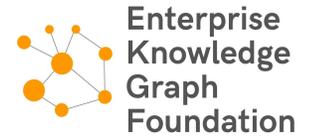


Enterprise Knowledge Graph Foundation

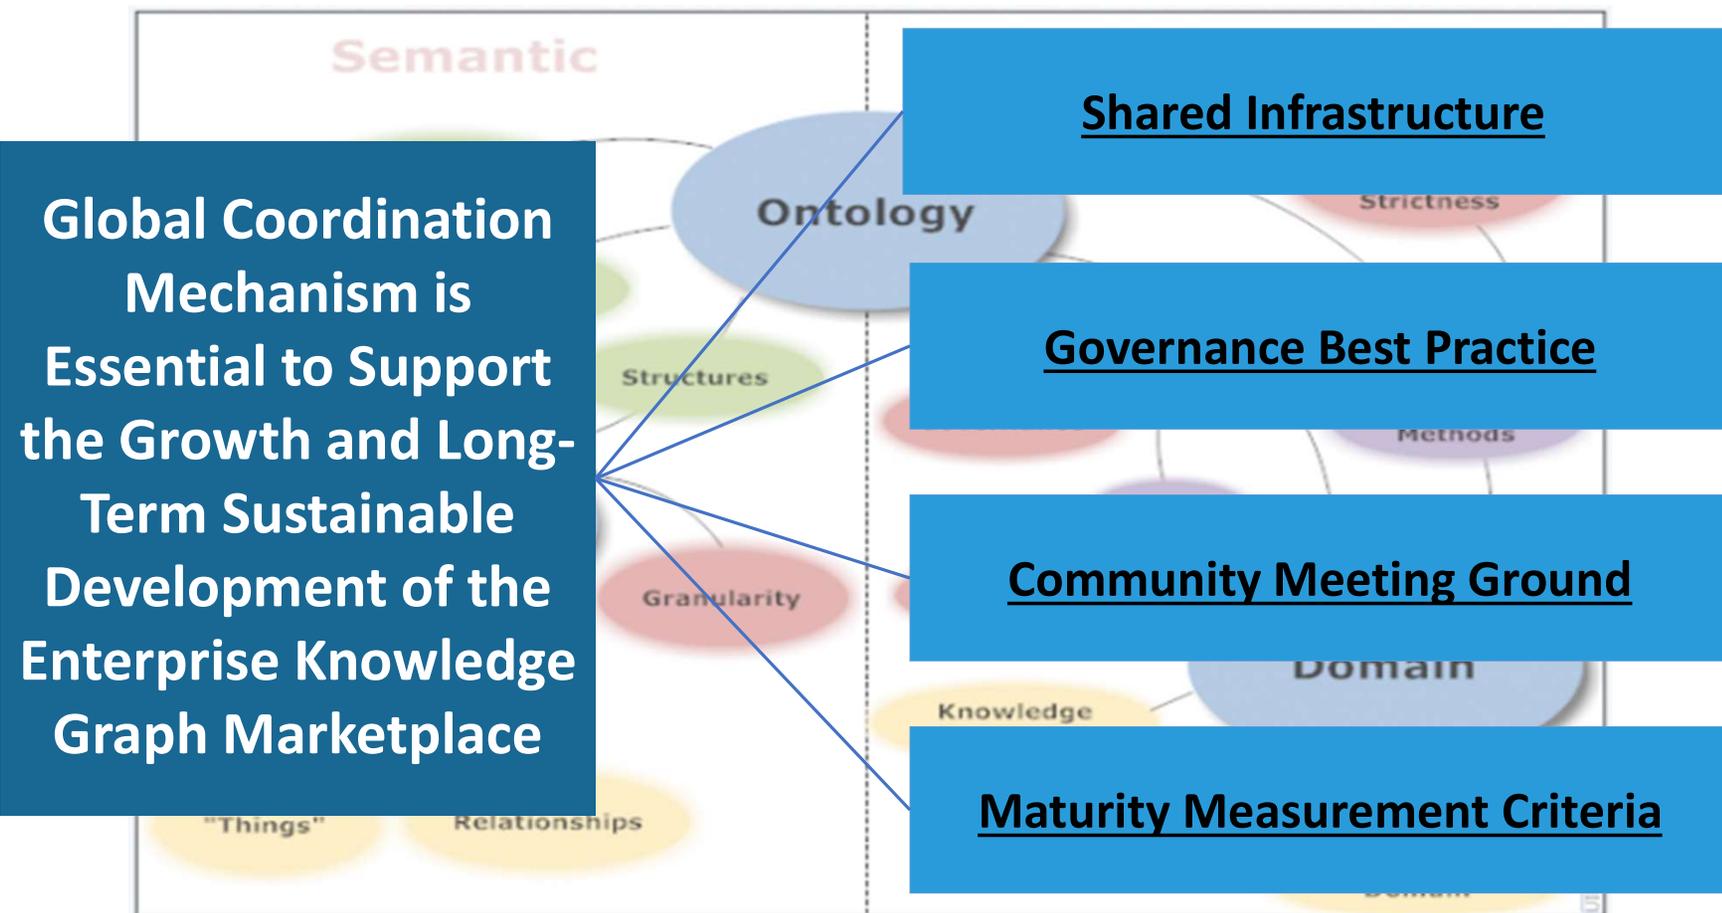
Global Advocates for Semantic Technology

