

Serverless Distributed Ledger



Chris Anderson

Director of Product, **Fauna**

Co-founder of Couchbase,
Architect of Couchbase Mobile

@jchris



Blockchain

Blockchain

Each new block carries a signature of the previous block. If you know the current block, you can read the entire history securely. Useful for data provenance, history tracking, etc.

Combined with proof-of-work makes for an immutable log.

Blockchain



Blockchain



#noBlockchain



Ledger

10

- Cash -

1953

May 1	On hand	415 28	May 16	Phatmo	9 00
" 16	Buchanan	80 00	" 20	AB 1360	350 00
" 18	D. J. Perin	20 00	" 21	F. Hoke	30 00
" 20	Thurman W. P.	300 00	" 25	Frt	4 69
" 25	" "	75 00	" 28	Bells Pay	200 00
" 27	" "	325 00			
" 30	Thurman	10 00			
		<u>1225 28</u>			<u>593 69</u>

Ledger

A **distributed ledger** is a consensus of replicated, shared, and synchronized digital data geographically spread across multiple sites, countries, or institutions.

Distributed Ledger



Distributed Ledger







<https://animal-exchange.neocities.org/>



Items for Sale



₪12



₪100



₪10



₪7



₪2

Players

Bob

58



Carol

62



Alice

48



Purchases

₪6 🐞 Alice -> Carol

₪3 🐛 Carol -> Alice

₪12 🐱 Alice -> Bob

₪9 🐕 Alice -> Carol

₪11 🐶 Alice -> Bob

₪10 🐎 Bob -> Carol

₪12 🐆 Bob -> Alice

₪12 🐱 Carol -> Alice

₪70 🐋 Carol -> Bob

₪30 🍌 Carol -> Bob

₪3 🐛 Bob -> Carol

₪6 🐞 Bob -> Alice

₪40 🐋 Alice -> Carol

₪10 🐘 Carol -> Alice

₪6 🐞 Carol -> Bob

₪11 🐶 Bob -> Alice

₪30 🍌 Alice -> Carol

₪3 🐛 Alice -> Bob

₪70 🐋 Bob -> Carol

<https://animal-exchange.neocities.org/>



Items for Sale



12



100



10



7



2

Players

Bob

58



Carol

62



Alice

48



Purchases

6 🐞 Alice -> Carol

3 🐛 Carol -> Alice

12 🐱 Alice -> Bob

9 🐕 Alice -> Carol

11 🐶 Alice -> Bob

10 🐎 Bob -> Carol

12 🐆 Bob -> Alice

12 🐱 Carol -> Alice

70 🐋 Carol -> Bob

30 🍌 Carol -> Bob

3 🐛 Bob -> Carol

6 🐞 Bob -> Alice

40 🐋 Alice -> Carol

10 🐘 Carol -> Alice

6 🐞 Carol -> Bob

11 🐶 Bob -> Alice

30 🍌 Alice -> Carol

3 🐛 Alice -> Bob

70 🐋 Bob -> Carol

Getting Serious

Getting Serious

Global Consensus

Getting Serious

Global Consensus

Distributed Ledger Transaction

Getting Serious

Global Consensus


Distributed Ledger Transaction

Serverless Security Model

Global Consensus



Global Consensus



Global Consensus

Each ledger member
runs a high-availability
FaunaDB cluster.



Architecture

Each high-availability
FaunaDB cluster
contains a full copy of the
dataset.

This can be partitioned
for horizontal scaling.




Architecture

Each node runs FaunaDB Enterprise.

Implemented in Scala, delivered as a .jar packaged for your environment.

The cluster can be dynamically resized while serving traffic.



