

**Building agents is easy.
Running them is not.
How Flock and Dapr make AI
agents production-ready**

Tilman Sattler & Martin Brandl


white duck


Who we are





Martin Brandl

(CTO, Cloud Solution Architect)

 +49 8031 230159-0

 martin.brandl@whiteduck.de

 @martin_jib

 www.linkedin.com/in/mbrandl



Tilman Sattler

(Software and AI Engineer)

 tilman.sattler@whiteduck.de

 -

 -

Agenda

- From Chatbots to Agents
- Agent Loops, Tools & State
- Flock: Blackboard-Based Orchestration
- Dapr: Cloud-Native Runtime Building Blocks
- Using Dapr State Stores in Flock
- Demo: From Local Development to Production-Ready Agents
- Key Takeaways

20 Years of Software Evolution

- **Cloud** changed where software runs.
- **Containers** changed how we package it.
- **Kubernetes** changed how we scale it.
- **DevOps** changed how we deliver it.

HOW
we build software

AI changes **WHAT** software can do.

Beyond Chatbots

AI changes what software can do:

**From answering questions
to completing tasks.**

Chatbots respond.

Agents work toward a goal.

The shift is not from chatbots to smarter chatbots.

The shift is from conversation to execution.

What Makes AI Agentic?

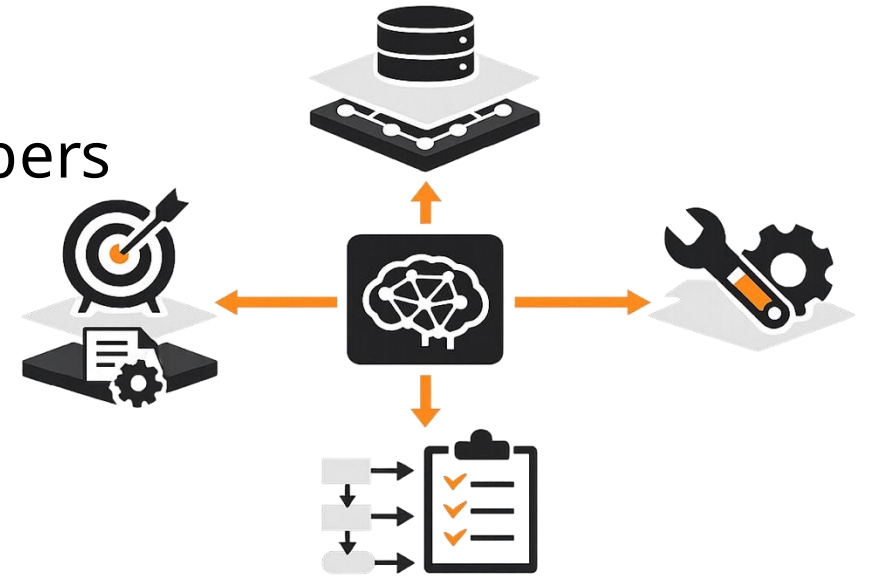
A model can answer.

Agentic AI can act - because it has:

- **Goal** — what should be achieved
- **Context & State** — what it knows and remembers
- **Tools** — what it can use
- **Instructions** — how it should behave

The model is only one part.

The system around it makes it agentic.



The Agent Loop

Agentic AI is not a single model call.

An agent repeatedly:

- **Observes** the current context
- **Reasons** about the next step
- Uses a **tool**
- Updates its **state**
- Continues or stops



A loop is not reliable by default.

Every step must be persisted, managed, and made recoverable.

Why State Becomes Critical

For a demo, state can live in memory.

For production, state must be:

- **Durable** — survive restarts
- **Traceable** — show what happened
- **Recoverable** — continue after failures
- **Secure** — protect data and secrets
- **Portable** — switch environments

Agents do not just need memory.
They need **operational state**.

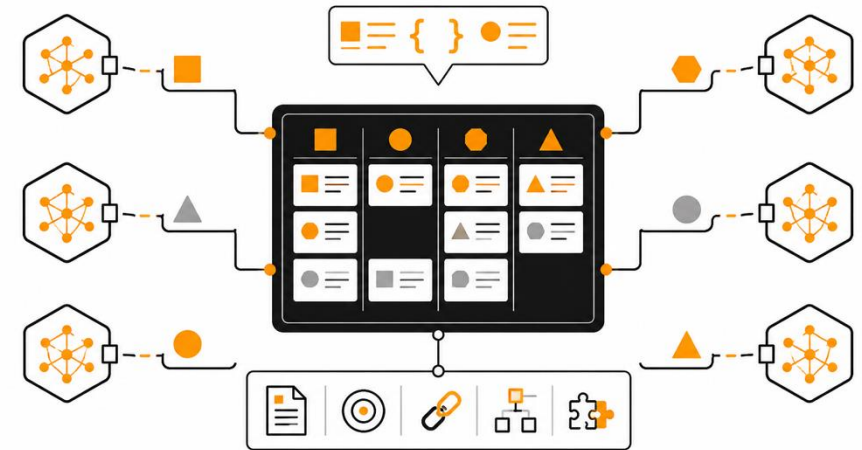


Flock: Blackboard-Based Orchestration

Agents do not call each other directly.

