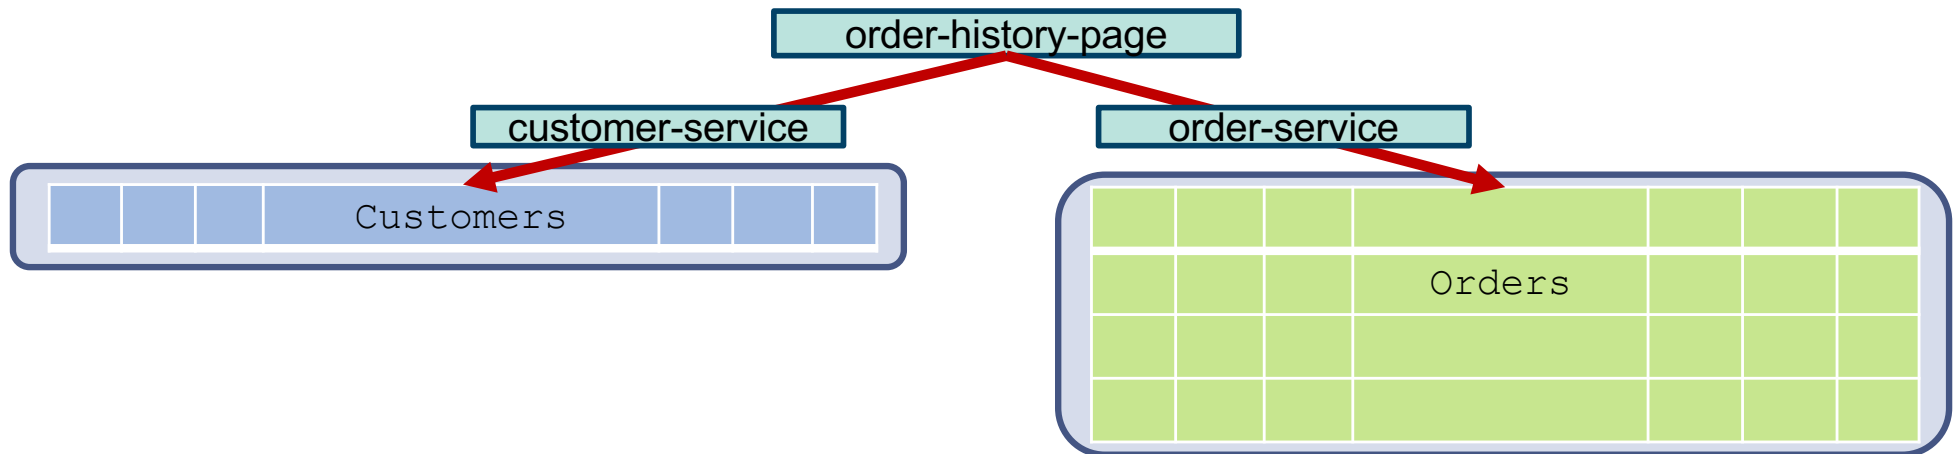
@randyshoup



# Joins

## Option 1: Join in Client Service

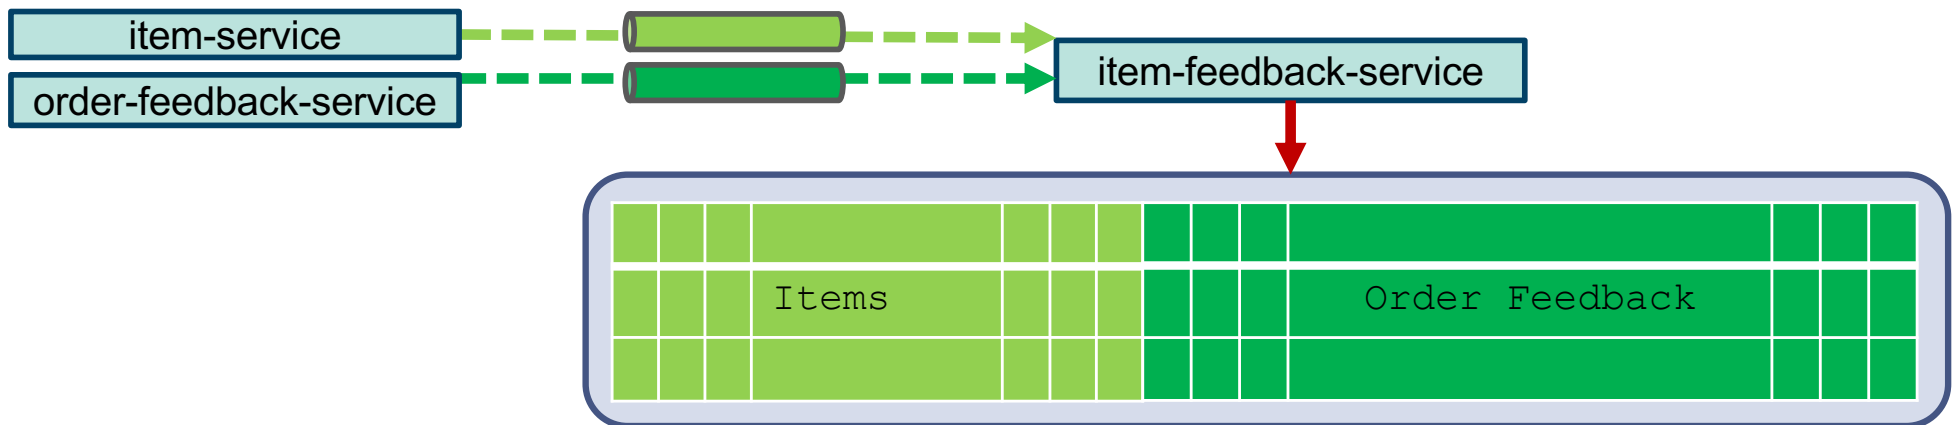
- Get a single customer from `customer-service`
- Query matching orders for that customer from `order-service`



# Joins

## Option 2: Service that “Materializes the View”


- Listen to events from `item-service`, events from `order-service`
- Maintain denormalized join of items and orders together in local storage



# Transactions

- Monolithic database makes transactions across multiple entities easy

**BEGIN; INSERT INTO A ...; UPDATE B...; COMMIT;**



- Splitting data across services makes transactions challenging

**“In general, application developers simply do not implement large scalable applications assuming distributed transactions.”**

-- Pat Helland

*Life After Distributed Transactions: An Apostate's Opinion, 2007*

“Grownups don’t use  
distributed transactions”

-- Pat Helland

# Workflows and Sagas

- Transaction → **Saga**
  - Model the transaction as a state machine of atomic events
- Reimplement as a workflow



- Roll back with compensating operations in reverse



# Workflows and Sagas

***Many real-world systems work like this***

- Payment processing
- Expense approval
- Software development process

# Intermediate States

***Explicitly expose intermediate states in the interface***

- *Payment started, pending, complete*
- *Expense submitted, approved, paid*
- *Feature developed, reviewed, deployed, released*



# Amazon Aurora

- Asynchronous redo log writes
  - Sent asynchronously to Aurora storage nodes