Calvin Protocol

<https://fauna.com/blog/distributed-consistency-at-scale-spanner-vs-calvin>

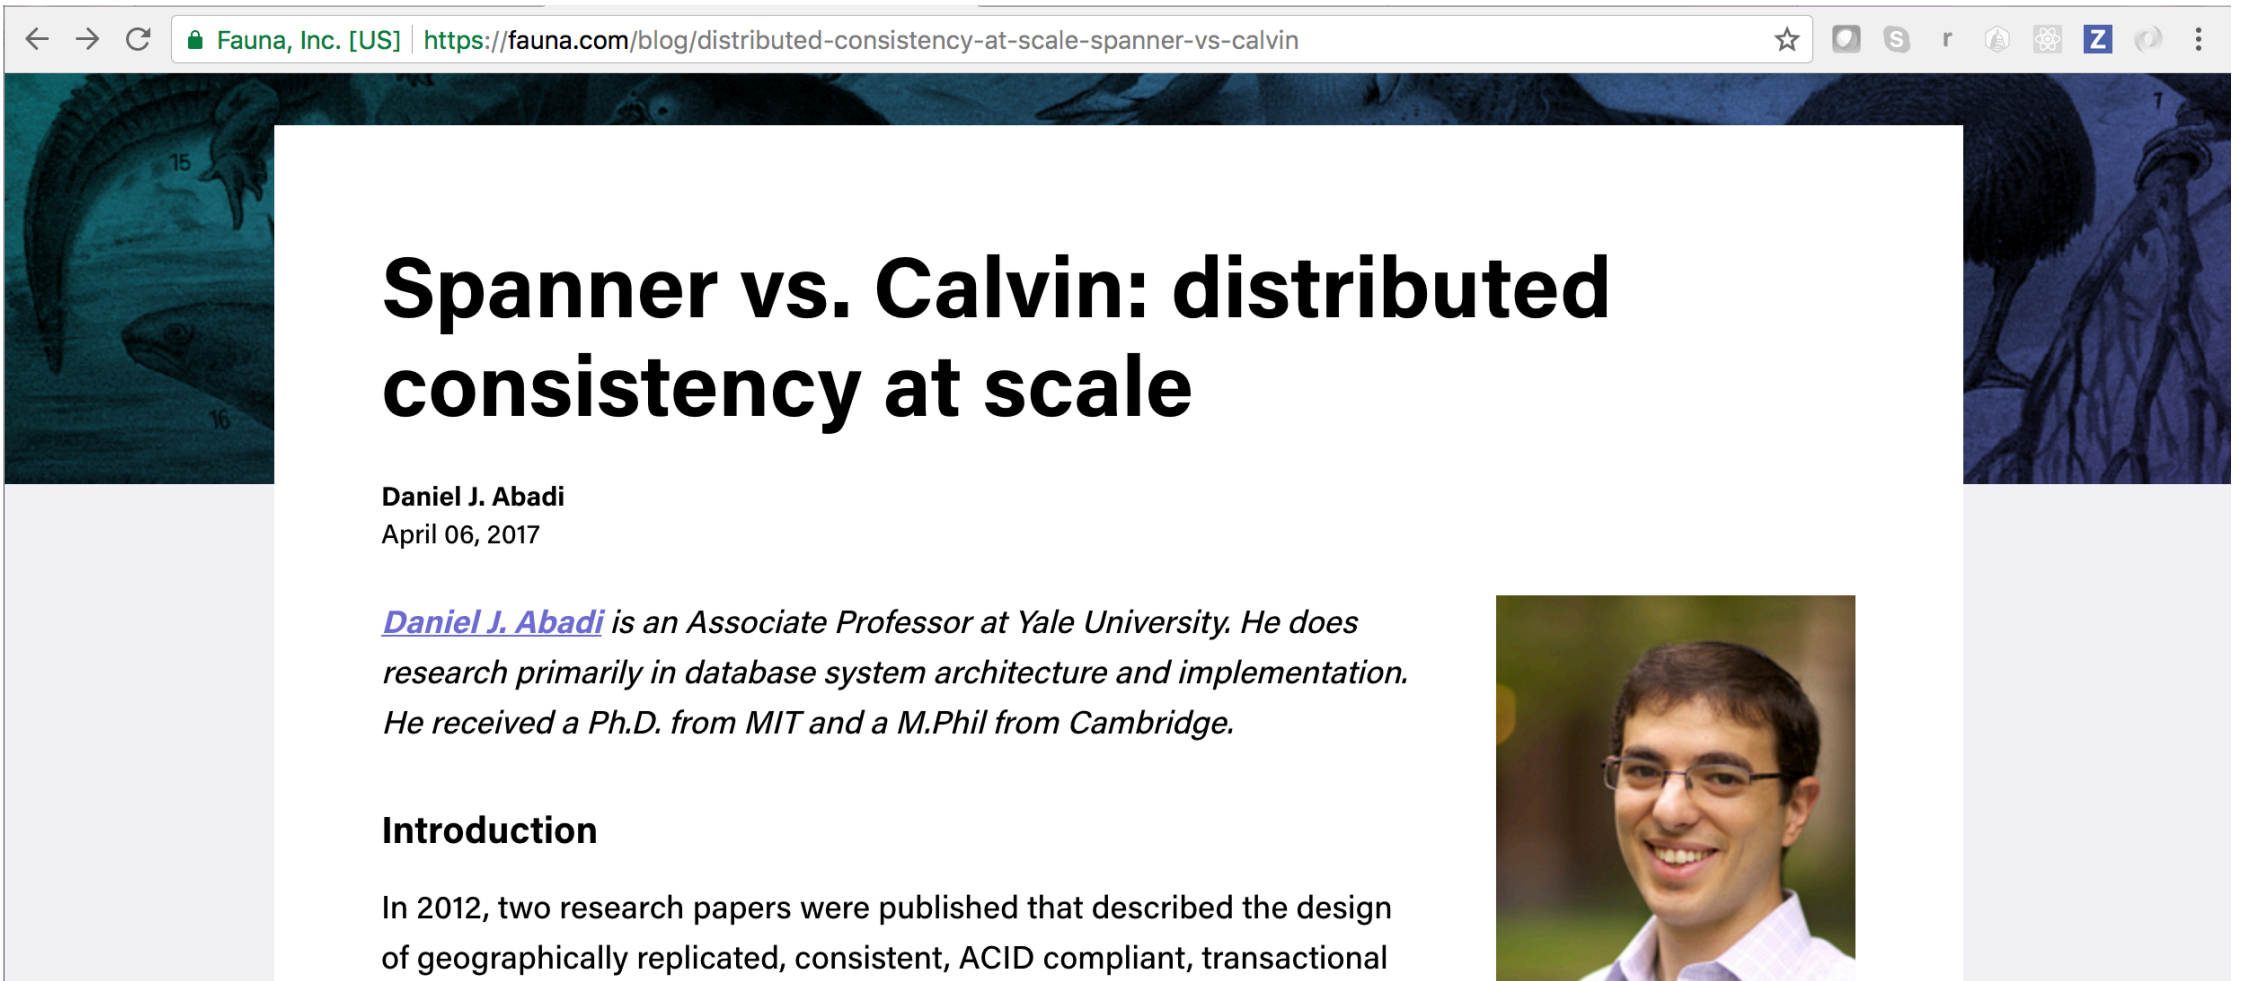
FaunaDB uses a distributed write-ahead-log to provide ACID transactions. In the presence of write conflicts transactions may be retried internally.