They collaborate through a blackboard:

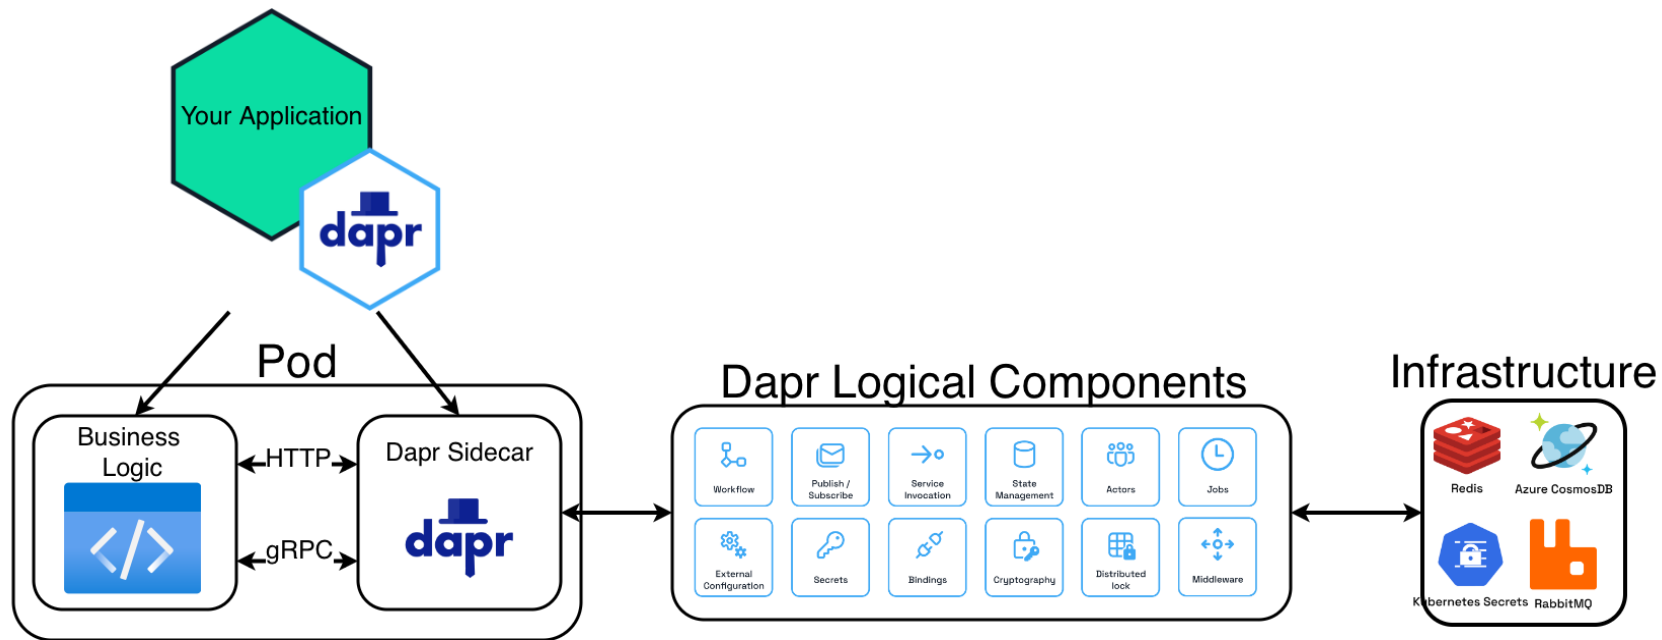
- **Typed contracts** — what agents produce
- **Semantic subscriptions** — who reacts
- **Loose coupling** — agents stay independent
- **Flexible workflows** — no rigid graph required