  - Acknowledged asynchronously to database instance
  - No distributed consensus round
  - Idempotent, immutable, monotonic
- Quorum acknowledgement
  - Log progresses forward once quorum of nodes acknowledges
- Reestablish consistency on crash recovery

Industry 3: DB Systems in the Cloud and Open Source SIGMOD'18, June 10-15, 2018, Houston, TX, USA

## Amazon Aurora: On Avoiding Distributed Consensus for I/Os, Commits, and Membership Changes

Alexandre Verbitski, Anurag Gupta, Debanjan Saha, James Corey, Kamal Gupta, Murali Brahmadassan, Raman Mittal, Sailesh Krishnamurthy, Sandor Maurice, Tengiz Kharatishvili, Xiaofeng Bao  
Amazon Web Services

**ABSTRACT**  
Amazon Aurora is a high-throughput cloud-native relational database offered as part of Amazon Web Services (AWS). One of the more novel differences between Aurora and other relational databases is how it pushes redo processing to a multi-tenant scale-out storage service, purpose-built for Aurora. Doing so reduces networking traffic, avoids checkpoints and crash recovery, enables failovers to replicas without loss of data, and enables fault-tolerant storage that heals without database involvement. Traditional implementations that leverage distributed storage would use distributed consensus algorithms for commits, reads, replication, and membership changes and amplify cost of underlying storage. In this paper, we describe how Aurora avoids distributed consensus under most circumstances by establishing invariants and leveraging local transient state. Doing so improves performance, reduces variability, and lowers costs.