Transactions commit across all datacenters.

Throughput oriented, each Calvin log segment may contain multiple transactions.

Calvin Protocol

<https://fauna.com/blog/distributed-consistency-at-scale-spanner-vs-calvin>



The screenshot shows a web browser window with the address bar displaying the URL <https://fauna.com/blog/distributed-consistency-at-scale-spanner-vs-calvin>. The page features a header image with various animals, including a dinosaur and a bird. The main content area has a white background with the article title 'Spanner vs. Calvin: distributed consistency at scale' in large, bold black text. Below the title, the author's name 'Daniel J. Abadi' and the date 'April 06, 2017' are listed. A bio paragraph follows, describing Daniel J. Abadi as an Associate Professor at Yale University. To the right of the bio is a portrait of Daniel J. Abadi. The article begins with an 'Introduction' section, which starts with the text: 'In 2012, two research papers were published that described the design of geographically replicated, consistent, ACID compliant, transactional'.

← → ↻ 🔒 Fauna, Inc. [US] <https://fauna.com/blog/distributed-consistency-at-scale-spanner-vs-calvin> ☆ 📄 S r 🏠 ⚙️ Z 🔍 ⋮


Spanner vs. Calvin: distributed consistency at scale

Daniel J. Abadi
April 06, 2017

[Daniel J. Abadi](#) is an Associate Professor at Yale University. He does research primarily in database system architecture and implementation. He received a Ph.D. from MIT and a M.Phil from Cambridge.

Introduction

In 2012, two research papers were published that described the design of geographically replicated, consistent, ACID compliant, transactional



Ledger Transaction

```

client.query(
  q.Let({
    buyer : q.Get(player.ref),
    item : q.Get(item.ref)
  }, q.Let({
    isForSale : q.Select(["data", "for_sale"], q.Var("item")),
    itemPrice : q.Select(["data", "price"], q.Var("item")),
    buyerBalance : q.Select(["data", "credits"], q.Var("buyer")),
    seller : q.Get(q.Select(["data", "owner"], q.Var("item")))
  }, q.If(q.Not(q.Var("isForSale")),
    "purchase failed: item not for sale",
    q.If(q.Equals(q.Select("ref", q.Var("buyer")), q.Select("ref", q.Var("seller"))),
      // buyer = seller, remove item from sale
      q.Do(
        q.Update(q.Select("ref", q.Var("item")), {
          data : {
            for_sale : false
          }
        })
      ),
      "item removed from sale"
    ),
    // check balance
    q.If(q.LT(q.Var("buyerBalance"), q.Var("itemPrice")),
      "purchase failed: insufficient funds",
      // all clear! record the purchase, update the buyer, seller and item.

```

```
// all clear! record the purchase, update the buyer, seller and item.
q.Do(
  q.Create(q.Class("purchases"), {
    data : {
      item : q.Select("ref", q.Var("item")),
      price : q.Var("itemPrice"),
      buyer : q.Select("ref", q.Var("buyer")),
      seller : q.Select("ref", q.Var("seller"))
    }
  }),
  q.Update(q.Select("ref", q.Var("buyer")), {
    data : {
      credits : q.Subtract(q.Var("buyerBalance"), q.Var("itemPrice"))
    }
  }),
  q.Update(q.Select("ref", q.Var("seller")), {
    data : {
      credits : q.Add(q.Select(["data", "credits"], q.Var("seller")), q.Var("itemPrice"))
    }
  }),
  q.Update(q.Select("ref", q.Var("item")), {
    data : {
      owner : q.Select("ref", q.Var("buyer")),
      for_sale : false
    }
  })
),
"purchase success" ) ) ) ) ) )
```


Ledger Transaction

```
client.query(
  q.Let({
    buyer : q.Get(player.ref),
    item : q.Get(item.ref)
  }, q.Let({
    isForSale : q.Select(["data", "for_sale"], q.Var("item")),
    itemPrice : q.Select(["data", "price"], q.Var("item")),
    buyerBalance : q.Select(["data", "credits"], q.Var("buyer")),
    seller : q.Get(q.Select(["data", "owner"], q.Var("item")))
  }, q.If(q.Not(q.Var("isForSale")),
    "purchase failed: item not for sale",
    q.If(q.Equals(q.Select("ref", q.Var("buyer")), q.Select("ref", q.Var(
      // buyer = seller, remove item from sale
    q.Do(
      q.Update(q.Select("ref", q.Var("item")), {
        data : {
          for_sale : false
        }
      }
    )),
    "item removed from sale"
  ),
  // check balance
  q.If(q.LT(q.Var("buyerBalance"), q.Var("itemPrice")),
    "purchase failed: insufficient funds",
    // all clear! record the purchase, update the buyer, seller and item.
```

```
// all clear! record the purchase, update the buyer, seller and item.
q.Do(
  q.Create(q.Class("purchases"), {
    data : {
      item : q.Select("ref", q.Var("item")),
      price : q.Var("itemPrice"),
      buyer : q.Select("ref", q.Var("buyer")),
      seller : q.Select("ref", q.Var("seller"))
    }
  }),
  q.Update(q.Select("ref", q.Var("buyer")), {
    data : {
      credits : q.Subtract(q.Var("buyerBalance"), q.Var("itemPrice"))
    }
  }),
  q.Update(q.Select("ref", q.Var("seller")), {
    data : {
      credits : q.Add(q.Select(["data", "credits"], q.Var("seller")), q.Var("itemPrice"))
    }
  }),
  q.Update(q.Select("ref", q.Var("item")), {
    data : {
      owner : q.Select("ref", q.Var("buyer")),
      for_sale : false
    }
  })
),
"purchase success" ) ) ) ) ) )
```

