

76-89-D Agent Specifications: Current vs Future

From Cognito-Static Design to Cognito-Dynamic Operations

The Hidden Assumption Behind Modern Agent Design

Most agent frameworks today focus on a relatively small set of questions:

- What is the agent trying to accomplish?
- What resources can it use?
- What authority does it possess?
- What constitutes success?

These questions are entirely reasonable. In fact, they mirror the way software systems have been specified for decades. The underlying assumption is - if the agent performs its intended mission correctly, the specification is complete.

This assumption worked well in deterministic environments. It becomes increasingly incomplete as systems evolve toward:

- autonomy,
- persistence,
- adaptation,
- collaboration,
- emergence.

The Great Discovery

A careful examination of agent specifications reveals something important.

- Not all specifications belong to the same category.
- Some characteristics are largely known before deployment.
- Others only become visible after deployment.
- This naturally divides agent specifications into two domains:

Cognito-Static Specifications
and
Cognito-Dynamic Specifications

Cognito-Static Specifications

These describe properties that are largely known before the agent executes.

They are designed, reviewed, approved, tested. Their values remain relatively stable throughout the life of the agent.

The Cognito-Static categories are:

- | | |
|-------------|--|
| • Mission | What is the agent trying to accomplish? |
| • Resources | What tools, APIs, databases, and models are available? |
| • Authority | What actions is the agent allowed to perform? |
| • Outcomes | How is success measured? |

Collectively these define: ***What the agent is supposed to do.***

Cognito-Dynamic Specifications

These describe properties that only become visible through operation. They emerge through interaction with the cognitive neighborhood:

- other agents,
- workflows,
- APIs,
- enterprises,
- sync zones,
- cognitive fields.

These properties cannot be fully understood during design. They must be observed, measured, and managed during runtime.

The Cognito-Dynamic categories are:

- Boundaries Where does the agent actually travel?
- Coupling Who does the agent actually interact with?
- Consequences What effects does the agent actually generate?
- Stability How much instability does the agent create or consume?
- Context What operational environment is the agent currently experiencing?
- Governance How does cognition actually flow through the ecosystem?

Collectively these define: ***How the agent coexists within a living cognitive ecosystem.***

The Ten Categories Revisited

Category	Type
Mission	Cognito-Static
Resources	Cognito-Static
Authority	Cognito-Static
Outcomes	Cognito-Static
Boundaries	Cognito-Dynamic
Coupling	Cognito-Dynamic
Consequences	Cognito-Dynamic
Stability	Cognito-Dynamic
Context	Cognito-Dynamic
Governance	Cognito-Dynamic

Design Time vs Runtime

A different way to view the same distinction is:

Category	Design-Time Visibility	Runtime Visibility
Mission	High	Low
Resources	High	Medium

Category	Design-Time Visibility	Runtime Visibility
Authority	High	Medium
Outcomes	High	Medium
Boundaries	Medium	High
Coupling	Low	High
Consequences	Low	High
Stability	Low	High
Context	Medium	High
Governance	High	High

This reveals something profound:

- The first four categories are primarily specified.
- The last six categories are primarily discovered.

Why Current Agent Frameworks Feel Complete

Most current agent platforms focus heavily on Cognito-Static specifications.

Examples include:

- prompts,
- objectives,
- permissions,
- tool access,
- success criteria.

These are all important. But they answer only one question: *What should the agent do?*
They largely ignore: *What becomes possible because the agent exists?*

The HCAS Perspective

As autonomous systems scale, many of the most important characteristics become runtime phenomena.

Examples include:

- | | |
|-----------------------------|---|
| • Boundary Pressure | Which boundaries are routinely approached? |
| • Passportless Coupling | Which interactions emerged that were never explicitly designed? |
| • Consequence Amplification | What downstream effects are being generated? |
| • Stability Consumption | How much stabilization capacity is being borrowed? |
| • Context Drift | How is the operating environment changing? |
| • Governance Topology | How is cognition flowing through the ecosystem? |

These are not design-time questions. They are operational questions.

Functional Specification vs Stability Specification

The Cognito-Static and Cognito-Dynamic categories naturally separate into two specification families.

Functional Agent Specification (FAS)

Defines: What the agent should do.

Includes:

- Mission
- Resources
- Authority
- Outcomes

This is the dominant focus of today's agent engineering.

Stability Agent Specification (SAS)

Defines: How the agent must coexist within a Hyper Complex Adaptive System.

Includes:

- Boundaries
- Coupling
- Consequences
- Stability
- Context
- Governance

This becomes the primary focus of Stability Engineering and Cognitive OT.

Why Cognitive OT Emerges

The Cognito-Static categories are largely managed by : *Agent Designers / Cognitive IT*

The Cognito-Dynamic categories are largely managed by: *Cognitive Operations / Cognitive OT*

This distinction explains why GUDIYA cannot simply be another feature embedded inside an agent framework. The two domains are fundamentally different.

- One concerns capability.
- The other concerns ecosystem health.

The Aviation Parallel

Early aircraft specifications focused on:

- speed,
- range,
- payload.

These were static design characteristics.

As aviation matured, additional concerns emerged:

- turbulence,
- flutter,
- stability margins,
- control authority,
- operational envelopes.

These became operational concerns.

The same transition is now occurring in autonomous cognition.

The Deep Insight

The Cognitive Era introduces a shift in engineering focus.

Traditional software engineering concentrated on

: Cognito-Static Reality.

The emerging discipline of Stability Engineering focuses on

: Cognito-Dynamic Reality.

This is where:

- emergence,
- coupling,
- instability,
- consequence propagation,
- cognitive drift,

actually live.

Final Insight

The first generation of agent specifications was sufficient for deterministic software and early autonomous systems.

The next generation must account for the reality that many of the most consequential properties of a cognitive ecosystem cannot be fully specified before deployment.

They emerge only through interaction.

Therefore:


- Mission, Resources, Authority, and Outcomes define capability.
- Boundaries, Coupling, Consequences, Stability, Context, and Governance define coexistence.

The first group is Cognito-Static.

The second group is Cognito-Dynamic.

And the transition from one to the other may ultimately mark the boundary between classical

- agent engineering and the
- emerging discipline of Stability Engineering for Hyper Complex Adaptive Systems.



Book Series Coming Soon ..

