

Torturing Databases for Fun and Profit

Mai Zheng, Joseph Tucek, Dachuan Huang, Feng Qin, Mark Lillibridge, Elizabeth S Yang, Bill W Zhao, Shashank Singh



































database





- ACID: atomicity, consistency, isolation, and durability
 - even under failures

SD (E

List of databases survived



File			Workload				
Database	System	W-1	W-2	W-3	W-4.1	W-4.2	W-4.3
TokyoCabinet	ext3	D	D	D	ACD	ACD	ACD
	XFS		D	D	ACD	D	ACD
MariaDB	ext3	D	D	D	D	D	D
	XFS	D	D	D	D	D	D
LightningDB	ext3						D
	XFS						
SQLite	ext3	D	D		D	D	D
	XFS			D	D	D	D
KVS-A	ext3			Hang			
	XFS						
SQL-A	ext3	D	D	D	D	D	D
	XFS	D	D	D	D	D	D
SQL-B	ext3	D	D	CD	CD	CD	CD
	XFS	CD	D	CD	CD	CD	CD
SQL-C	NTFS	D	D	D	D	D	D

Everything is broken under simulated power faults

SD (15

Power faults cannot happen nowadays, right?





1584

0111

2013:"... POWER OUTAGE stops Super Bowl for 34 minutes ..."

200

.....

.....

.....



Database Torture 101



Fault Model: Clean termination of I/O stream



2015 Storage Developer Conference. ©New Mexico State University. All Rights Reserved.

15

Fault Model: Clean termination of I/O stream



15

Why not introduce corruption/dropping/reordering?

- Unreasonable to require databases to handle arbitrary bad behavior introduced in the lower layers
- Simulated bad behavior w/o verification by real failures may be unrealistic
 - I/O path in kernel & device puts constraints on failure states





How do we test DBs on multiple OSes w/ high fidelity?



SD (15 Sto

How do we test DBs on multiple OSes w/ high fidelity?



S

15

How do we test DBs on multiple OSes w/ high fidelity?



S

15

Framework Overview





Framework Overview

















		key	value
ſ	_	THR-1-TXN-1	v-init-THR-1-TXN-1
meta		THR-1-TXN-2	v-init-THR-1-TXN-2
rows		THR-2-TXN-1	v-init-THR-2-TXN-1
		THR-2-TXN-2	v-init-THR-2-TXN-2
work rows		k-1	v-init-1
		k-2	v-init-2
		k-3	v-init-3
		k-4	v-init-4
		k-5	v-init-5
		k-6	v-init-6
		k-7	v-init-7
		k-8	v-init-8
		DB	table



	key	value	
Γ	THR-1-TXN-1	v-init-THR-1-TXN-1	
meta	THR-1-TXN-2	v-init-THR-1-TXN-2	
rows	THR-2-TXN-1	v-init-THR-2-TXN-1	
	THR-2-TXN-2	v-init-THR-2-TXN-2	
	k-1	v-init-1	
	k-2	v-init-2	
	k-3	v-init-3	
work	k-4	v-init-4	
rows	k-5	v-init-5	
	k-6	v-init-6	
	k-7	v-init-7	
	k-8	v-init-8	
	DB	table	-







	key	value	
Γ	THR-1-TXN-1	k-2-k-5-TS-00:01	save work-row keys
meta	THR-1-TXN-2	v-init-THR-1-TXN-2	& timestamp
rows	THR-2-TXN-1	v-init-THR-2-TXN-1	right before commit
	THR-2-TXN-2	v-init-THR-2-TXN-2	
Γ	k-1	v-init-1	
	k-2	v-THR-1-TXN-1	
	k-3	v-init-3	save
work	k-4	v-init-4	transaction ID
rows	k-5	v-THR-1-TXN-1	
	k-6	v-init-6	
	k-7	v-init-7	
	k-8	v-init-8	
	DE	s table	



















A power fault just happened during our workload ...



	key	value		
Γ	THR-1-TXN-1	k-2-k-5-TS-00:01		
meta	THR-1-TXN-2	k-6-k-8-TS-00:13		
rows	THR-2-TXN-1	k-7-k-6-TS-00:03		
Ĺ	THR-2-TXN-2	v-init-THR-2-TXN-2		
ſ	k-1	v-init-1		
	k-2	v-THR-1-TXN-1		
	k-3	v-THR-2-TXN-2		
work	k-4	v-init-4		
rows	k-5	v-THR-1-TXN-1		
	k-6	v-THR-1-TXN-2		
	k-7	v-THR-2-TXN-2		
	- k-8	v-THR-1-TXN-2		
recovered DB table				



