MySQL Backups

presented at the

Boston MySQL Users Group

January 9th, 2006

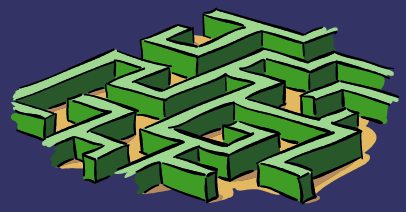

Introduction

- Sheeri Kritzer, MySQL DBA
  - awfief@gmail.com
  - http://www.sheeri.com

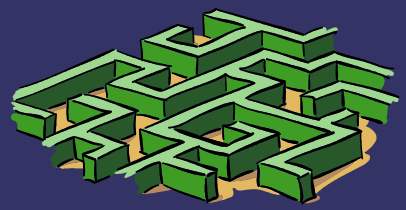

First Things First. . .

- Not an expert
- Not any smarter than you
- Am not affiliated with MySQL AB

- Do not have the breadth of experience of the group

- Feel free to ask questions
- Correct me

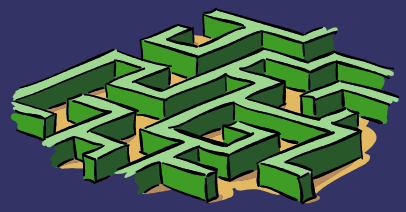

Qualifications

- Master's degree in Computer Science Specializing in Databases (2001)
- MySQL DBA
- MySQL 4.1 Core Certified
- Over 5 years experience with MySQL

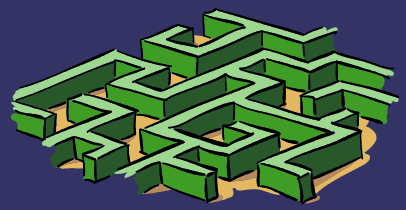

Most Importantly . . .

I volunteered!

Anyone can present a topic at these meetings.
Just contact me, awfief@gmail.com

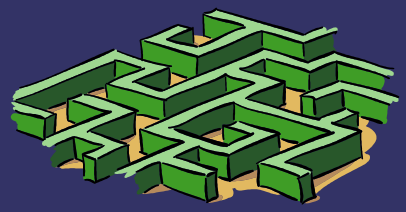

Why?

- Database migration
- Archive
- Reporting
- Analysis/Liability
- Recovery

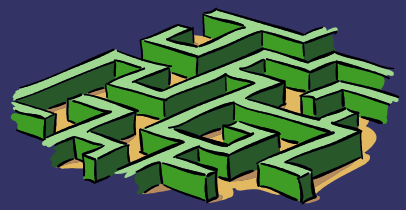

Recovery

- Recovering from data corruption
- Recovering from data loss
  - Not losing ANY data? (bank)
  - 10 minutes' worth?
  - An hour's worth?

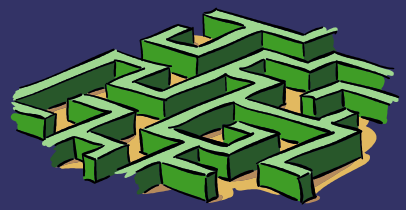

Who?

- Ultimately, head of IT

- Backup service?
- Sysadmin?
- DBA?
- Operators?

- **Make duties clear**

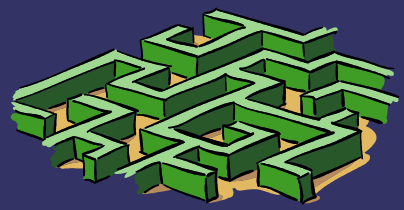

What?

- Do you need to back all your data up?
- Commands? (SOX, HIPAA. . .)
- TOS
- Expectations

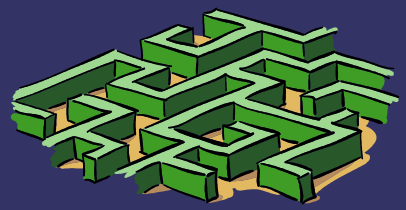

When?

- Static data backed up less frequently
- Not during high usage/reporting

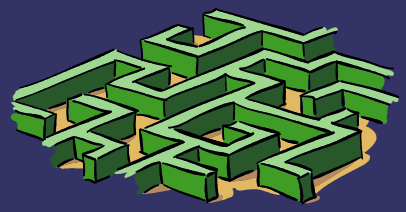

Where?

- Where will the backups be stored?
  - On the db server (different partition!)
  - Copied to another server (on/off site)
  - Backed up to tape (on/off site)

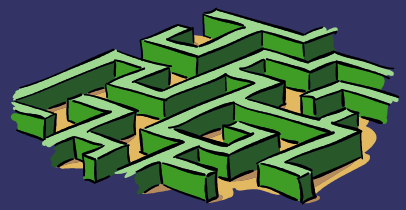

Disaster Zones

- Given earthquakes, mudslides, monsoons, wildfires, flooding, hurricanes, blizzards, tornadoes, tsunami, sandstorms, eruptions, avalanches, lahars, terrorism, extreme heat, and extreme cold, it seems there is no GOOD place to keep information.
- It's never going to be perfect!