Ledger Transaction

```
client.query(  
  q.Let(  
    buyer : q.Get(player.ref),  
    item : q.Get(item.ref)  
  ), q.Let(  
    isForSale : q.Select(["data", "for_sale"], q.Var("item")),  
    itemPrice : q.Select(["data", "price"], q.Var("item")),  
    buyerBalance : q.Select(["data", "credits"], q.Var("buyer")),  
    seller : q.Select(["data", "owner"], q.Var("item"))  
  ), q.If(q.Not(q.Var("isForSale")),  
    "purchase failed: item not for sale",  
    q.If(q.Equals(q.Select("ref", q.Var("buyer")), q.Select("ref", q.Var("seller"))),  
      // buyer = seller, remove item from sale  
      q.Update(q.Select("ref", q.Var("item")), {  
        data : {  
          for_sale : false  
        }  
      }  
    ),  
    "item removed from sale"  
  )  
)  
// Check buyer balance  
q.If(q.LT(q.Var("buyerBalance"), q.Var("itemPrice")),  
  "purchase failed: insufficient funds",  
  // all clear! record the purchase, update the buyer, seller and item.
```

Ensure item is for sale

Buyer != seller

Check buyer balance

// all clear! record the purchase, update the buyer, seller and item.

```
q.Do(  
  q.Create(q.Class("purchases"), {  
    data : {  
      item : q.Select("ref", q.Var("item")),  
      buyer : q.Select("ref", q.Var("buyer")),  
      seller : q.Select("ref", q.Var("seller"))  
    }  
  }  
),  
  q.Update(q.Select("ref", q.Var("buyer")), {  
    data : {  
      credits : q.Subtract(q.Var("buyerBalance"), q.Var("itemPrice"))  
    }  
  }  
),  
  q.Update(q.Select("ref", q.Var("seller")), {  
    data : {  
      credits : q.Add(q.Select(["data", "credits"], q.Var("seller")), q.Var("itemPrice"))  
    }  
  }  
),  
  q.Update(q.Select("ref", q.Var("item")), {  
    data : {  
      owner : q.Select("ref", q.Var("buyer")),  
      for_sale : false  
    }  
  }  
),  
  "purchase success" ) ) ) ) ) ) )
```

Write a purchase record

Deduct from buyer balance

Add to seller balance

Update item owner

Update Buyer Balance

```
q.Update(q.Select("ref", q.Var("buyer")), {  
  data : {  
    credits : q.Subtract(q.Var("buyerBalance"), q.Var("itemPrice"))  
  }  
})
```

Queries are composed on the client, and sent to the server as an abstract syntax tree encoded as JSON.

Client Library in Your Language

```
update( select('ref', var('buyer')),  
  data: {  
    credits: subtract(var('buyerBalance'), var('itemPrice'))  
  })
```

Java Javascript Scala Ruby C# Python Go Swift


ACID Transactions

Not just for distributed ledgers

Enhance developer productivity

Simplify applications

Address mission-critical use cases at scale




Serverless Security

Layered access approach.