		кеу	value
ſ	THE	R-1-TXN-1	k-2-k-5-TS-00:01
meta	THE	R-1-TXN-2	k-6-k-8-TS-00:13
rows	THE	R-2-TXN-1	k-7-k-6-TS-00:03
Į	THF	R-2-TXN-2	v-init-THR-2-TXN-2
[k-1		v-init-1
	k-2		v-THR-1-TXN-1
	k-3		v-THR-2-TXN-2
work rows	k-4		v-init-4
	k-5		v-THR-1-TXN-1
	<mark>k-6</mark>		v-THR-1-TXN-2
	k-7		v-THR-2-TXN-2
	k-8		v-THR-1-TXN-2
recovered DB table			














Is there any ACID violation after recovery?

		key	١	/alue								
ſ	_	THR-1-TXN-1	k-2-k-5-	TS-00:01								
meta		THR-1-TXN-2	k-6-k-8-	TS-00:13								
rows		THR-2-TXN-1	k-7-k-6-	TS-00:03								
	_	THR-2-TXN-2	v-init-Th	HR-2-TXN-	.2							
		k-1	v-init-1									
		k-2	v-THR-	1-TXN-1								
		k-3	v-THR-2	2-TXN-2								
work		k-4	v-init-4									
rows		k-5	v-THR-	1-TXN-1								
		k-6	v-THR-	1-TXN-2								
		k-7	v-THR-2	2-TXN-2								
		k-8	v-THR-	1-TXN-2								
	recovered DB table											

allow checking time & order related properties

SD (E

Framework Overview





Framework Overview





Capturing I/O trace without kernel modification





Capturing I/O trace without kernel modification





Constructing a post-fault disk image



2015 Storage Developer Conference. ©New Mexico State University. All Rights Reserved.

SD

Constructing a post-fault disk image





SD

Checking the post-fault DB



2015 Storage Developer Conference. ©New Mexico State University. All Rights Reserved.

SD

Checking the post-fault DB



2015 Storage Developer Conference. ©New Mexico State University. All Rights Reserved.

SD

Testing different fault points easily



2015 Storage Developer Conference. ©New Mexico State University. All Rights Reserved.

Testing different fault points easily



2015 Storage Developer Conference. ©New Mexico State University. All Rights Reserved.

The framework is not good enough

- Sometimes need several days
 - too many mini operations, too many potential fault points





The framework is not good enough

- Sometimes need several days
 - too many mini operations, too many potential fault points
- We tried sampling
 - but only a few fault points trigger ACID violations





The framework is not good enough

- Sometimes need several days
 - too many mini operations, too many potential fault points
- We tried sampling

 but only a few fault points trigger ACID violations
- Don't know why



Enhanced Design

Framework Overview





Framework Overview





Original trace provides little semantics





Enhancing w/ more context



What makes some fault points special?



checking result



15

• REP_p: repetitive LBA (Logical Block Address)

checking result Worker's multi-layer trace

00

11

...

00

•

•



fixed location for important metadata & in-place update



MMAP_p: unintended update to mmap'ed blocks

		op#	LBA	file	syscall
	•	•••	•••	•••	
<u></u>		463	1012	x.db	fsync(x.log)
~~	•		•••	•••	
-	•		•••		
	•		•••		
		564	1012	x.db	msync(x.db)
	•	•••	•••		



MMAP_p: unintended update to mmap'ed blocks

		op#	LBA	file	syscall
	•		•••	•••	
		463	1012	x.db	fsync(x.log)
	•	•••	•••		
-	:				
9	•				
		564	1012	x.db	msync(x.db)
	:				



• MMAP_p: unintended update to mmap'ed blocks





• MMAP_p: unintended update to mmap'ed blocks





• MMAP_p: unintended update to mmap'ed blocks



implicit flush of dirty blocks by kernel or FS under heavy transactions



• MMAP_p: unintended update to mmap'ed blocks



• MMAP_p: unintended update to mmap'ed blocks



Three more patterns: JUMP_p, HEAD_p, TRAN_p

Add fault injection policy to determine where to inject faults



2015 Storage Developer Conference. ©New Mexico State University. All Rights Reserved.

(15)