**KEYWORDS**  
Databases; Distributed Systems; Log Processing; Quorum Models; Fault tolerance; Quorum Sets; Replication; Recovery; Performance

**ACM Reference Format:**  
Alexandre Verbitski, Anurag Gupta, Debanjan Saha, James Corey, Kamal Gupta, Murali Brahmadassan, Raman Mittal, Sailesh Krishnamurthy, Sandor Maurice, and Tengiz Kharatishvili. 2018. Amazon Aurora: On Avoiding Distributed Consensus for I/Os, Commits, and Membership Changes. In SIGMOD'18, 2018 International Conference on Management of Data, June 10-15, 2018, Houston, TX, USA. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/318713.319097>

### 1 INTRODUCTION

IT workloads are increasingly moving to public cloud providers such as AWS. Many of these workloads require a relational database. Amazon Relational Database Service (RDS) provides a managed service that automates database provisioning, operating system and database patching, backup, point-in-time restore, storage and compute scaling, instance health monitoring, failover, and other capabilities. Our experience managing hundreds of thousands of

database instances in RDS led to the design requirements for Aurora, a high-throughput cloud-native relational database.

In our earlier paper [12], we provided an overview of the design considerations behind Aurora. A key contribution of that paper is to show that, on a fleet-wide basis, it is insufficient to treat failures as independent. At a minimum, it is necessary to consider the correlated impact of the largest unit of failure in addition to the background noise of on-going independent failures. In AWS, the largest unit of failure a system may need to tolerate is an Availability Zone (AZ). An AZ is a subset of a Region that is connected to other AZs through low-latency networking links, but is isolated for most faults, including power, networking, software deployments, flooding, and other phenomena. Aurora supports "AZ+1" failures, resulting in six copies of data, spread across three AZs, a 4/6 write quorum, and a 3/6 read quorum as illustrated in Figure 1. Aurora implements quorum membership changes to handle unexpected failures, heat management, as well as planned software upgrades.

**Figure 1: Why are 6 copies necessary?**


Quorum models, such as the one used by Aurora, are rarely used in high-performance relational databases, despite the benefits they provide for availability, durability, and the reduction of latency jitter. We believe this is because the underlying distributed algorithms typically used in these systems – two-phase commit (2PC), Paxos commit, Paxos membership changes, and their variants – can be expensive and incur additional network overheads. The commercial systems we have seen built on these algorithms may scale well but have order-of-magnitude worse cost, performance, and peak to average latency than a traditional relational database running on a single node against local disk.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting, storing, and distributing this work for other than personal or classroom use, or for general or specific promotional, marketing, or sales promotion, or for resale, or for any other form of copying, reproduction, or redistribution, is prohibited. For more information, contact the ACM Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923, USA. <http://www.copyright.com>

SIGMOD'18, June 10-15, 2018, Houston, TX, USA  
© 2018 Copyright held by the owner(s). Publication rights licensed to the Association for Computing Machinery.  
ACM ISBN 978-1-4503-4701-3/18/06...\$15.00  
<https://doi.org/10.1145/318713.319097>