Lambdas use keys that only have privileges to run predefined functions.

User defined functions use keys that cannot modify schema rules or old temporal snapshots.



AWS Lambda

Function as a Service

JavaScript runs in response to events.

Authenticate users, process resources,
etc.

For distributed ledger, this is the code that
reacts to user events by submitting queries
that call predefined functions.

Code can run on premise.



Predefined Function

FaunaDB user defined functions API where query fragments can be stored and executed by other queries.

Only objects with the call permission on a function can call it, so in the distributed ledger use case the Lambdas are granted keys that authenticate into the access-control graph in a place where they only have permissions to call the UDF.

```
{  
  "call": { "function": "create_entry" },  
  "arguments": [  
    "First Post Title",  
    "This is my first blog post!" ]  
}
```



Temporal Data

FaunaDB stores data in temporal snapshots, and has APIs for updating old snapshots, for instance to fix data-entry mistakes. Old snapshots are cleaned up after a configurable TTL.

<https://fauna.com/tutorials/timeline>

<https://fauna.com/blog/time-traveling-databases>

For distributed ledger, the UDFs run in a role limited to the current snapshot, so any snapshot editing can only be done from an administrative interface.





FaunaDB: Serverless Database Table Stakes

Runs in the cloud(s)

- “Not my server, not my problem, that’s what I say.” / “Around the world.”

Friendly to JSON / NoSQL


- Schema enforcement is optional, we ♥ rich nested data structures

Relational queries and constraints

- Proper database features are BACK and they SCALE

Event feeds and temporal features

- So you can build streams and triggers



FaunaDB: The First Serverless Database

What makes Fauna different!

Object level security

- Model your business rules in the database.

Escape the provisioning trap

- No need to fear traffic spikes, or pay in advance for speculatively high throughput.

Hierarchal multi-tenancy

- Makes creating new logical datasets cheap and easy. Serverless processes can scale your business without operator intervention.

Stateless client

- Your Lambdas aren't paying setup and teardown costs for nothing.

FaunaDB: The First Serverless Database

What makes Fauna different!

Object level security

- Model your business rules in the database.

Escape the provisioning trap

- No need to fear traffic spikes, or pay in advance for speculatively high throughput.

Hierarchal multi-tenancy

- Makes creating new logical datasets cheap and easy. Serverless processes can scale your business without operator intervention.

Stateless client

- Your Lambdas aren't paying setup and teardown costs for nothing.




FaunaDB - General Purpose Database

Key Features

- Transactions with global consistency
- Rich query support
- Serverless ease-of-use or on-premise
- Hierarchical Multi-tenancy


Use Cases

- Distributed Ledger
- Social Graph Content
- Single Page Applications



A simplified developer experience


- Expression-oriented, flexible, safe
- Simplify multi-part queries into simple questions:
 - In the rental car fleet, which make and model built between 2013 and 2015 has parts from manufacturers X and Y?
- Increase developer productivity
 - Isolate from complexity of different data models
 - Prevent context switches when moving among query languages and data sets
- Extensible: support data domains such as geographic indexing, full-text search, iterative machine learning.



Communal resources allotted across teams


- Safe sharing across multiple teams, projects or companies in a single FaunaDB cluster
- Dynamically tune resource allocation across tenants
- Align resource utilization with business priorities
 - Prioritize customer data over batch analytics
- Amortize infrastructure costs across multiple services

The more diverse applications, datasets, and workloads are hosted in a single FaunaDB cluster, the better the price/performance becomes compared to a traditional, statically provisioned siloed data architecture.



Shared, hierarchical database infrastructure

- Databases within databases
- Shared resources across teams, projects, applications
- Delegated administration
- Security through isolation
- Can reflect organizational structure, physical structure, etc.



A globally shared resource pool

- Native geo-replication
- Physical cluster spans all data centers
- Logical databases assigned by business priority
- No impact on operational overhead
- Increases compute elasticity
- Enables:
 - Low-latency real-time data
 - Geographical data compliance (safe harbor)