Worker's multi-layer trace

Scoreboard

total

											totai
op#	LBA	cmd#	file	syscall		$MMAP_{p}$	REP_{p}	$JUMP_{p}$	$HEAD_{p}$	$TRAN_{p}$	score
1	348	1	x.db	msyc(x.db)	1						
2	352	2	x.log	fsync(x.log)	2						
3	356	2	x.log	fsync(x.log)	3						
4	360	2	x.log	fsync(x.log)	4						
5	364	2	x.log	fsync(x.log)	5						
6	370	3	x.log	fsync(x.log)	6						
7	348	4	x.db	fsync(x.log)	7						
8	906	5	fs-j	fsync(x.log)	8						

SD (

Worker's multi-layer trace

Scoreboard

total

op#	LBA	cmd#	file	syscall		$MMAP_{p}$	REP_{p}	$JUMP_{p}$	$HEAD_{p}$	$TRAN_{p}$	score
1	348	1	x.db	msyc(x.db)	1	0					
2	352	2	x.log	fsync(x.log)	2	0					
3	356	2	x.log	fsync(x.log)	3	0					
4	360	2	x.log	fsync(x.log)	4	0					
5	364	2	x.log	fsync(x.log)	5	0					
6	370	3	x.log	fsync(x.log)	6	0					
7	348	4	x.db	fsync(x.log)	7	1					
8	906	5	fs-j	fsync(x.log)	8	1					

Worker's multi-layer trace

SD (15

Scoreboard

total

op#	LBA	cmd#	file	syscall		$MMAP_{p}$	REP_{p}	$JUMP_{p}$	$HEAD_{p}$	$TRAN_{p}$	score
1	348	1	x.db	msyc(x.db)	1	0	1				
2	352	2	x.log	fsync(x.log)	2	0	0				
3	356	2	x.log	fsync(x.log)	3	0	0				
4	360	2	x.log	fsync(x.log)	4	0	0				
5	364	2	x.log	fsync(x.log)	5	0	0				
6	370	3	x.log	fsync(x.log)	6	0	0				
7	348	4	x.db	fsync(x.log)	7	1	1				
8	906	5	fs-j	fsync(x.log)	8	1	0				

Worker's multi-layer trace

Scoreboard

total

op#	LBA	cmd#	file	syscall		$MMAP_{p}$	REP_{p}	$JUMP_{p}$	$HEAD_{p}$	$TRAN_{p}$	score
1	348	1	x.db	msyc(x.db)	1	0	1	0			
2	352	2	x.log	fsync(x.log)	2	0	0	0			
3	356	2	x.log	fsync(x.log)	3	0	0	0			
4	360	2	x.log	fsync(x.log)	4	0	0	0			
5	364	2	x.log	fsync(x.log)	5	0	0	0			
6	370	3	x.log	fsync(x.log)	6	0	0	1			
7	348	4	x.db	fsync(x.log)	7	1	1	1			
8	906	5	fs-j	fsync(x.log)	8	1	0	1			



Worker's multi-layer trace

Scoreboard

total

Worker's multi-layer trace

Scoreboard

				.,							
op#	LBA	cmd#	file	syscall		MMAP _p	REP_{p}	JUMP _p	$HEAD_{p}$	$TRAN_{p}$	score
1	348	1	x.db	msyc(x.db)	1	0	1	0	0	1	
2	352	2	x.log	fsync(x.log)	2	0	0	0	1	1	
3	356	2	x.log	fsync(x.log)	3	0	0	0	0	0	
4	360	2	x.log	fsync(x.log)	4	0	0	0	0	0	
5	364	2	x.log	fsync(x.log)	5	0	0	0	0	0	
6	370	3	x.log	fsync(x.log)	6	0	0	1	0	1	
7	348	4	x.db	fsync(x.log)	7	1	1	1	0	1	
8	906	5	fs-j	fsync(x.log)	8	1	0	1	0	1	



