NAIROBI / AUDIT & EVIDENCE REFERENCE

Connected city data

A published perspective on the data gap, city context and intended use.

Source and scope

This document reformats the complete published source. Historical results and current-tense wording belong to that source context. It is not a new rerun or independent certification. People, histories and model counts are synthetic.

Read before use

The connection gap

Cross-domain questions need identities and recorded relationships across systems.

City context

Synthetic residents connect to institutions, services and dated histories.

Use within scope

Analysis, testing and evaluation depend on published coverage and agreed rights. The document is a perspective, not empirical certification.

Request the full supporting notes from Fidelra

Full published source

Nairobi · Audit report

FCTE · Evidence Library

Published audit report · Read the scope with the result

Public positioning view.

This describes what the platform is and who it's for. Internal construction methods remain controlled.

The gap

Single-domain synthetic data answers single-domain questions

Most synthetic data today is generated one domain at a time - a synthetic health dataset here, a synthetic employment dataset there, each internally consistent but with no real relationship to the others. That's fine for testing a single system in isolation. It falls apart the moment the real question spans domains: does losing formal employment predict losing health coverage? Does housing informality change utility-connection odds? Does a household's composition affect land-dispute risk? Those questions need the same synthetic person to exist, consistently, across every domain their real counterpart would touch - which is what "connected" is doing the work of, in "connected synthetic city infrastructure."

Why a whole city

A city is the natural unit for connected data

A person's health, employment, housing, land, education and household aren't independent facts about them - they're facts about one life lived in one place, shaped by that place's real institutions, history and infrastructure. Modeling a "citizen" without also modeling the city they live in (which schools exist, which utilities reach which wards, what a real employer looks like) produces synthetic people who don't behave like real ones. A City Twin™ is the platform's answer to that: generate the city and its population together, from the same historical reference data, so the resulting connections are structurally real rather than coincidental.

Who it's for

Analysis, testing and decision-making without real individuals

Three groups get direct use from this: teams building or testing systems that need realistic, connected, at-scale data without any privacy exposure (no real person is ever represented, at any stage); researchers and analysts who want to explore cross-domain patterns - the kind of question that needs real relationships between records, not just realistic-looking individual tables; and evaluators who need a released dataset's rigor to be inspectable, not just claimed - which is the reason this Evidence Library, the Certification Explorer and the public audit trail exist as first-class parts of the product, not an afterthought.

This is a positioning document, not a technical specification - see Research Papers for the underlying methodology and Certification for the live, checkable results.

Keep the supporting record

Ask Fidelra for the report context, applicable release, source references and full supporting notes. Availability and rights for a paid data package are confirmed in the written quote.

Request full notes and supporting artifacts

fidelra.africa@fidelra.ai

stripeddonkey.com

Original HTML source is attached for reference. Public data-quality datasets are also attached where applicable.

Source page: https://stripeddonkey.com/cities/nairobi/audits/white-papers

Captured source SHA-256: 92477f31f1969f256385f68eb46715f18769667cbfd3d42c4b67d221575362fd