Why Data Contracts Matter

And what separates a contract that works from a document nobody reads

Every data leader knows the pattern. An upstream team renames a column or quietly changes business logic. Nothing fails loudly, then three days later the revenue dashboard is wrong, a forecast model has drifted, and analysts spend the week tracing the break the issue instead of delivering value for the organisation. The cost always lands downstream, on the people least able to prevent it. The root cause is not bad engineering; it is undefined interfaces to your data.

A Data Contract is a formal, versioned agreement between the producer of a dataset and its consumers, defining the schema, meaning, quality guarantees, service levels and ownership of the data, typically written in a machine readable form to support automated enforcement.

Why it matters now

- **Fewer incidents, caught earlier.** Changes are validated against the contract as part of continuous integration and development before they reach production. Breaking changes are blocked in development instead of only being discovered after the change has happened.
- **Clear accountability.** Every contracted dataset has a named owner and registered consumers. When something breaks, the right team knows within minutes.
- **Faster, safer change.** Producers can evolve systems confidently because downstream impact is visible and negotiated in advance, with notice periods for breaking changes.
- **Trust and reuse.** Data with no explicit guarantees often gets rebuilt, duplicated and second-guessed. Contracts are the foundation of dependable data products.
- **Audit and AI readiness.** Reliable, well-described data is the prerequisite for any serious AI initiative. Documented ownership, lineage and quality controls directly support regulatory obligations.

What makes a good data contract

A good data contract is:

- **Specific.** It covers schema and types, the business meaning of each field, quality rules (freshness, completeness, validity), SLAs, and a named accountable owner.
- **Negotiated.** It is an *agreement*, not a document published by IT and left on a sharepoint folder to be forgotten. Data consumers state what they need, producers commit only to what they can deliver.
- **Enforced.** It lives in version control as code, is validated automatically on every producer change, and is monitored continuously in production.
- **Versioned.** It has explicit rules for change: backward-compatible changes ship freely; breaking changes require a new major version, advance notice and a migration path.
- **Owned end to end.** Violations are treated like outages and should be logged, alerted, and resolved by the producing team with consumers notified automatically.

G Poeti Consulting · giovanni@poeti.co.uk · www.poeti.co.uk

The contract itself is the easy part. The value comes from the ecosystem around it: the ownership model, change-management processes, tooling and ways of working that need to be introduced to make data contracts an effective way for teams to govern and exchange data.

Organisations that start with a focused pilot on their highest-pain datasets typically see incident volumes fall within a quarter, allowing adoption to spreads on results rather than mandate.

Ready to make your data dependable?

I help organisations design and implement data contracts.

From defining the standard and operating model to a creating a working pilot with automated enforcement. Book a free 30-minute consultation to discuss where contracts would deliver the fastest return in your data estate.

G Poeti Consulting · giovanni@poeti.co.uk · **www.poeti.co.uk**