Worker's multi-layer trace

Scoreboard

total

op#	LBA	cmd#	file	syscall		MMAP _p	REP_{p}	JUMP _p	$HEAD_{p}$	TRAN _p	score
1	348	1	x.db	msyc(x.db)	1	0	1	0	0	1	2
2	352	2	x.log	fsync(x.log)	2	0	0	0	1	1	2
3	356	2	x.log	fsync(x.log)	3	0	0	0	0	0	0
4	360	2	x.log	fsync(x.log)	4	0	0	0	0	0	0
5	364	2	x.log	fsync(x.log)	5	0	0	0	0	0	0
6	370	3	x.log	fsync(x.log)	6	0	0	1	0	1	2
7	348	4	x.db	fsync(x.log)	7	1	1	1	0	1	4
8	906	5	fs-j	fsync(x.log)	8	1	0	1	0	1	3
Alternative to Exhaustive: Pattern-based Ranking

Worker's multi-layer trace

S

15

Scoreboard

total

op#	LBA	cmd#	file	syscall		MMAP _p	REP_{p}	JUMP _p	$HEAD_{p}$	$TRAN_{p}$	score
1	348	1	x.db	msyc(x.db)	1	0	1	0	0	1	2
2	352	2	x.log	fsync(x.log)	2	0	0	0	1	1	2
3	356	2	x.log	fsync(x.log)	3	0	0	0	0	0	0
4	360	2	x.log	fsync(x.log)	4	0	0	0	0	0	0
5	364	2	x.log	fsync(x.log)	5	0	0	0	0	0	0
6	370	3	x.log	fsync(x.log)	6	0	0	1	0	1	2
7	348	4	x.db	fsync(x.log)	7	1	1	1	0	1	4
8	906	5	fs-j	fsync(x.log)	8	1	0	1	0	1	3
	1 st -ra 2 nd -ra 3 rd -ra 4 th -ra	nk: ank: nk: ank:	7 ← 8 1 2 3 4	2 <u>6</u> 4 <u>5</u>		predict	ed m	ost er	ror-pr	one	

Alternative to Exhaustive: Pattern-based Ranking



2015 Storage Developer Conference. ©New Mexico State University. All Rights Reserved.

S

Alternative to Exhaustive: Pattern-based Ranking





Diagnosis Support

Worker's multi-layer trace helps understand what happened at fault time





15

Add function-call tracing to disclose more semantics for diagnosis



2015 Storage Developer Conference. ©New Mexico State University. All Rights Reserved.

15

Enable same tracing during checking to see why recovery didn't work



Results

Experimental Environment

- 8 databases
 - Open-source: TokyoCabinet, MariaDB, LightningDB, SQLite
 - Commercial: KVS-A, SQL-A, SQL-B, SQL-C
- 4 workloads
- 3 file systems
 - ext3, XFS, NTFS
- Several operating systems
 - Linux: RHEL 6, Debian6, Ubuntu 12 LTS
 - Windows 7 Enterprise

Not a single DB can survive all tests

DB	FS	W-1	W-2	W-3	W-4.1	W-4.2	W-4.3	Α	С	I	D
TakyoCabinat	ext3	D	D	D	ACD	ACD	ACD	0.15%	0.14%	0	16.05%
токуосартнет	XFS		D	D	ACD	D	ACD	<0.01%	0.01%	0	4.38%
MariaDD	ext3	D	D	D	D	D	D	0	0	0	1.36%
IVIdIIdDB	XFS	D	D	D	D	D	D	0	0	0	0.49%
LightningDD	ext3						D	0	0	0	0.05%
	XFS							0	0	0	0
SOL ita	ext3	D	D		D	D	D	0	0	0	19.15%
SQLILE	XFS			D	D	D	D	0	0	0	10.60%
	ext3			Hang				0	0	0	0
KV3-A	XFS							0	0	0	0
	ext3	D	D	D	D	D	D	0	0	0	3.31%
SQL-A	XFS	D	D	D	D	D	D	0	0	0	0.92%
	ext3	D	D	CD	CD	CD	CD	0	8.96%	0	3.24%
JUL-D	XFS	CD	D	CD	CD	CD	CD	0	7.77%	0	3.90%
SQL-C	NTFS	D	D	D	D	D	D	0	0	0	8.08%



