



Geodatabase Replication in the Real World

Implementing an Enterprise-level Geodatabase 2-Way Replication model
across poor ICT Infrastructure

Dr. Mark Norris-Rogers

July 2012

Introduction



- Presentation Overview
 - Mondi at a glance
 - Mondi's GIS – Old and New
 - Mondi's Replication Configuration
 - Mondi's Database Design Issues and System Limitations
 - Mondi's Critical Success Factors
 - Careful Database Design
 - Rigid Workflow Procedures
 - Thorough User Training

Mondi at a Glance



- International vertically integrated forestry, pulp and paper company

- **Mondi South Africa:**

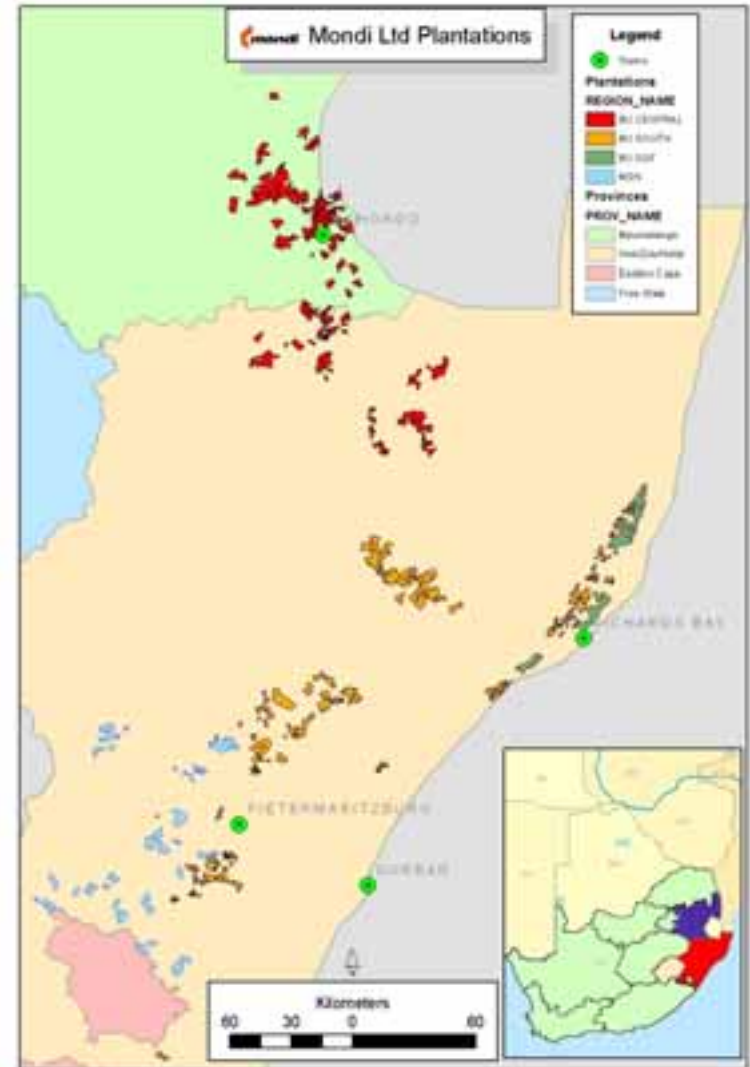
- 320 000 ha: 6 areas in KZN and Mpumalanga

- **Forestry plantations situated in rural areas**

- Poor ICT infrastructure
 - Low bandwidth lines
 - Copper rather than Fibre optic backbone
 - Line theft
 - Very high costs

- **Mondi's GIS History**

- 1994-2001: Arc/Info (Unix); ArcSTORM; ArcView
 - Distributed databases – little integration
 - One question = Different Answers.
 - 2001-2011: Customised MapObjects App
 - Written for Mondi
 - Centralised SDE Database
 - Edits on thick client via Citrix – Very slow
 - Line drops = data corruption
 - 2011: ArcGIS 10.0 Replicated GDB



ArcGIS 10.0 Replicated Geodatabase



● Requirements for New GIS

○ “Of-the-Shelf” Software

- Minimal customisation required
- “Plug and play” implementation
- Industry standard – adequate support

○ **Disconnected Editing, with Archiving Capability**

- Work locally - Removes network constraints
- Maintains advantages of centralised database
- Allows quality-control processes on edits

○ **Topology Validation Capability**

- Essential to maintain spatial data integrity

○ **Each of above requirements due to problems experienced with previous systems**

- System has achieved the desired results!

