

The Provenance Manifesto

Decision provenance in the age of AI

Addressing organizational context amnesia through the art of provenance.

"The history of every decision is the architecture of our future."

The [Provenance Manifesto](#) explores a simple but increasingly important observation: software organizations are very good at preserving outputs, but very bad at preserving the reasoning that created those outputs.

For a long time, that gap was inconvenient but tolerable. In the age of AI, it becomes strategic. As machines help us produce more code, more plans, more architectures, and more decisions, the question is no longer only what we can build faster. The question is how we preserve the intent, context, and accountability behind what gets built.

That is the space where the idea of **decision provenance** comes from.

About Me

My name is **Yauheni Kurbayeu**.

- software engineering leader with decades of experience building and operating complex systems
- experienced in managing large engineering teams
- worked on enterprise platforms, large-scale SaaS products, and distributed engineering organizations
- increasingly focused not only on how systems are built, but on how the decisions behind them are made, preserved, and evolved

Throughout my career, I became increasingly interested not only in delivery itself, but in what happens to the reasoning that shapes systems after the work is done. Again and again, the same pattern appeared: organizations kept the implementation, but lost the decision history behind it.

Provenance is an exploration of that idea.

Why Provenance?

The Problem

Software teams preserve:

- code
- tickets
- deployments
- infrastructure
- documentation

But they rarely preserve:

- why a decision was made
- which trade-offs were accepted
- which assumptions shaped it
- who owned it

We keep the artifact. We lose the reasoning.

This is the core problem. Most organizations preserve almost everything around delivery, but the most important part often disappears: the logic behind the decisions.

What trade-off was accepted? What assumption was true at the time? Who made the call? What alternatives were rejected? Months later, the artifact is still there, but the reasoning is gone.

As a result, teams inherit systems without inheriting the logic that shaped them.

→ We keep the artifact

We lose the reasoning.

Why This Matters Now

For a long time, this was a delivery problem.

Now AI is turning it into a governance problem.

AI is making the **"what"** cheaper:

- code
- plans
- architectures
- outputs

So the **"why"** becomes more valuable:

- intent
- constraints
- accountability
- decision ownership

AI accelerates execution. Provenance preserves accountability.

Teams have always lost context, repeated discussions, rebuilt the same workarounds, and rediscovered past decisions. AI changes the stakes because it dramatically lowers the cost of producing artifacts.

When code, plans, and solution options become cheaper, the scarce layer shifts upward. Intent, trade-offs, constraints, and ownership become more important than ever. Without them, an organization can generate more output while understanding less and less about why that output exists.

The Core Idea

Decisions should become first-class artifacts.

A meaningful decision should preserve:

- the problem
- the context
- the assumptions
- the alternatives
- the reasoning
- the owner
- the later evolution

Decision provenance explains why a system exists in its current form.

Today, decisions usually live as side effects of conversations. They appear in meetings, chats, design reviews, and maybe later in documentation, but rarely as durable, structured knowledge in their own right.

A meaningful decision should preserve more than the final statement. It should retain the problem being solved, the context and constraints, the alternatives considered, the reasoning behind the choice, the owner of the decision, and how that decision later evolved.

Decision provenance is therefore about making the **"why"** of the system durable, not only the **"what."**

Decisions should become *first-class artifacts.*

Undocumented intuition

Traceable decisions

What Provenance Changes

It moves us:

- from undocumented intuition to traceable decisions
- from private memory to institutional memory
- from hidden assumptions to transparent reasoning
- from static documentation to evolving decision history
- from uncontrolled automation to governed human-AI collaboration

This is not just better documentation. It is a memory layer for reasoning.

Once decisions are treated as first-class artifacts, a deeper shift happens. Knowledge stops depending on who happens to remember the story. Instead, reasoning becomes durable, traceable, and reusable.

This is especially important in AI-assisted environments. The challenge is no longer only to store what happened. It is to preserve how conclusions were formed, under which assumptions, and with whose authority.

Provenance creates that shift from fragile personal memory to structured institutional memory.

→ *This is not just better documentation.*
It is a memory layer for reasoning.

The Principles

1. **Decisions are first-class artifacts**

Every architecture, product behavior, operational process, and incident response originates from decisions. Organizations should treat decisions with the same rigor as source code.

2. **Decisions must carry context**

A decision without its assumptions, alternatives, risks, and reasoning is incomplete. Context transforms a decision from a statement into knowledge.

3. **Decisions evolve but are never erased**

Decisions may be revised, superseded, or branched, but their history must remain preserved. Organizational intelligence grows through the evolution of reasoning, not through rewriting history.

4. **Decisions must be queryable**

Organizations should be able to ask why a system was designed this way, what assumptions justified it, and which risks were accepted. Institutional memory must be structured so those answers can be retrieved instantly.

5. **Decisions must be attributable**

Every meaningful decision must have ownership. Accountability enables trust, governance, and responsible change.

6. **AI must operate within decision governance**

Artificial intelligence can generate code, architecture, and solutions, but it must operate within a traceable decision framework where reasoning, assumptions, and approvals remain visible. AI accelerates execution. Provenance preserves accountability.

7. **Institutional memory is a strategic asset**

Organizations that preserve the reasoning behind their systems move faster, avoid repeating mistakes, and make better decisions. Memory compounds. For modern organizations, decision provenance becomes intellectual capital.

Memory compounds.

