

On the historic roots of Agentic AI in Semantic Web Services

Axel Polleres, Frederik Bauer, Daniel Dobriy, Tobias Käfer, Lukas Kubelka, Andreas Harth, and Thomas Wehr



The 2001 “Semantic Web” vision was based on agents:



The Semantic Web

A new form of Web content that is meaningful to computers will unleash a revolution of new possibilities

by [TIM BERNERS-LEE](#), [JAMES HENDLER](#) and [ORA LASSILA](#)

The entertainment system was belting out the Beatles' "We Can Work It Out" when the phone rang. When Pete answered, his phone turned the sound down by sending a message to all the other *local* devices that had a *volume control*. His sister, Lucy, was on the line from the doctor's office: "Mom needs to see a specialist and then has to have a series of physical therapy sessions. Biweekly or something. I'm going to have my agent set up the appointments." Pete immediately agreed to share the chauffeuring. At the doctor's office, Lucy instructed her Semantic Web agent through her handheld Web browser. The agent promptly retrieved information about Mom's *prescribed treatment* from the doctor's agent, looked up several lists of *providers*, and checked for the ones *in-plan* for Mom's insurance within a *20-mile radius* of her *home* and with a *rating of excellent or very good* on trusted rating services. It then began trying to find a match between available *appointment times* (supplied by the agents of individual providers through their Web sites) and Pete's and Lucy's busy schedules. (The emphasized keywords indicate terms whose semantics, or meaning, were defined for the agent through the Semantic Web.)



In a few minutes the agent presented them with a plan. Pete didn't like it—University Hospital was all the way across town from Mom's place, and he'd be driving back in the middle of rush hour. He set his own agent to redo the search with stricter preferences about *location* and *time*. Lucy's agent, having *complete trust* in Pete's agent in the context of the present task, automatically assisted by supplying access certificates and shortcuts to the data it had already sorted through.

*"[...] The **agent** promptly retrieved information about Mom's prescribed treatment from the doctor's agent, looked up several lists of providers, and checked for the ones in-plan for Mom's insurance within a 20-mile radius of her home and with a rating of excellent or very good on trusted rating **services**[...]"*

- *Appointment detection in emails*
- *services*
- *Discovery (Semantic Search)*

Happy anniversary!

... has then mostly become the basis for the "Web of Data"...

*"If HTML and the Web made all the online documents look like one huge **book**, RDF, schema and inference languages will make all the data in the world look like **one huge database**"*

Tim Berners-Lee, Weaving the Web, 1999

... and has been popular since under the term of “Knowledge Graphs”

25 years later: “Agentic AI”



NBC NEWS WORLD CUP POLITICS U.S. NEWS WORLD LOCAL BUSINESS SHOPPING TIPLINE

ARTIFICIAL INTELLIGENCE

Bot web traffic has overtaken human web traffic, data shows

Cloudflare says 57.4% of requests to a selection of websites it hosts are now automated bot requests, while 42.6% are human-generated.


So, are we done?

90's

Web "pages"

HTML

Web "forms"



How did we get here?




2026

Web "pages"

Web APIs

JSON-RPC




90's

Web "pages"

HTML

Web "forms"



(What) can we learn?


"tool-calling" LLMs

MCP
A2A

Web "pages"


Web APIs

JSON-RPC



2026

Rewind to ~2000...



90's

Web "pages"

HTML

Web "forms"

Early 2000's

Semantic Web "idea"

XML

Web Services

XML-RPC/SOAP

WSDL

UDDI

BPEL

Generation 1 — WS-* : standardised plumbing



- *Built on XML and Web standards. Strong on messaging and syntactic interfaces, weak on automation. Driven by Microsoft, IBM, HP, Oracle, SAP via W3C and OASIS.*
- **Messaging — SOAP.** XML remote-procedure calls tunnelled over HTTP POST (W3C Rec. 2003/2007).
- **Description — WSDL:** operations + message grammar (XML Schema) + transport bindings. Syntactic, no “semantic” description of the service, mainly interface description.
- **Discovery — UDDI:** registry of white / business / technical pages; public registries all shut down before 2010.
- **Composition — BPEL:** executable *orchestrations* by a single requester; *choreography* (WS-CDL) stalled.
- **Mediation:** mostly hard-coded (e.g. XSLT in BPEL) — no automated reconciliation.

Also, important other **conceptualizations**:


e.g. What is a **service**? [Preist, 2004] “*the provision of a **value** or product in the context of the domain of application **through interactions**, by **one party to another**”*”

Generation 1 — WS-* : What did we learn?

- Messaging — SOAP ... Maybe too much overhead?
- Description — WSDL ... description of the interface needed
- Discovery — UDDI ... a trusted registry is probably a good idea!
- Composition — BPEL/WS-CDL... important distinction:
orchestration vs. *choreography*
- Mediation — XSLT ... “hard-coded” mediation did not really scale.




vs.