Logic for using Replication



○ 2-Way Replication System

○ Recommended by Esri Geodatabase Experts in Redlands

- Original replication design required two separate 1-way replications:
 - Default to Local
 - Local to QA Version hanging off Default (Central Server)
 - Reconcile/Post operation from QA to Default
 - Delete QA Version
 - Compress Default to Base Tables
- Esri experts thought there could be unintended consequences using this method (not design to work this way)

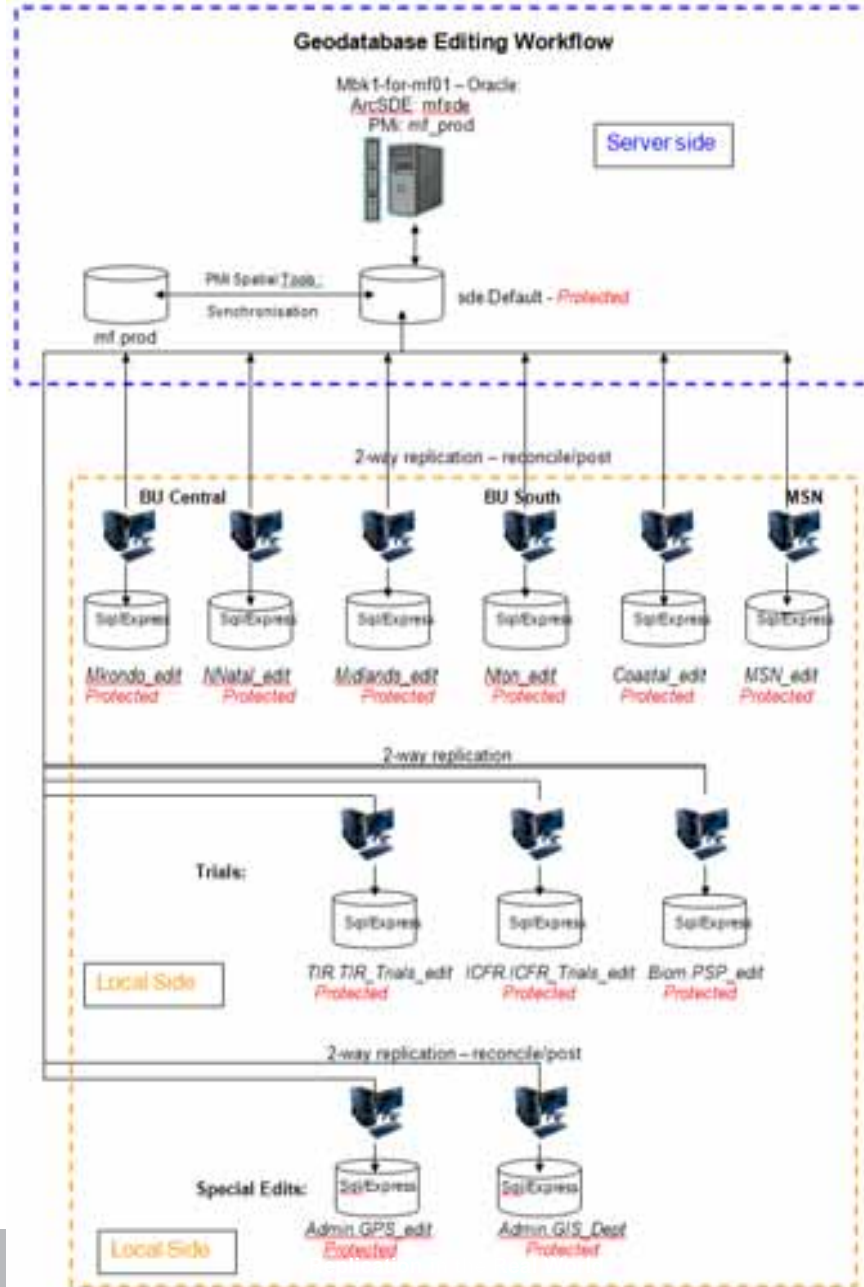
○ Very Poor Network Comms with Areas

- Most Area offices are in rural areas with very slow network speed.
- Delta tables are small enough to not create problems across network
- Major time savings/much more efficient editing can occur if done locally