789

# Netflix Viewing History



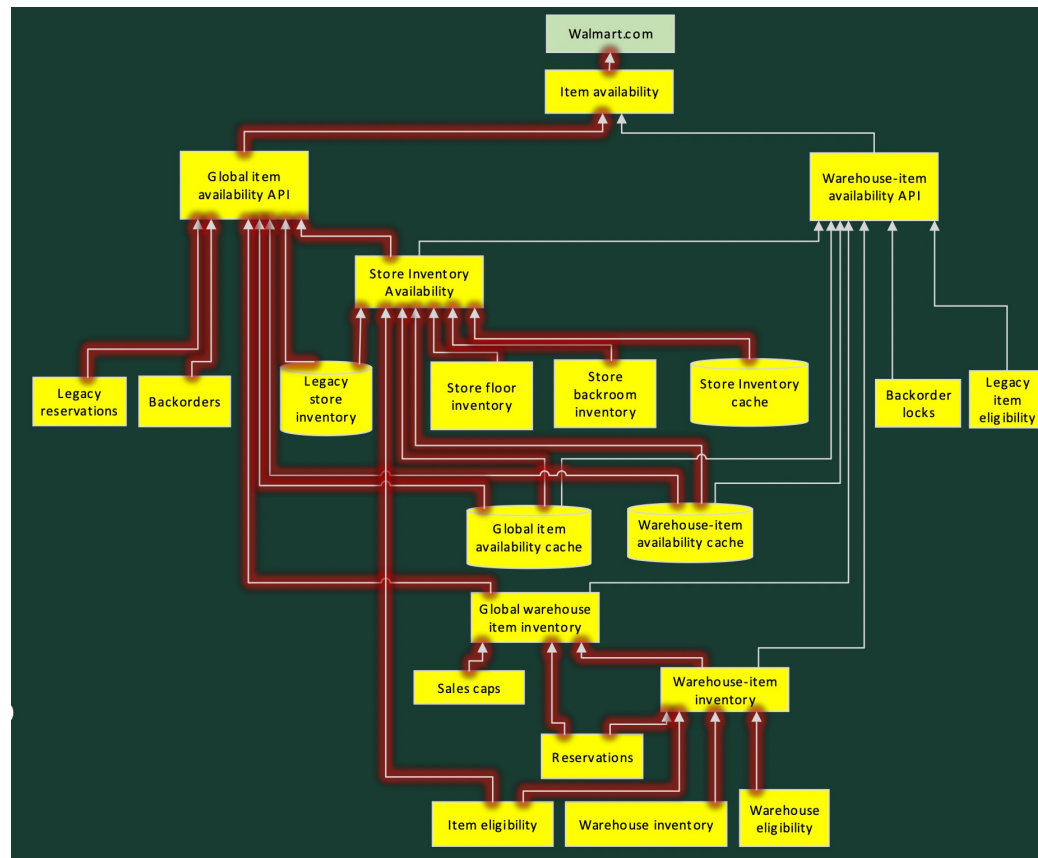
- Store and process member's playback data
  - 1M requests per second
  - Used for viewing history, personalization, recommendations, analytics, etc.
- Original synchronous architecture
  - Synchronously write to persistent storage and lookup cache
  - Availability and data loss from backpressure at high load
- Asynchronous rearchitecture
  - Write to durable queue
  - Async pipeline to enrich, process, store, serve
  - Materialize views to serve reads

# Walmart Item Availability



- Is this item available to ship to this customer?
  - Customer SLO 99.98% uptime in 300ms
- Complex logic involving many teams and domains
  - Inventory, reservations, backorders, eligibility, sales caps, etc.
- Original synchronous architecture
  - Graph of 23 nested synchronous service calls in hot path
  - Any component failure invalidates results
  - Service SLOs 99.999% uptime with 50ms marginal latency
  - Extremely expensive to build and operate

