Release Status & Limitations

SEPAN v0.1

What the Kandy conference release is, what is ready to use, and what still requires local validation.

Release position: SEPAN v0.1 is a working conference release: a tested information structure, public demonstration, deployment-oriented Starter database and reference documentation. It is not a claim of national deployment, epidemiological validation or improved management outcomes.

1. What v0.1 is

- A Saint Lucia-derived framework for organising human-snake encounter information while preserving source reports, reviewed events, uncertainty and linked follow-up records.
- A working Record and Explore interface for demonstrating how the information can be entered, reviewed and analysed.
- A clean Starter database intended for configuration and testing by programmes that want to examine the approach in their own setting.
- A versioned reference point for discussion, pilot testing and collaboration after the World Snake Congress.

2. What is included

- Separate Reports and reviewed Encounters, with links between them.
- Responses, captures/checks, individual identifiers, releases/relocations, bites and media references as connected records rather than mandatory fields in every encounter.
- Uncertainty-aware dates, times, locations and species identification.
- Configurable localities, species lists and controlled code lists.
- Role-aware review concepts, traceable linked records and a deployment-oriented Grist Starter.
- Reference documents describing the information model, schema, setup and use.

3. Important boundaries

- Demonstration data is synthetic. They must not be treated as estimates of real-world encounter frequency, snakebite burden or geographic risk.
- Hosted interface components require internet access unless an organisation deliberately uses a pinned or offline build.
- Precise locations and sensitive records require local governance. The data model does not determine who should be permitted to see them.

4. Validation status

The current release has been developed and checked as an information-system design. That is different from validating a programme in real-world operation.

Already examined in v0.1	Still required in a real deployment
Report/Encounter separation and linkage logic	Usability with actual reporters, reviewers and responders
Synthetic multi-record workflows and linked follow-up records	Permissions, privacy, governance and access decisions
Configuration, code lists and uncertainty handling	Operational backup, recovery and local administration

Descriptive exploration questions and data-quality checks	Performance under local connectivity and workload
Starter deployment structure	Any required institutional, clinical or external-system integrations

5. How to treat the conference release

- Use the public demonstration to understand the workflow and analytical possibilities.
- Use the Starter only as a clean starting point for configuration and testing, not as a substitute for local implementation decisions.
- Test with fictional records before entering operational data.
- Agree on and document the data-try choices (in drop-down lists), permissions and governance before launch.
- Back up before structural changes. Future database migrations should be explicit, versioned and tested.
- Treat interface improvements and database migrations as different kinds of change: the interface may evolve without replacing accumulated operational data.