○ Replication done as Geodata Service in ArcGIS Server

- Recommended that Replication be done as a Geodata Service in ArcGIS Server to overcome network speed problems.
- Not yet implemented – normal SDE database connection works well.
- Another option: Direct Connects – requires Oracle client on user PC

GIS Architecture: 2-Way Replication



GIS Functionality in Mondi



- Spatial and Attribute Data Maintenance Requirements in Mondi
 - Each Area is responsible for maintaining specific spatial and attribute data sets relating to their individual geographic entities. These include:
 - Compartments; Roads; Infrastructure (Fire Towers; Water points etc.)
 - The forest operations attribute data, which is linked to the spatial data via a Unique ID
 - The Forestry attribute data is managed using Syndicate's MicroForest Application
 - Spatial Data Maintenance:
 - Base Data (*e.g. national roads; rivers; contours*): Maintained by GIS Unit (Pmb)
 - Area Data (*Compartments; Roads; Dams; Infrastructure*): Maintained by Area Data Controllers (DBAs) – Area Foresters own the data; initiate changes. Data Controllers synchronise spatial data with attribute data in MicroForest
 - Data Storage:
 - Spatial data is stored in enterprise ArcSDE on Oracle (MFSDE Schema) on a server in Durban - DEFAULT version
 - Attribute data is stored in Microforest (MF) on Oracle (MF Schema) on same server
 - Spatial and Attribute data are synchronised using MF Spatial Tools application.
 - Link is by unique key LID – generated by MF and posted into spatial data during synchronisation process.

Data Maintenance Workflow



- 2-Way Replication System
 - **Step 1: Default to Local (Only one way of 2-way replication run)**
 - Area Editable Feature Classes are replicated to Local PC (Personal SDE on SQL Express)
 - Spatial data is subset to extract only data relevant to particular Area (Select by Graphic)
 - There are separate database connections, each unique to its Area or User Replica
 - **Step 2: Spatial Edits done locally**
 - Spatial edits are done as required to update spatial changes.
 - NB. Only spatial edits are done at this stage – no attribute edits!
 - Polygon and Line topology rules built into database are replicated down to Local version
 - Editors must run these topology checks prior to completing edits
 - **Step 3: Local to Default (Only one way of 2-way replication run)**
 - Delta tables are replicated back to Default on central server
 - Use made of Multiversioned Views
 - **Step 4: Synchronise Spatial and Attribute Data**
 - Using the MF Spatial Tools application, the spatial changes are synchronised with MF
 - Attributes in MF are updated/added/deleted as required.
 - LIDs are written to spatial data as required.
 - **Step 5: Steps 1 – 4 are then repeated.**

Database Design Issues











● Spatial Databases designed on “clean slate” basis

○ New System – Opportunity to redesign databases

- Applied successful principles of previous databases
- Enhanced design to improve data integrity; performance
- Added many new layers

○ Designed on Esri Geodatabase principles and recommendations:

- “The ESRI Guide to Geodatabase Design”- Michael Zeiler, Esri Press
- Use of Feature Datasets; Feature Classes
- Feature Datasets only used where required – Features with Topology; Relationship Classes
- NB! Everything in a Feature Dataset is included if any 1 feature in FDS is part of a replica
- Feature Classes grouped appropriately with “tag” in FC name

-  Mondi_Base_Data.DBO.HARVIMP_HARDSETTING_HAZARD
-  Mondi_Base_Data.DBO.HARVIMP_HARVEST_IMPACTS
-  Mondi_Base_Data.DBO.HARVIMP_WET_SOIL_HAZARD
-  Mondi_Base_Data.DBO.SA_FORESTRY_CO
-  Mondi_Base_Data.DBO.SA_WOOD_PROCESS
-  Mondi_Base_Data.DBO.SITE_CLIMATIC
-  Mondi_Base_Data.DBO.SITE_DEFAULT_SPECIES
-  Mondi_Base_Data.DBO.SITE_GENOZONE

○ Replicating Tables:

- Requires a Relationship class to be set up (in a Feature Dataset!)

System Issues & Limitations



- **Synchronising data with Third-Party Application/Database**
 - **Spatial Data has to link to Third-Party Attribute Database**
 - Third-party Vendor had written synchronisation module
 - Third-party App could only read Base Tables, not Adds/Deletes Tables
 - Issue resolved by using Multiversion View Tables of relevant layers to be synchronised
 - But, current MVV tables do not allow for Archiving on Default
 - Work-around: Archiving only enabled on local replicas – Most critical place!
 - **Compress to Base Table Issues:**
 - All Users access Database via “common” user name
 - Prevented ‘*Compress to Base Table*’ operations completing successfully
 - Resolved by using Multiversion View Tables
 - All edits were visible to third-party application without running a Compress
 - Did try using models to run Compress, but often failed due to network failures
 - **“Knock-on Effect” on other Systems and Databases:**
 - Need to review potential impacts on other Company systems and databases!
 - Interface software/links - new layer/table names
 - Involve IT Department!

Rigid Workflows, User Training Essential



● Strict adherence to published Workflows critical

○ Develop and Publish Detailed (Step-by-Step) Workflows

- Need to clearly understand Replication Process and Organisational Workflows
- Break each process down to a “Step-by-Step” Workflow
- Clearly document these – screen-shots of each process
- Decision-Tree process flows useful where choices are required
- Mondri wrote 75 Page Manual detailing every step required in the Editing/Replication/Synchronisation Workflow
- ~90% of problems due to not correctly following workflows

● Thorough User Training critical

○ Initial Training, Follow-up Training, Refresher Training, Ongoing Training!

- System in daily use for last 8 months – Still holding training sessions!
- Don't under-estimate the need for repeat training
- Initial 5 Day customised training course - Esri Instructor
- Refresher training on-site at the time of actual implementation
- Weekly visits to staff by Mondri Trainers (GIS Staff) for first two months
- Monthly visits to staff ongoing.
- However, amount of training may depend on how different new system is compared to system it is replacing

● Processes/Tools that are useful

○ Replication Log File

- An xml log file is generated each time a synchronisation is run
- ReplicaLog.dat
- Stored at C:\Users\user_name\AppData\Local\Temp\ReplicaLog.dat
- Details each layer synchronised and the time of replication
- Very useful to check if Replica succeeded or where problem occurred

○ Distributed Database Toolbar

- Do not provide this to Users – “Synchronise Changes” icon inserted in Editing Toolbar
- “Manage Replicas” Icon – Useful tools to manage Replicas
- Right-click on a Replica Name – “View Log”: Quick way to check if Replica succeeded
- However, can give “false positive”, i.e. Reports Succeeded, when it failed.
- Check ReplicaLog to make sure!
- Biggest risk with “Manage Replicas” – Users can inadvertently UNREGISTER Replica!
- Deletes Replica and must recreate from scratch
- Have not found a way to customise Toolbar with “Manage Replicas” Icon

○ Create Focussed Replicas

- “View-only” Replicas – One-way replication of layers for viewing purposes only
- Replicas involving sub-set of data or specific layers – one-way or two-way
- Replication process is flexible – can meet different needs across organisation

Acknowledgements:

Esri South Africa – Chris Byren; Dave Terblanche; Vassilo Walluschnig; Rudolf de Munnik
Mondi SA – Johan Wiese; Rob Woolley; Riette Richardson

Thank you very much

FORWARD - LOOKING STATEMENTS

It should be noted that certain statements herein which are not historical facts, including, without limitation those regarding expectations of market growth and developments; expectations of growth and profitability; and statements preceded by “believes”, “expects”, “anticipates”, “foresees”, “may” or similar expressions, are forward-looking statements. Since these statements are based on current knowledge, plans, estimates and projections, they involve risks and uncertainties which may cause actual results to materially differ from those expressed in such forward-looking statements. Various factors could cause actual future results, performance or events to differ materially from those described in these statements. Such factors include in particular but without any limitation: (1) operating factors such as continued success of manufacturing activities and the achievement of efficiencies therein, continued success of product development plans and targets, changes in the degree of protection created by Group’s patents and other intellectual property rights, the availability of capital on acceptable terms; (2) industry conditions, such as strength of product demand, intensity of competition, prevailing and future global market prices for the Group’s products and raw materials and the pricing pressures thereto, financial condition of the customers, suppliers and the competitors of the Group, potential introduction of competing products and technologies by competitors; and (3) general economic conditions, such as rates of economic growth in the Group’s principal geographical markets or fluctuations of exchange rates and interest rates.

Mondi does not

- a) assume any warranty or liability as to accuracy or completeness of the information provided herein
- b) undertake to review or confirm analysts’ expectations or estimates or to update any forward-looking statements to reflect events that occur or circumstances that arise after the date of making any forward-looking statements.