



Applying the Saga Pattern Caitie McCafffrey



Conference: May 11-12 / Workshops: 13-14

Caitie McCaffrey Distributed Systems Engineer



@Caitie <u>CaitieM.com</u>





- Why Sagas?
- Sagas Paper
- Distributed Sagas
- Sagas in Halo 4

Systems Used to Be Simple



Serializability & ACID

SOA/Microservices



Two Phase Commit

2PC: Prepare Phase



2PC: Commit Phase



2PC Doesn't Scale

- O(n^2) Messages
- Coordinator: Single Point of Failure
- Reduced Throughput

Spanner

Google's Globally-Distributed Database Corbett et. al. "Spanner is Google's scalable, multi-version, globally distributed, and synchronouslyreplicated database. It is the first system to distribute data at global scale and support externally-consistent distributed transactions."

-Corbett et al.

"The key enabler of these properties is a new **TrueTime API** and its implementation...using multiple modern clock references (**GPS and atomic clocks**)."

-Corbett et al.

Spanner is Expensive & Proprietary

- Custom Hardware & Data-Centers
- Synchronization Not Solved

Distributed Transactions are Hard & Expensive

Can We Do Better?



Can We Do Better?



Sagas

Hector Garcia-Molina, Kenneth Salem Princeton University 1987

Sagas are Long Lived Transactions

"A Saga is a Long Lived Transaction that can be written as a sequence of transactions that can be interleaved.

All transactions in the sequence complete successfully or compensating transactions are ran to amend a partial execution."

A Saga is a Collection of Sub-Transactions

T₁, T₂ ... T_n

Each Sub-Transaction has a Compensating Transaction

C1, C2...Cn

Cn Semantically Undoes Tn

Saga Guarantee

Either

T₁, T₂ ... T_n or
T₁, T₂ ... T_j, C_j, ... C₂, C₁

Trade-Off: Atomicity for Availability

Sagas are a Failure Management Pattern

Large Single Transaction









Book Hotel (T1)



• Cancel Hotel (C1)

- Book Car (T₂)
- Book Flight (T₃)



• Cancel Car (C₂)

Cancel Flight (C₃)

Saga Execution Coordinator (SEC)

Saga Log

- Begin Saga
- End Saga
- Abort Saga

- Begin Ti
- End Ti
- Begin Ci
- End Ci



Begin Saga



Begin Saga



Start Book Hotel (T₁)

Begin Saga



Start Book Hotel (T₁)

End Book Hotel (T₁)



Begin Saga



Start Book Hotel (T1)

End Book Hotel (T₁)

Start Book Car Rental (T₂)



Begin Saga



Start Book Hotel (T1)

End Book Hotel (T₁)

Start Book Car Rental (T2)

End Book Car Rental (T₂)




Successful Saga

Begin Saga



Start Book Hotel (T1)

End Book Hotel (T₁)

Start Book Car Rental (T2)

End Book Car Rental (T₂)

Start Book Flight (T₃)





Successful Saga

Begin Saga



Start Book Hotel (T1)

End Book Hotel (T₁)

Start Book Car Rental (T2)

End Book Car Rental (T₂)

Start Book Flight (T₃)

End Book Flight (T₃)







Successful Saga

Begin Saga



Start Book Hotel (T1)

End Book Hotel (T₁)

Start Book Car Rental (T2)

End Book Car Rental (T₂)

Start Book Flight (T₃)

End Book Flight (T₃)

End Saga







Unsuccessful Saga Backwards Recovery





Start Book Hotel (T₁)

Begin Saga



Start Book Hotel (T₁)

End Book Hotel (T₁)



Begin Saga



Start Book Hotel (T₁)

End Book Hotel (T1)

Start Book Car Rental (T₂)

Begin Saga



End Book Hotel (T1)

Start Book Car Rental (T₂)

Abort Saga





Begin Saga

Start Book Hotel (T₁)

End Book Hotel (T1)

Start Book Car Rental (T2)

Abort Saga

Start Compensate Car Rental (C₂)





Begin Saga



Start Book Hotel (T₁)

End Book Hotel (T1)

Start Book Car Rental (T2)

Abort Saga

Start Compensate Car Rental (C2)

End Compensate Car Rental (C₂)

Begin Saga



Start Book Hotel (T₁)

End Book Hotel (T1)

Start Book Car Rental (T2)

Abort Saga

Start Compensate Car Rental (C2)

End Compensate Car Rental (C₂)

Start Compensate Book Hotel (C1)

Begin Saga



Start Book Hotel (T₁)

End Book Hotel (T1)

Start Book Car Rental (T2)

Abort Saga

Start Compensate Car Rental (C2)

End Compensate Car Rental (C₂)

Start Compensate Book Hotel (C1)

End Compensate Book Hotel (C1)

Begin Saga

Start Book Hotel (T₁)

End Book Hotel (T₁)

Start Book Car Rental (T2)

Abort Saga

Start Compensate Car Rental (C2)

End Compensate Car Rental (C₂)

Start Compensate Book Hotel (C1)

End Compensate Book Hotel (C1)





Sagas in Distributed Systems

"Due to space limitations, we only discuss Sagas in a centralized System, although clearly they can be implemented in a distributed database system."