EKG FOUNDATION MEETING AGENDA

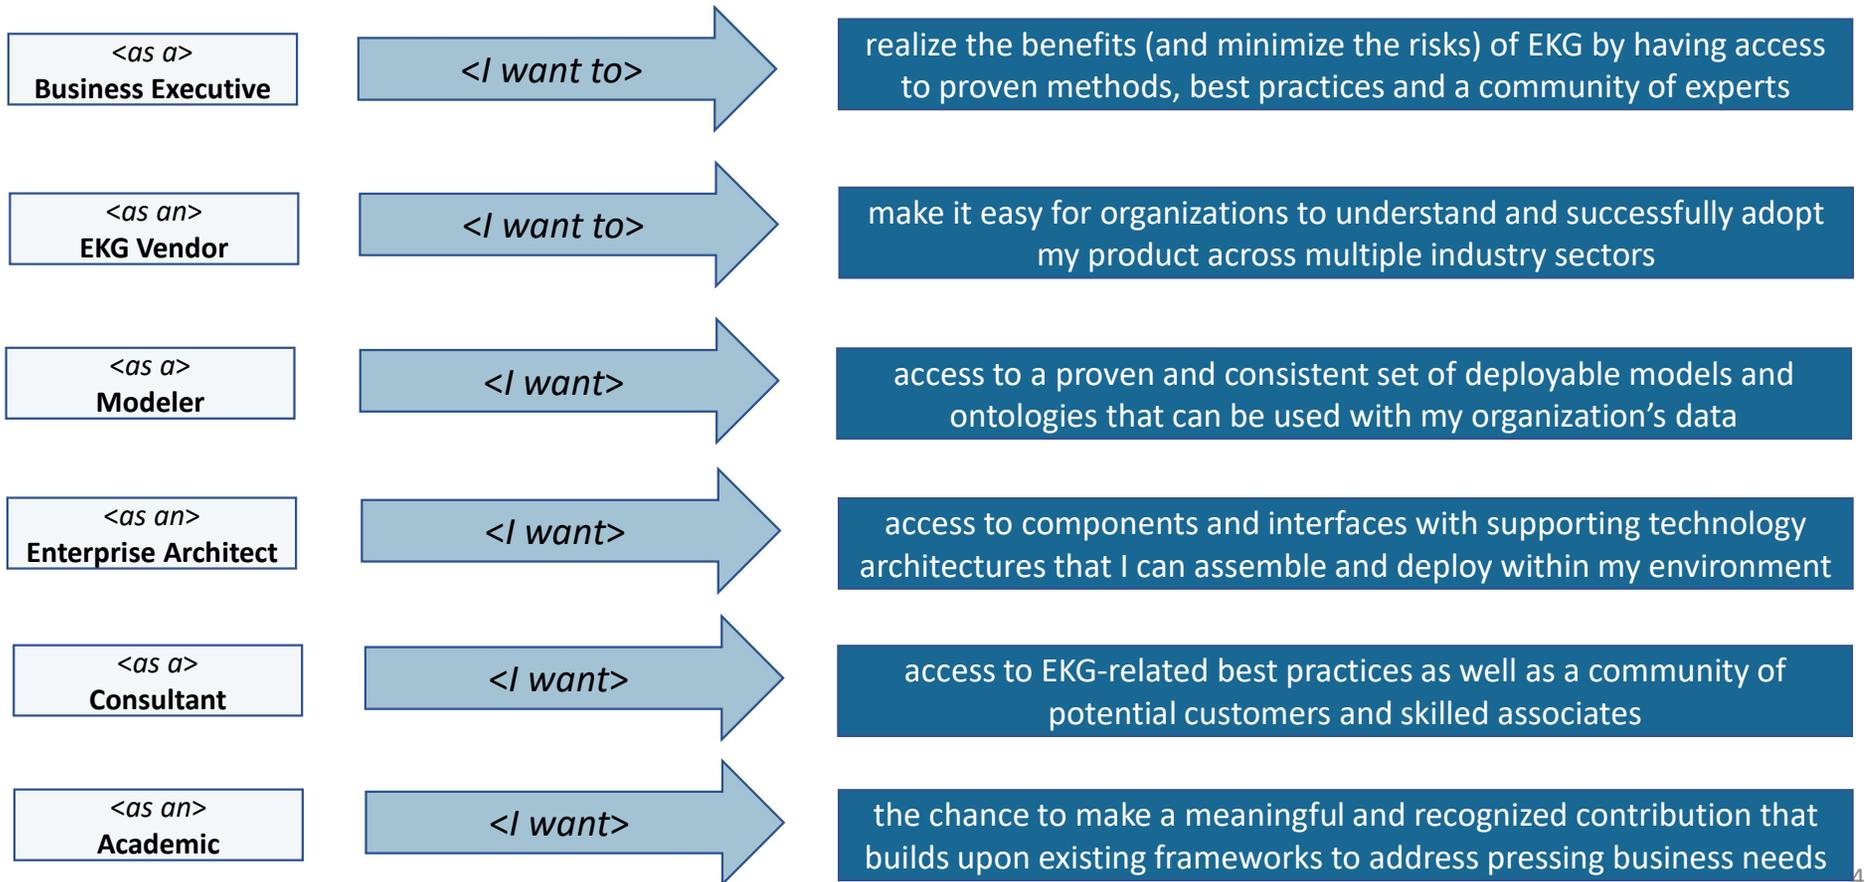
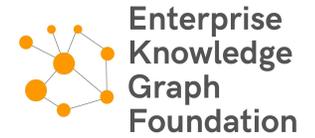


