

Best Practices for Database Administrators

Sheeri K. Cabral

Database Administrator

The Pythian Group, www.pythian.com

cabral@pythian.com

2008 MySQL User Conference & Expo



MIRE

- Make It Really Easy
- Automate
- Document
- As for your brain.....

Use your brain for CPU, not storage!

(use a request tracking system!)

Monitoring Basics

- Graph (Cacit)
- Alert (Nagios)
- Check your checks

Tradeoffs

- Tradeoffs always exist
- Think about them

Most Commonly Given Advice

- Backup
- Restore
- What do you use backups for?

In Case of Failure.....

- Restore from backup
- Master/Slave
- Master/Master
- Cluster



DR? HA!

- Test your DR/HA plans
- What scenarios do they cover?
- What scenarios do they NOT cover?

Ounces of Prevention

- Configuration vs. reality
- Possible memory usage
- Disk space / tablespace size

Pounds of Cure

- Error Logs
- Slow Query Logs
- Query Review



Query Profile

- mk-query-profiler
- mysqlreport
- SHOW STATUS before and after
- Is every index being used?

What Does This Query Do?

- `SELECT amount, created FROM payments where user='sheeri'`



What Does This Query Do?

- `SELECT amount, created FROM payments where user='sheeri'`

VS

- `SELECT /* find payments for customer */ amount, created FROM payments where user='sheeri'`



Data Profile

- PROCEDURE ANALYSE()
- What does each column, table do?
- What's in a name?

Does It Make Sense?

- Can you “read” the data?
- Question “best practices”
- Then, make some!



Schema Profile

- Start normal
- Denormalize if necessary
 - Descriptive foreign keys can prevent denormalization
- Stored procedures for developers

Replication

- Be careful of TRIGGERS
 - And any DML from SP or UDF
- Sync often with mk-table-sync
- Handle duplicates carefully

Maintenance

- Partition
- Archive
- Purge
- Static data

Manage User Expectations

- Constant report refreshes
- Aggregate data once every hour or 15 minutes
- Split off processing:
 - customer/non-customer
 - internal/external



Creative Suggestions – Forums

- COUNT(*) for paging
- LIMIT 0,n+1
- Cache each page of forum
 - First page can have up to n+5 entries

Creative Suggestions – Calculations

- Limited set of calculations?
 - Distance between zip codes in the US
- Calculate the same thing more than once?
- Calculating easy constants?

Foreign Keys

- Application still has to handle problems



Summary



When I Walk In the Door

- alerting
- request tracking
- graphing
- documentation



My Toolbox

- EXPLAIN
- Ma'atkit
- hackmysql.com (mysqlsla, mysqlreport)
- MySQL Manual

Also Useful

- mytop
- innotop

Testing, Testing...is this thing on?

- Functional
- Load
- Start now!



Consider

- master/master vs. read_only slave
- Learn the foundations
- Semi-dynamic data


Redundancy

- Make everything replaceable
 - Kickstart
 - Automated config management
- Including yourself!
 - Don't you want a vacation?

Be a Good DBA






- Prove your work to yourself
- Be clear





FoxTrot

by Bill Amend



Sheeri Cabral

cabral@pythian.com

www.sheeri.com