# Walmart Item Availability



@randyshoup

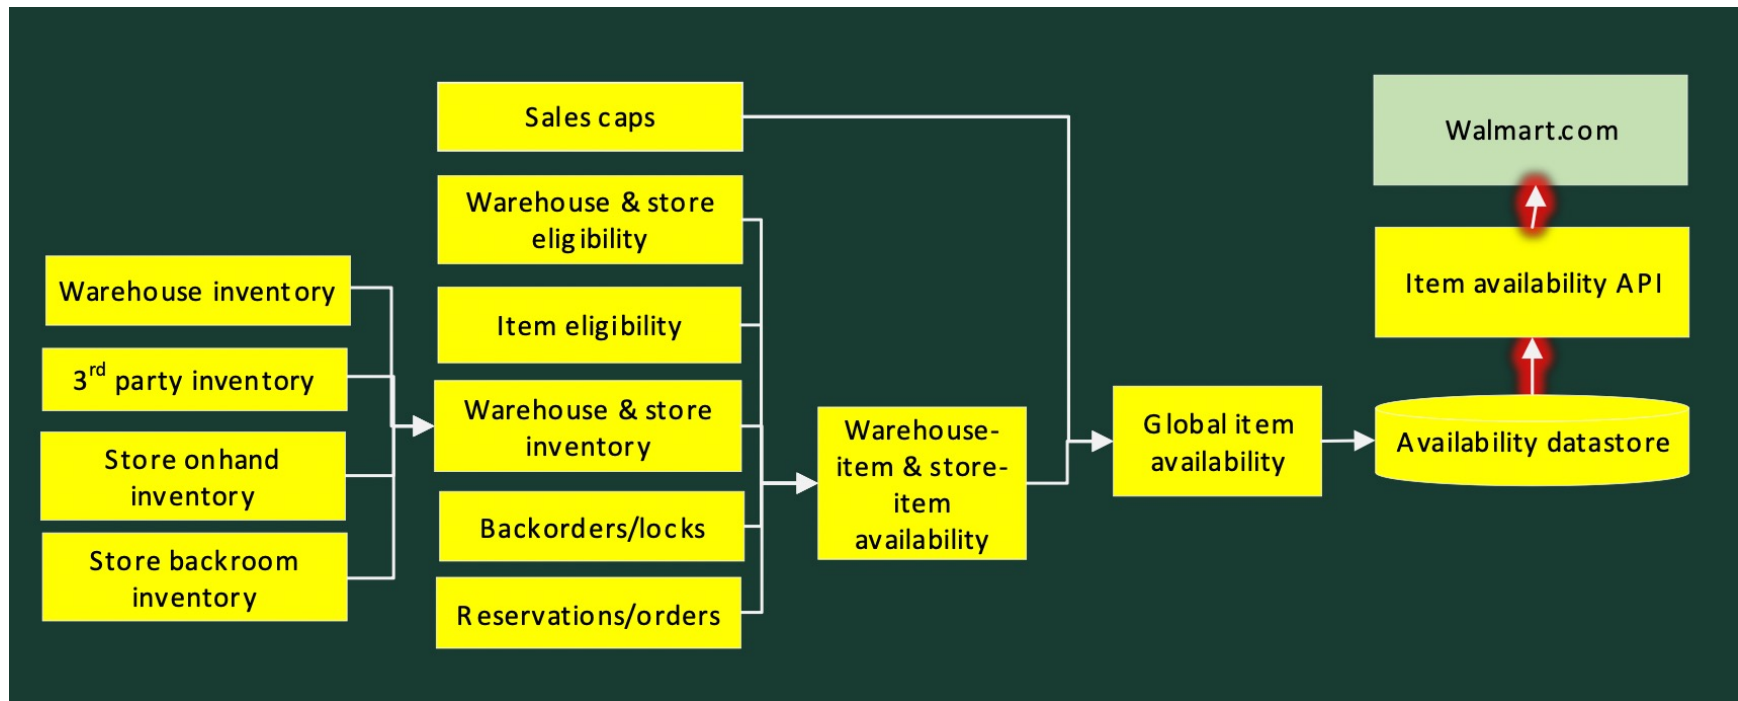
Scott Havens, 2019, [Fabulous Fortunes, Fewer Failures, and Faster Fixes from Functional Fundamentals](#), DOES 2019.

# Walmart Item Availability



- Invert each service to use async events
  - Event-driven “dataflow”
  - Idempotent processing
  - Event-sourced immutable log
  - Materialized view of data from upstream dependencies
- Asynchronous rearchitecture
  - 2 services in synchronous hot path
  - Async service SLOs 99.9% uptime with latency in seconds or minutes
  - More resilient to delays and outages
  - Orders of magnitude simpler to build and operate

# Walmart Item Availability



# Large-Scale Architecture



- Simple Components



- Simple Interactions



- Simple Changes



- Putting It All Together

# Thank you!



@randyshoup



[linkedin.com/in/randyshoup](https://www.linkedin.com/in/randyshoup)



[medium.com/@randyshoup](https://medium.com/@randyshoup)