Durability violation is most common

DB	FS	W-1	W-2	W-3	W-4.1	W-4.2	W-4.3	Α	С	I	D
TakyoCabinat	ext3	D	D	D	ACD	ACD	ACD	0.15%	0.14%	0	16.05%
lokyocabinet	XFS		D	D	ACD	D	ACD	<0.01%	0.01%	0	4.38%
MariaDB	ext3	D	D	D	D	D	D	0	0	0	1.36%
Manada	XFS	D	D	D	D	D	D	0	0	0	0.49%
LightningDP	ext3						D	0	0	0	0.05%
	XFS							0	0	0	0
SOLita	ext3	D	D		D	D	D	0	0	0	19.15%
SQLITE	XFS			D	D	D	D	0	0	0	10.60%
	ext3			Hang				0	0	0	0
KV3-A	XFS							0	0	0	0
	ext3	D	D	D	D	D	D	0	0	0	3.31%
JQL-A	XFS	D	D	D	D	D	D	0	0	0	0.92%
	ext3	D	D	CD	CD	CD	CD	0	8.96%	0	3.24%
JQL-D	XFS	CD	D	CD	CD	CD	CD	0	7.77%	0	3.90%
SQL-C	NTFS	D	D	D	D	D	D	0	0	0	8.08%

SD (6

Some violations are difficult to trigger

DB	FS	W-1	W-2	W-3	W-4.1	W-4.2	W-4.3	Α	С	I	D
TakyoCabinat	ext3	D	D	D	ACD	ACD	ACD	0.15%	0.14%	0	16.05%
lokyocabillet	XFS		D	D	ACD	D	ACD	<0.01%	0.01%	0	4.38%
MariaDD	ext3	D	D	D	D	D	D	0	0	0	1.36%
IVIdIIdDD	XFS	D	D	D	D	D	D	0	0	0	0.49%
LightningDP	ext3						D	0	0	0	0.05%
	XFS							0	0	0	0
SOLita	ext3	D	D		D	D	D	0	0	0	19.15%
SQLITE	XFS			D	D	D	D	0	0	0	10.60%
	ext3			Hang				0	0	0	0
KVJ-A	XFS							0	0	0	0
	ext3	D	D	D	D	D	D	0	0	0	3.31%
JUL-A	XFS	D	D	D	D	D	D	0	0	0	0.92%
	ext3	D	D	CD	CD	CD	CD	0	8.96%	0	3.24%
JUL-D	XFS	CD	D	CD	CD	CD	CD	0	7.77%	0	3.90%
SQL-C	NTFS	D	D	D	D	D	D	0	0	0	8.08%



Some violations are difficult to trigger

DB	FS	W-1	W-2	W-3	W-4.1	W-4.2	W-4.3	Α	С	I	D
TalwaCabinat	ext3	D	D	D	ACD	ACD	ACD	0.15%	0.14%	0	16.05%
lokyocabillet	XFS		D	D	ACD	D	ACD	<0.01%	0.01%	0	4.38%
MariaDD	ext3	D	D	D	D	D	D	0	0	0	1.36%
MANADB	XFS	D	D	D	D	D	D	0	0	0	0.49%
LightningDD	ext3						D	0	0	0	0.05%
	XFS							0	0	0	0
	ext3	D	D		D	D	D	0	0	0	19.15%
SQLITE	XFS			D	D	D	D	0	0	0	10.60%
	ext3			Hang				0	0	0	0
KV3-A	XFS							0	0	0	0
	ext3	D	D	D	D	D	D	0	0	0	3.31%
SQL-A	XFS	D	D	D	D	D	D	0	0	0	0.92%
	ext3	D	D	CD	CD	CD	CD	0	8.96%	0	3.24%
JQL-B	XFS	CD	D	CD	CD	CD	CD	0	7.77%	0	3.90%
SQL-C	NTFS	D	D	D	D	D	D	0	0	0	8.08%



• Failure symptoms

faults injected in a region of operations cause:

- A violation: a transaction is partially committed
- D violation: some rows are irretrievable
- C violation: retrievable rows by range query and point queries are different



Why recovery didn't work?

tchdbopenimpl(x.tcb)	tchdbopenimpl(x.tcb)
open(x.tcb) = 3	open(x.tcb) = 3
read(x.tcb) = 256	read(x.tcb) = 256
//op#, LBA, content, file	//op#, LBA, content, file
op#1, 480,100, x.tcb	op#1, 480,101, x.tcb
	tchdbwalrestore()
tcbdbget()	tcbdbget()
Checker's trace	Checker's trace

hen ACID violations were found Checker's trace when no violation was found

Delta Debugging [Zeller, SIGSOFT'02/FSE-10]