These principles translate the core idea into a governance model for modern engineering organizations. They move the conversation away from isolated documents and toward living, attributable, queryable reasoning that can evolve without disappearing.

→ *Memory compounds.*

The Value

Decision provenance helps organizations:

- reduce rework and repeated rediscovery
- onboard people faster
- make architectural change safer
- expose hidden assumptions
- improve incident, audit, and compliance reasoning
- give AI access to the reasoning layer, not only the artifact layer

The more AI generates, the more provenance matters.

In practice, this creates value at multiple levels. It reduces waste caused by rediscovering old context. It makes onboarding easier because new people can understand why a system looks the way it does. It makes change safer because teams can revisit the original reasoning instead of guessing. It also improves auditability and incident analysis by preserving not only what happened, but why certain paths were chosen.

In AI contexts, the value becomes even clearer. Models can often explain what the code does or what the docs say, but without provenance they usually cannot explain why a trade-off was accepted, why a workaround exists, or why one option was rejected.

**The more AI generates,
*the more provenance matters.***

What This Looks Like In Practice

Start by preserving major decisions around:

- architecture trade-offs
- scope and product choices
- incident responses
- policy and risk boundaries
- AI-assisted implementation plans

For each important decision, capture:

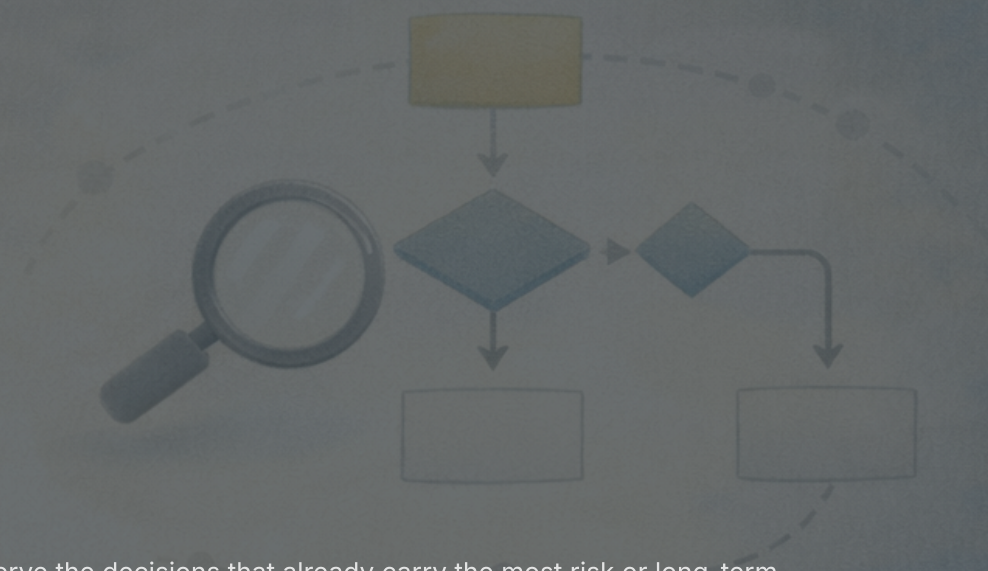
- why it was made
- what alternatives were considered
- what assumptions and risks existed
- who owns it
- how it changed over time

This does not require waiting for a perfect platform. A practical starting point is to preserve the decisions that already carry the most risk or long-term impact.

In simple terms:

1. Capture the decision.
2. Capture its context, assumptions, risks, and alternatives.
3. Assign ownership.
4. Preserve the history when the decision changes.
5. Make it queryable later.

The exact format can vary. Markdown, structured records, graphs, ADR-like documents, or other systems can all help. The important change is not the format. It is the intention to preserve reasoning in a reusable way.



Remembering
decisions

Deciding with
context

The Real Shift

The hardest part is not technical.

It is cultural.

Provenance is a mindset shift:

- decisions should not disappear into conversations
- reasoning should not remain private
- AI should not operate outside visible governance

Provenance is not primarily about tools. It is about how organizations remember.

This is why provenance should not be misunderstood as just another tooling proposal. The technical implementation matters, but the deeper challenge is deciding that reasoning itself deserves to become part of the system.

That means moving from hidden context to shared context, from undocumented authority to traceable ownership, from static artifacts to evolving reasoning, and from outputs without lineage to accountable decisions.

In many ways, the technical systems are the easier part. The harder part is changing how organizations think about knowledge, decisions, and responsibility.

Provenance clears the path to progress.

Closing

The real divide is not:

- human-written vs AI-generated

The real divide is:

- **opaque outputs** vs **traceable reasoning**

If we cannot reconstruct how a conclusion was formed, then even a polished result is still just a well-presented guess.

This applies to human work, AI work, and increasingly to the hybrid systems we are all building together.

Without preserved reasoning, AI can optimize the system while the organization gradually loses the ability to explain itself. At that point, humans do not necessarily lose their jobs first. They lose their authority over the **"why."**

Discussion

- Where does your organization lose decision context today?
- Which decisions are most worth preserving?
- How should AI participate in governed organizational memory?

Let's build the future with provenance.

Thank You

Thank you for exploring the Provenance Manifesto.

The goal is not to present a finished product, but to invite a broader conversation about how engineering organizations can preserve memory, govern reasoning, and remain accountable in the age of AI.



Learn more at
ProvenanceManifesto.org