



The Trust Infrastructure

HOW ENTERPRISE CONTEXT MAKES AI SAFE, USEFUL, AND REPEATABLE





Gabe Keith
Director – Content Intelligence
Engineering



Rafal Hryniewicz
Manager – Content Intelligence
Product



The Trust Infrastructure

Agents NEED Context

Breakout – Backwards Build A Bot

Share learnings and wrap up



Who has agents running in production?

- Add Live Poll updates



In the world of AI, **agents**
make decisions with guided
autonomy rather than
programmed workflows



How can agents get the
right data, at the **right time**
to make **effective decisions**?



AI agents rely on context to make better decisions



Relevant documents, data tables, previous messages, background information, constraints, policy

What is context?

- Information you provide to an AI model alongside your question or instruction
- Acts as the AI's "working memory" for that specific conversation

Why does context matter?

- ✓ **Accuracy:** Better inputs lead to better outcomes
- ✓ **Relevance:** Helps AI focus on what matters
- ✓ **Control:** Grounds responses in real information
- ✓ **Consistency:** Keeps answers aligned with your rules



Making Context Concrete

Context is the trust layer that sits on top of your content.



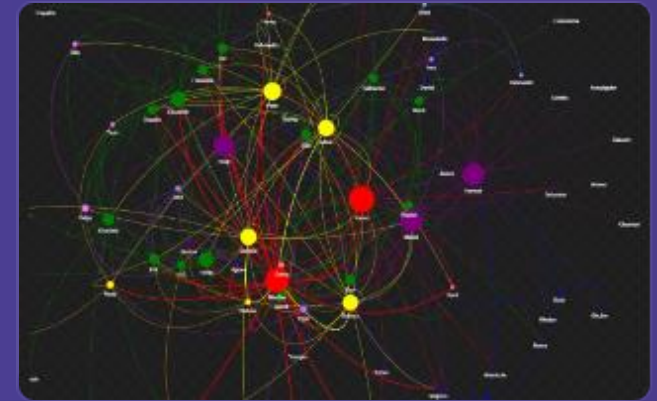
Business Definitions

Standardizing business logic, glossaries, and metadata across disparate silos to establish a single source of truth.



Permissions & Rules

Integrating ACLs, versioning, retention policies, and regional compliance constraints natively.



Relationships

Mapping the invisible threads that connect critical documents to specific customers, cases, and products.



AI Trust Maturity Model: context makes AI safe, useful, and repeatable

“Rewrite this email”

Blind Assistant



LEVEL 1

Public LLM – Smart model
Hallucination Risk
Security Risk

“What is our PTO Policy?”

Semantic AI



LEVEL 2

Getting Smarter!
Flattened Context
Prep via pipeline
Limited Access Control
Can't Handle Complex Questions

“What documents are Available for this applicant?”

Tool Agents




LEVEL 3

Getting Agential!
Common Connectors Available
Can Retrieve Like a Human
Basic Authentication
No Unified Content View
Brittle Integrations
Still Security Gaps

“What is this applicant's relationship and history with our Company?”

Knowledge Graph Agents



LEVEL 4

You are in the top 5 percent!
Centralized Knowledge Graph
Business Logic Embedded
Complex Relationships Identified
Multi-Doc Reasoning Unlocked
Agent Access Restrictions
User ACL's Not Present
Agent Traversal Complications

“Do my normal research For John Doe”

KG + ACL + Agent Memory



LEVEL 5

The Holy Grail of Trust!
Access Control Lists
Permission Inheritance
User IAM/Access mappings
Agent Learns the Graph
Agent Understands Preferences
Enterprise Security
Data Governance
User Trust Skyrockets
Industry Compliance

Complexity and Trust Increase



Breakout Session

Backwards Build-A-Bot




1

Your Company



2

What problem do you want to solve?




3

Define how it works at Level 5 AI Trust Maturity?




4

Define how it works at Level 4 AI Trust Maturity?



5

Level 3




BASE	1000	100	25	200	100
1	44	3.2	8	33	100
CO	10			1000	100
1	2	21	01	970	100



6

Level 2




HYPERCUBE



7

Level 1



How does Hyland Solve this Problem?



Hyland Enterprise Context Engine

Trusted AI that cites sources, takes safe actions and compounds learning, leading to faster decisions, fewer errors and lower cost.

Enterprise knowledge graph

Ontologies that map the business

Short- and long-term episodic memory

Semantic memory

Questions?
Visit us at Booth _____