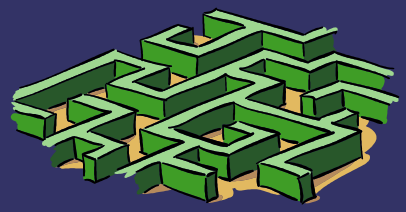

Lahar

How?

- SELECT . . . INTO OUTFILE
- mysqldump
- BACKUP TABLE
- Replication

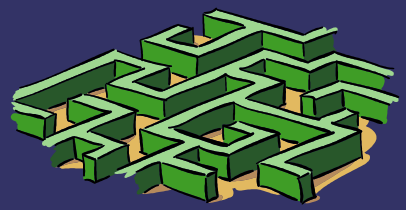

How? (con't)

- OS level copy of data files
- mysqlhotcopy
- InnoDB hot backup

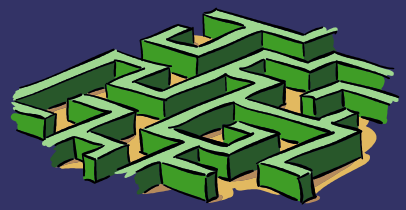

Are Backups Used for Recovery?

- If not a snapshot, recovery harder
- Binary logging for DML after snapshots
- Exception: Replication

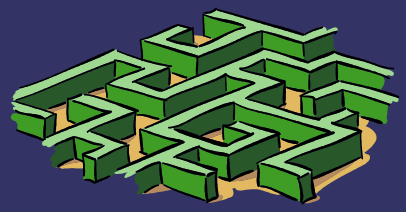

Why So Many Options?

- Organic growth
- Different storage architectures
- OS level tools
- MySQL client commands

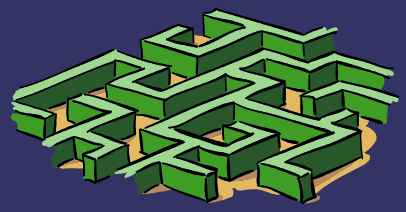

SELECT . . . INTO OUTFILE

- MySQL extension to ANSI standard SQL
- Works with all storage engines
- Very flexible

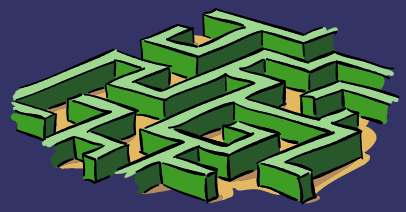

SELECT . . . INTO OUTFILE (con't)

- Constrained by SELECT
  - Does not get schema/DDL in versions < 5.0
  - Extra code to get all tables

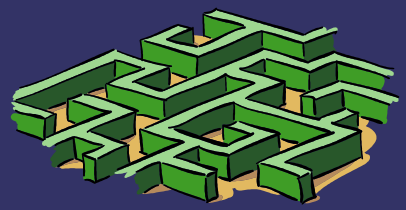

mysqldump

- Operating System binary
- In wide usage
- For all storage engines
- Easy to spot corruption in text files

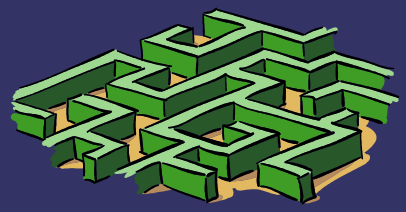

mysqldump's Many Options

- Can backup one or all tables/databases
- Thorough documentation online
- --quick
- --compatible

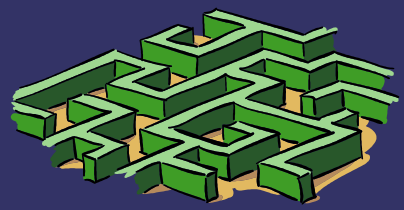

mysqldump's Many Options (con't)

- --lock-tables, --single-transaction
- --master-data=1
- --opt has most of what you need
- --skip-opt or skip-*feature*

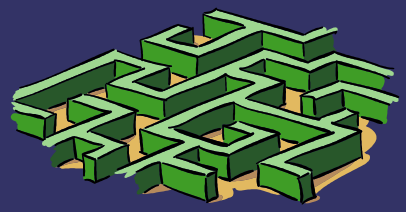

Replication

- Replication creates an automatic sync
- Can use for backup, HA, reporting
- Slave can use different engines than master
- 1 master can replicate to many slaves