Generation 1 — WS-* vs. “Agentic AI”

- Messaging — SOAP ... JSON-RPC
- Description — WSDL ... “example-based” calls, “natural language” descriptions of RPC
- Discovery — UDDI ... **not (yet) standardized? Not open: “App-store” model?**
- Composition — BPEL ... MCP: basically still only “*orchestration*” (by LLM+ReACT)”
- Mediation — XSLT *Solved?... works surprisingly well*
... but potentially prone to “injections”!



Generation 1 — WS-* vs. “Agentic AI”



- Mediation — XSLT *Solved?...* works surprisingly well
... but potentially prone to “injections”!

90's

Web "pages"

HTML

Web "forms"

Early 2000's

Semantic Web "idea"

XML +**RDF**+**OWL**

Web Services


~ 2005

Semantic Web Services

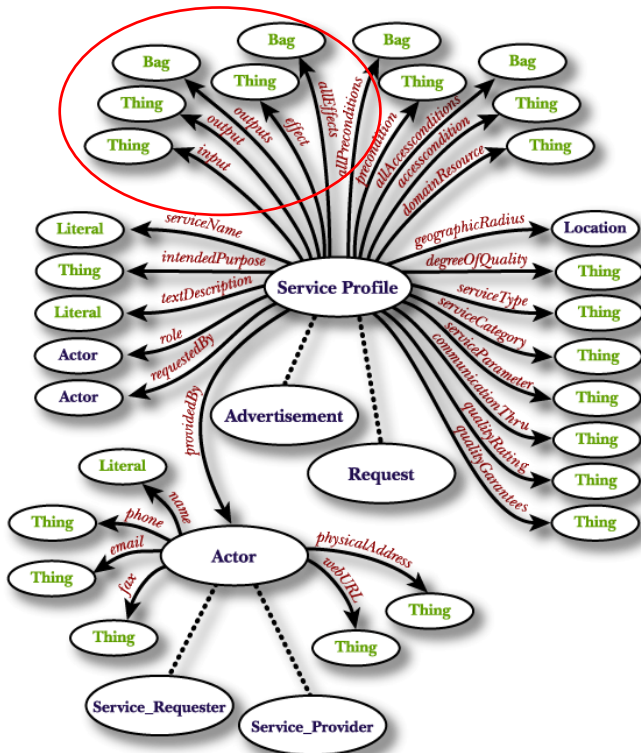
OWL-S
WSMO

Generation 2 — Semantic Web Services


- **Semantic Web Services:** We cannot only describe Websites and Metadata, but also Services and APIs! (~2002-2007)



- Main Idea:
 - Describing **Services & Agents** in a **declarative manner**
 - using ontologies...
 - ... should enable **automated composition, execution**



On hindsight: automated composition by purely symbolic inference ... was maybe maybe ahead of its time...



Generation 2 — Semantic Web Services (SWS)

- *Same SOAP/WSDL base, enriched with formal semantics so agents could reason. Funded by DARPA (DAML) and the EU. Aimed squarely at automation.*
- **Description — OWL-S & WSMO:** semantic descriptions of **capabilities** plus **inputs, outputs, pre-conditions, effects (IOPEs)** in OWL / WSML (W3C Member Submissions, 2004–05).
- **Discovery:** ontological reasoning matches goals to capabilities by logical subsumption.
- **Composition:** formal descriptions feed symbolic AI planners that generate orchestrations automatically.
- **Mediation:** WSMO's reusable mediator components; in OWL-S a mediator is just another service.
Assumption: Solvable by declarative reasoning

Generation 2 — SWS : What did we learn?

- Description ... cold-start-problem: Who provides/generates formal descriptions?
- Discovery ... (complex) logical reasoning for discovery did not scale
- Composition ... OWL-S essentially just a Semantically enriched “BPEL” (orchestration only), but IOPEs should allow for some degree of **automated planning!**
- Mediation WSMO’s reusable mediator components; in OWL-S a mediator is just another service.
Assumption: Solvable by declarative reasoning — **in theory... but not in practice** 😞

In summary, mostly academic: it stayed a research endeavour; industry adoption never followed.

Generation 2 — SWS : What did we learn?

Composition: there's more than sequential ReACT-based plans:

Sequential planning: STRIPS, etc.

Hierarchical Planning: hierarchical task network planning (HTN): typically ordered decompositions of complex goals into sequences of sub-goals, with planners such as e.g. SHOP and its successors.