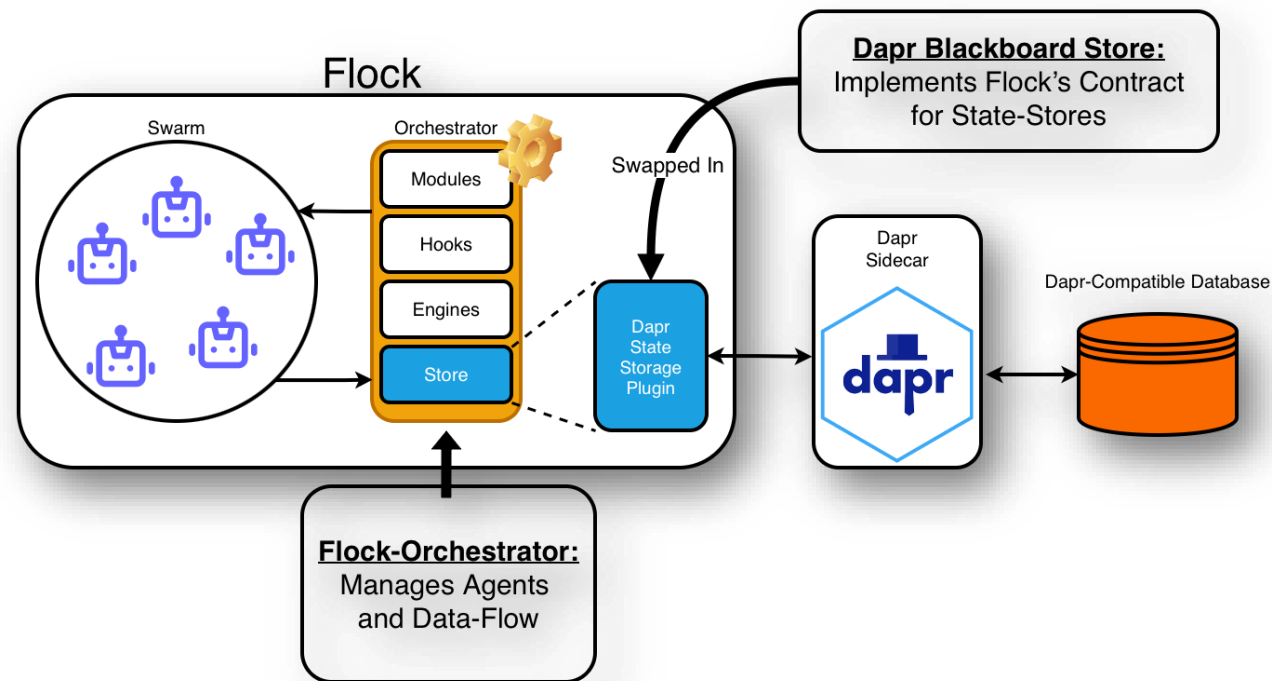
Blackboard = shared context for agent collaboration.

Dapr: Conceptual Overview

- **Dapr** = **D**istributed **A**pplication **R**untime
- Portable, event-driven, modular
- Building blocks handle common cloud-native tasks
- Provides a **unified interface** for underlying services



How Flock and Dapr go Hand-in-Hand



Flock's Blackboard models:

- Message-flow from agent to agent.
- The state s_t of an agent-team / agent swarm at a point in time t .
- Flock's blackboard represents the **state** of the entire application!

Dapr State-Stores are the perfect fit.

Using Dapr State Stores in Flock

- Dapr and Flock share the same modular design-philosophy.
- Thank's to Flock's modular architecture, integration is a breeze.
- A new component just needs to **adhere to Flock's contract-system.**

```
from flock import (
    Flock,
    DaprStateBlackboardStore,
    DaprStateBlackboardConfig,
)
from dapr.clients import DaprClient


# Load Secrets from Dapr
dapr_secrets = load_dapr_secrets()

# Configure the store
store_config = DaprStateBlackboardConfig(dapr_secrets)

# Initialize the Blackboard Store
dapr_store = DaprStateBlackboardStore(config=store_config)










# Intialize the swarm
flock = Flock(
    model="azure/gpt-5.5",
    store=dapr_store,
)
```

What you get out of the Box



Component	CRUD	Transactional	ETag	TTL	Actors	Workflow	Status	Component version	Since runtime version
Aerospike	✓	☐	✓	☐	☐	☐	Alpha	v1	1.0
Apache Cassandra	✓	☐	☐	✓	☐	☐	Stable	v1	1.9
CockroachDB	✓	✓	✓	✓	✓	☐	Stable	v1	1.10
Couchbase	✓	☐	✓	☐	☐	☐	Alpha	v1	1.0
etcd	✓	✓	✓	✓	✓	✓	Beta	v2	1.12
Hashicorp Consul	✓	☐	☐	☐	☐	☐	Alpha	v1	1.0
Hazelcast	✓	☐	☐	☐	☐	☐	Alpha	v1	1.0
In-memory	✓	✓	✓	✓	✓	✓	Stable	v1	1.9
JetStream KV	✓	☐	☐	☐	☐	☐	Alpha	v1	1.7
Memcached	✓	☐	☐	✓	☐	☐	Stable	v1	1.9
MongoDB	✓	✓	✓	✓	✓	✓	Stable	v1	1.0
MySQL & MariaDB	✓	✓	✓	✓	✓	✓	Stable	v1	1.10
Oracle Database	✓	✓	✓	✓	✓	✓	Beta	v1	1.7

What you get out of the Box

1. Zero code-changes between development and prod. 
2. Integration with all Dapr-supported State Stores. 
3. Distributed state for multiple Flock-swarms. 
4. Choice of consistency for state (Strong vs. Eventual). 
5. Transactional updates for Blackboard-State. 
6. Encryption of Data-at-Rest via Dapr. 
7. Apply Dapr Resiliency-Policies for your underlying storage. 
8. Secrets stored in supported Secret-Stores. 
9. Vendor-Agnostic Business-Logic. 

Demo



What are these Agents working on?

Find out more at:

<https://agent-challenge.whiteduck.de>



Thank You!