Why recovery didn't work?

 tchdbopenimpl(x.tcb)	 tchdbopenimpl(x.tcb)
open(x.tcb) = 3	open(x.tcb) = 3
read(x.tcb) = 256	read(x.tcb) = 256
//op#, LBA, content, file	//op#, LBA, content, file
op#1, 480,100, x.tcb	op#1, 480,101, x.tcb
//no tchdbwalrestore()	tchdbwalrestore()
tcbdbget()	tcbdbget()

Checker's trace when ACID violations were found

Checker's trace when no violation was found

Delta Debugging [Zeller, SIGSOFT'02/FSE-10]



Why recovery didn't work? ... tchdbopenimpl(x.tcb) ... open(x.tcb) = 3 read(x.tcb) = 256 //op#, LBA, content, file op#1, 480, ...100, ., x.tcb //no tchdbwalrestore() tcbdbget() ...

Checker's trace when ACID violations were found Checker's trace when no violation was found

Delta Debugging [Zeller, SIGSOFT'02/FSE-10]



What happened at fault time? Why recovery didn't work? Worker's trace tchdbopenimpl(x.tcb) tchdbopenimpl(x.tcb) around the bug-triggering . . . fault points #30 - #90 open(x.tcb) = 3open(x.tcb) = 3(x.tcb) = 256read(x.tcb) = 256#, LBA, content, file //op#, LBA, content, file mmap(8192, x.tcb, 0) op#1, 480, ...**101**..., x.tcb L, 480, ...**100**..., x.tcb tchdbwalrestore() tchdbwalrestore() fsync(x.tcb.wal) bget() tcbdbget() //op#, LBA, content, file, syscall op#26, 630,, x.tcb.wal, fsync(x.tcb.wal) op#27, 960,, fs-j , fsync(x.tcb.wal) op#28, 964,, fs-j , fsync(x.tcb.wal) op#29, 480, ...100..., x.tcb , fsync(x.tcb.wal) msync(x.tcb) //op#, LBA, content, file, syscall op#91, 480, ...101..., x.tcb, msync(x.tcb)















Patterns reduce required test points greatly while achieving similar coverage



2015 Storage Developer Conference. ©New Mexico State University. All Rights Reserved.

15

Patterns reduce required test points greatly while achieving similar coverage





Patterns reduce required test points greatly while achieving similar coverage



- A wake-up call
 - traditional testing methodology may not be enough for today's complex storage systems
 - thorough testing requires purpose-built workloads and intelligent fault injection techniques



- A wake-up call
 - traditional testing methodology may not be enough for today's complex storage systems
 - thorough testing requires purpose-built workloads and intelligent fault injection techniques
- Different layers in OS can help in different ways
 - iSCSI: fault injection w/ high portability & high fidelity
 - LBA & syscall: generic behavior patterns
 - combined multi-layer info: clear whole picture of complicated scenarios



• We should bridge the gaps of understanding/assumptions!

between DB & OS

between User & DB



• We should bridge the gaps of understanding/assumptions!

Pop Quiz: true or false?

between	"mmap'ed files are not updated until msync()"	
DB & OS	"file-length update are persistent after fdatasync()"	
between	"durability is provided by the default configuration"	
User & DB	"transactions are durable after COMMIT"	



• We should bridge the gaps of understanding/assumptions!

Pop Quiz: true or false?

between	"mmap'ed files are not updated until msync()"	false
DB & OS	"file-length update are persistent after fdatasync()"	
between	"durability is provided by the default configuration"	
User & DB	"transactions are durable after COMMIT"	



• We should bridge the gaps of understanding/assumptions!

Pop Quiz: true or false?

between	"mmap'ed files are not updated until msync()"	false
DB & OS	"file-length update are persistent after fdatasync()"	depends!
between	"durability is provided by the default configuration"	depends!
User & DB	"transactions are durable after COMMIT"	depends!



• We should bridge the gaps of understanding/assumptions!

Pop Quiz: true or false?

between	"mmap'ed files are not updated until msync()"	false
DB & OS	"file-length update are persistent after fdatasync()"	depends!
between	"durability is provided by the default configuration"	depends!
User & DB	"transactions are durable after COMMIT"	depends!

Thank you!



