

“Change Bad!”

Coping with change in the data you store

Andrew J. Todd esq.

OSDC – November, 2009

What?

A journey through identifying, understanding
and managing change in database schemas

Who?

Application developers

Who?

Database designers/developers

Who?

Database Administrators

Databases



INGRES



ORACLE



Even these ones

Amazon SimpleDB™ BETA



mongoDB



CouchDB
relax

Terminology

- Database
 - A service that allows an application to store data
- Database Schema
 - A logical collection of objects within a database
- Database Object
 - A construct within a database that stores similar items
- Domain Model
 - A representation of our understanding of a problem area

Change

“Data matures like wine, applications like fish”

The Causes of Change

User requirements maturing

The Causes of Change

New domains

The Causes of Change

YAGNI

When Do Things Change?

Data models have a lower rate of change than the applications that work with them

What Changes?

Tables

What Changes?

Columns

What Changes?

Views

What Changes?

Stored Procedures

What Changes?

Indexes

How Do We Track Changes?

Put the DDL in version control

How Do We Track Changes?

Put a database independent model in version control

How Do We Track Changes?

But ...

Database = DDL + Data

How Do We Track Changes?

Actually ...

Database = DDL + Data + Implementation

Strategies For Coping

Do Nothing

Strategies For Coping

Store complete DDL for each schema version

Strategies For Coping

Store complete initial DDL and subsequent deltas

Strategies For Coping

Store complete DDL for every version
and subsequent deltas

Strategies For Coping

Don't forget reference data

Strategies For Coping

Or schema (user) creation scripts

What To Do When Things Go Wrong

You have a backup, right?

What To Do When Things Go Wrong

The secret to success is testing

What To Do When Things Go Wrong

Where possible complement 'upgrade' with
'downgrade'

What To Do When Things Go Wrong

Each migration should maintain data integrity

What To Do When Things Go Wrong

From “Database Refactoring”:

- Prior to the migration take a copy of the table being updated and store it in a staging area
- Validate that the copy and the original object have the same structure and contents.
- Perform the migration on the original object
- Validate that the migration is successful - by running the migration tests
- After a suitable period of time, usually not less than a week, remove the copy of the table from the staging area

What To Do When Things Go Wrong

High Availability is a different kettle of fish

Finis

Version Control is your friend but you need to think about how you use it ...