-Molina et. al

SOA/Microservices



Requests instead of Transactions

Book Hotel (T1)



Cancel Hotel (C1)

- Book Car (T₂)
- Book Flight (T₃)

Cancel Car (C2)



Cancel Flight (C₃)

A Distributed Saga is a Collection of Sub-Requests

$T_1, T_2 \dots T_n$

Each Sub-Request has a Compensating Request

C1, C2...Cn

Successful Distributed Saga

Begin Saga

Start Book Hotel Request (T1)

End Book Hotel Request (T₁)

Start Book Car Rental Request (T₂)

End Book Car Rental Request (T₂)

Start Book Flight Request (T₃)

End Book Flight Request (T₃)

End Saga







Saga Log Durable & Distributed

Saga Execution Coordinator (SEC)

- Interprets & Writes to Saga Log
- Applies Saga Sub-Requests
- Applies Saga Compensating Requests when Necessary



















Apply Compensating Requests

- Aborted Saga Response
- Start Request Fails
- SEC Crashes (non-safe state)







What Happens when Compensating Requests Fail?

Compensating Requests Must Be Idempotent
What Happens when SEC Fails?

Safe States

- All Executed Sub-Requests are Complete (Start Ti & End Ti both logged)
- Saga has been Aborted, Proceed with Compensating Transactions

Un-Safe State

- Start Ti logged, no End Ti logged
 - Abort Saga
 - Start Compensating Requests

Request Messaging Semantics

- Sub-Requests (Ti): At Most Once
- Compensating Requests (Ci): At Least Once

Distributed Saga Guarantee

Either

T₁, T₂ ... T_n or
T₁, T₂ ... T_j, C_j, ... C₂, C₁

Distributed Sagas

- Distributed/Durable Saga Log
- SEC Process
- Compensating Requests: Idempotent



Halo Statistics

CORE TELLASIAN INTERNET DISCHARGERO HALO

SERVICE RECORD

Θ

0

۲

ų

Q



SRUKOP PR PT EN SKRETK

	WAR G	AMES MA	тонма			
	2294 945 094000	1443	130			
	WAR GAMES CUSTOM					
SR Q	35 evid central	18	03H			
- 14	SPARTAN OPS					
	0 / 50 000000 000	0/50	0011			
Hall Name	CAMPAIGN					
	0/8					

RECE

HOST VING

MOS

975

сти

ING

nн

IIM

	DECEMT O	-						-	
T CAMES	RECENT GAMES						ę	OLDER	
NAVED I STATS	0	HAVE	N 10.52495-1995	THC47	160		12		411
NUTATION	10	SHLT			215		14		21
ACTINA				250			17		6TB rus
					ar 100		14		511 1.043
	10	SHUT	OUT N SHIEF TO		125		10		3R
	Consectore and								
STATS SUMMARY									
OVERVIEW									
	130 06H 17H	1	11.22.12		92 2000 m			229 09440	
		100%	HORATONS	9 19000 1					
	WAR GAMES:	матенни	4KING						
	()	3D 0.3H (04M	2294 10% (MG) 1010		1443		7934	9
WAR GAMES: CUSTOM									
	•	00 03H 23M		35 Total Carriel A Level		18		753	
	SPARTAN OPS								
	6		ам			0 / 54		0/	50

CAMPAIGN

0D 00H 00M 0

MOST PLAYED VARIANT STATS

MATCHMAKING: SLAYER



0/8

SHUTOUT





TEAM RESULTS						OUTCOME	COMENT K/
1	۳	TOSSEM	сяк 26	270 xxxx	17 1811	34 Hora	CSR 26
2		TTL L ASKAN	594 27	215 xxxx	14 141.5	29 10143	CSR 26
3	٢	TIL YETI	сая 17	200	11 	29 	CSR 15
4	(66)	TTL KYLI3	618 26	145	8	21 HTML	CSR 24
1	\diamond	SAVACE SILVA	648 34	240	18 	29 	CSR 34
5		TRULYLEGENDARYX	сяя 21	155	<u>Ш</u>	80 ••••••	CSR 21
3	0	POT WOLTEY	сяк 23	130 xvrt	7	19 *0***	CSR 23
4	\otimes	ICOLIC BC	632 24	105 xxxx	7 1010	14	CSR 22

TEAM COMPARISON THE & ADDRESS CAR



MEDAL DISTRIBUTION THE REPORTS DON'T RESULT

CUSTOM: SLAYER

Halo Statistics Service



Halo Statistics Service with Sagas



Forward-Recovery



Forward Recovery Sub-Requests Must Also Be Idempotent







- Long Lived / Distributed Transactions
- Trade Atomicity for Availability
- Failure Management Pattern

Thank You

- @pbailis
 @aphyr
- @randommood
- @jmhodges
- @tsantero
 @clemnsv





Questions?

Please remember to evaluate via the GOTO Guide App





Conference: May 11-12 / Workshops: 13-14