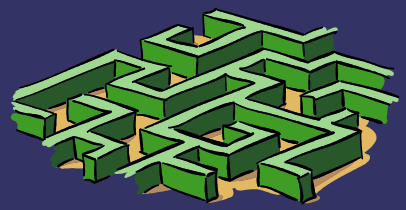

Replication Considerations

- No need for locking
- A true snapshot
- Does not recover from corrupt data
- Requires another MySQL instance

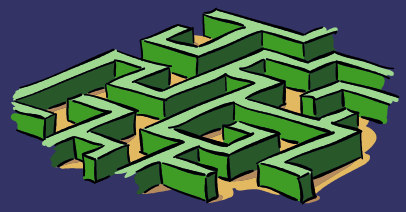

BACKUP TABLE

- Deprecated – use mysqlhotcopy

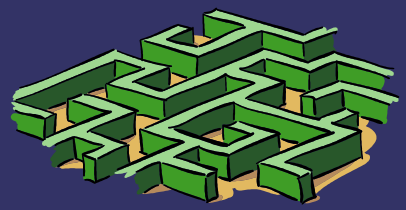

OS Level Copying of Files

- Lock tables separately
- External code
- Backup corruption harder to catch
- Can be used for migration
  - OS independence with MyISAM, InnoDB

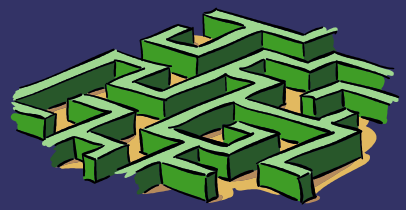

mysqlhotcopy

- Only for MyISAM
- Perl script for Unix and NetWare
- Regexp (perl) for db/table selection
- --dryrun

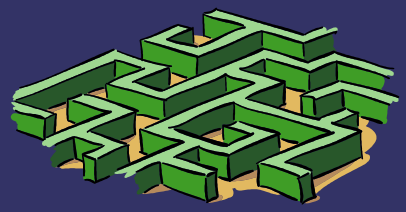

InnoDB Hot Backup

- Commercial -- $500 per year per server
  - $1300 per server for lifetime
- No locks
- Copies data files
- --apply-log to make it current
  - Or to any point in time

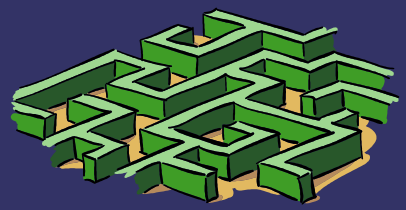

InnoDB Hot Backup

- Can compress with control
- Extra (free) tool to backup MyISAM
  - Consistent with InnoDB
  - Uses LOCK TABLES and OS level copy

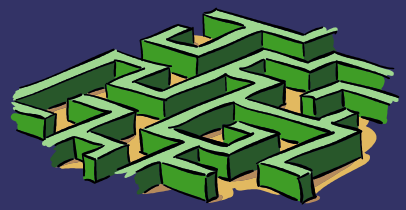

Comparison

Method	No Locking	DDL	Snapshot	Remote
SELECT . . . INTO OUTFILE	Engine Dependent	No	No	Yes
mysqldump	No	Option	No	Yes
Replication	Yes	No	Yes	Yes
OS level copy	No	Yes	No	No
mysqlhotcopy	Yes	Yes	Yes	No
InnoDB Hot Backup	Yes	Yes	Yes	No

Method	Free	All Engines	All Tables	Text File	Recover Corruption
SELECT . . . INTO OUTFILE	Yes	Yes	No	Yes	Yes
mysqldump	Yes	Yes	Yes	Yes	Yes
Replication	Yes	Yes	Yes	No	No
OS level copy	Yes	No	Yes	No	Yes
mysqlhotcopy	Yes	No	Yes	No	Yes
InnoDB Hot Backup	No	No	Yes	No	Yes

Best Practices

- Use more than one
- Check during a non-emergency

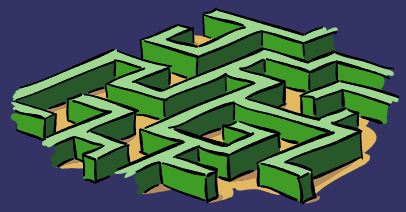

My Company

- Replication for immediate failover
- InnoDB Hot Backup (plus MyISAM backup)
- mysqldump to protect against corruption
- Copied to other servers
- Uses backups + binary logs to make new slaves without stopping the master

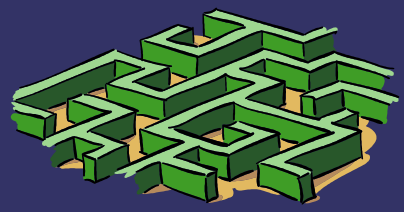

How about you?

- Current practices
- Desires
- Past mistakes
- Comments/suggestions for MySQL AB