Conditional Planning (e.g. CNLP, MBP) plans conditional trees, allowing for non-deterministic actions with contingencies,

Conformant planning, i.e. only returning deterministic, non-branching plans that work under all possible contingencies (where possible),

*Remarkable, one early conformant planning tool was called “**GPT**” (General Planning Tool).*


B. Bonet and H. Geffner. 2001. GPT: a tool for planning with uncertainty and partial information. In Proc. IJCAI-01 Workshop on Planning with Uncertainty and Partial Information, A. Cimatti, H. Geffner, E. Giunchiglia, and J. Rintanen (Eds.). Seattle, WA, 82–87.

Marco Pistore, Fabio Barbon, Piergiorgio Bertoli, Dmitry Shaparau, and Paolo Traverso. 2004. Planning and monitoring web service composition. In Proceedings of the 11th International Conference on Artificial Intelligence: Methodology, Systems, and Applications (AIMSA). Springer, 106–115.

Why Semantic Web Services stalled (~2008)

- The last OWL-S and WSMO releases date from around 2008. Reference engines such as WSMX never left the prototype stage. Why?
- **High authoring cost:** writing the required formal semantic descriptions was hard, expert work — a steep tax on every service.
- **Too few descriptions:** sparse annotations meant registries and reasoners had little to work with; the network effect never ignited.

Fast forward to 2025...



Generation 3 - Today's agentic stack: MCP & A2A

MCP — Model Context Protocol (Anthropic)

- let a model call and orchestrate external tools.
- tool descriptions in *natural language* + JSON Schema for I/O — close in spirit to WSDL, but informal.
 - **Architecture not too different to WS-* (message passing).**

A2A — Agent2Agent (Google)

- Agents expose 'cards' and 'skills' so they can call one another (choreography flavor!).
- Again: natural-language capability + JSON Schema parameters.

Common thread:

messaging is still JSON-RPC;

Natural language descriptions are expressive but ***semantically ambiguous***.

Three generations, five dimensions:

	WS-*	Semantic Web Services	MCP / A2A (today)
Messaging	SOAP / XML over HTTP	SOAP / WSDL base	JSON-RPC, message passing
Description	WSDL — syntactic only	OWL-S / WSMO — formal, IOPEs	NL + JSON Schema — ambiguous
Discovery	UDDI registries (defunct)	Reasoning-based matchmaking	Mostly a fixed, predefined tool set
Composition	BPEL orchestration	AI planning over descriptions	ReAct / Plan-and-Solve — central, opaque
Mediation	Hard-coded (XSLT)	Explicit mediator components	Implicit, handled inside the LLM

LLMs “broke the bottleneck” — at a price

What we gained

- Automation never reachable before — **discovery, composition** and **mediation** without hand-authored formal descriptions.
- Cheap **descriptions**: a few natural-language examples replace ontologies.
- **Mediation** on the fly: the model bridges mismatched tools from a few examples.

What it cost

- **Reduced transparency**: why a tool was chosen, or how a plan formed, is hard to inspect.
- **Lack of control**: no declarative contracts; no clean way to debug or renegotiate failed / redundant calls. No guarantees:
 - when tools are called,
 - whether mediation works
 - how “discovery” works
- **Ambiguity**: prose capabilities can be misread; the model’s choices stay implicit.

Open challenges & routes ahead

Safe dynamic discovery

- Agents mostly use a fixed tool set — below UDDI’s ambition. Yet, dynamic discovery would need trusted, searchable repositories.

Transparency & contracts

- Bring back declarative control: explicit capability semantics and contracts.*

<https://arxiv.org/abs/2605.19186>

Discoverable Agent Knowledge — A Formal Framework for Agentic KG Affordances (Extended Version)*

Terry R. Payne¹[0000-0002-0106-8731], Valentina Tamma¹[0000-0002-1320-610X],
and Enrico Daga²[0000-0002-3184-5407]

Beyond central orchestration

- ReAct / Plan-and-Solve assume a single controller. Decentralised, “planning peer-to-peer choreography” of agents with distinct goals is still open.

Decentralised coordination?

- The Web is permissionless — anyone can link without asking: resource-centred, hypermedia-driven, decentralized coordination (discovery, composition, policy negotiatio) among peers is a road not yet taken.

Open challenges & routes ahead

Safe dynamic discovery

- Agents mostly use a fixed tool set — below UDDI’s ambition. Yet, dynamic discovery would need trusted, searchable repositories.

Transparency & contracts

- Bring back declarative control: explicit capability semantics and contracts.*

Beyond central orchestration

- ReAct / Plan-and-Solve assume a single controller. **Decentralised**, “planning peer-to-peer choreography” of agents with distinct, (conflicting?) goals is still open.

The Web?

- The Web is permissionless — anyone can link without asking: resource-centered, hypermedia-driven, decentralized coordination (discovery, composition, policy negotiation) among peers is a road not yet taken.

Thank you! *Questions & discussion welcome.*

In memory of Dieter Fensel

who envisioned key principles of agentic AI through his pioneering work on Semantic Web Services, decades ago.



The presenter was funded by:

