



An intelligent own-cloud

INDIGO

What is INDIGO?

A foundational repository system

Robust

Adaptive

Scalable

Origins

Developed by Archive Analytics Ltd

Valued partnerships with:

University of Liverpool

University of Maryland

Why Indigo?

- The archive exists ... applications access it.
- Agile Evolution of Capabilities & Services
- Low start-up costs
- Scalable + Cascadable

A Platform for -

- Preservation
 - Curation
 - Analytics
 - Discovery
 - Data Access
 - Teaching

“Open / Transparent / Extensible”

Future Proof

- Standards Based
- Open & Transparent ... no secrets
- Extensible
- Policy Capable
 - Triggers + Rules — choose your language

Accessible

Data presented as:

Objects

Rows (tables)

Facts (triples)⁻

...

— whatever is most natural

A Canvas for data sharing

e.g. Assemble a cloud of Repositories

- With their own rules
- Migrating and transforming data
- Shared or private
- Accessible via custom web-apps
- Distributed where-ever you want (or not)



Robust, Adaptive, Scalable

John Burns

john.burns@archiveanalytics.com

ANY QUESTIONS?

Details – Indigo.

- Based on modern , web-scale technologies — no secrets
- Default DB / storage is Cassandra — highly scalable
- Can utilise pre-existing Cassandra, or install its own.
- Hierarchical object store in conformance with the CDMI model.
- Directory of objects and metadata stored in DB.
- Default storage is compressed blocks in DB
- External object storage via URLs, incl. file:/ , others TBD.
- Rules / Action scripts stored in DB
- Rule execution on originating Node (locality of code/data).
- Tables appear as query-able objects in the repository,
 - presented in natural formats