1. Background and Introductions
2. Lessons from the Financial Industry Business Ontology (FIBO)
3. Mission and Core Principles of the EKG Foundation
4. Enterprise Knowledge Graph Maturity Modeling
5. EKG Foundry (Use Case, Ontology, Best Practice and Tools portals)
6. Participant Interaction (Q&A)

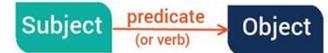
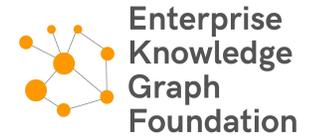
EKG FOUNDATION



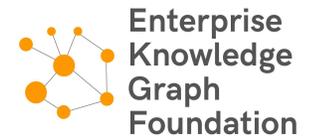
EKG Foundation Marketplace



BACKGROUND



EKG FOUNDATION LEADERSHIP



Jacobus Geluk: *Chief Technology Officer, agnos.ai. Delivered the first EKG platform at scale into production in the financial industry.*



Dean Allemang: *Twenty years of research and development in semantic solutions, knowledge management and rule-based systems. FIBO lead ontologist.*



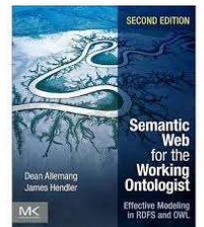
Dennis Wisnosky: *Former Chief Technology Officer in the US Department of Defense. Standards implementation process for FIBO.*



Pete Rivett: *Chief Technology Officer of Adaptive Inc. OMG Architecture Board, Ontology Definition Metamodel and Meta Object Facility. FIBO Leadership team.*



Michael Atkin: *Thirty years as strategic analyst to financial institutions and regulators. Founder Enterprise Data Management Council. Faculty, Columbia University.*

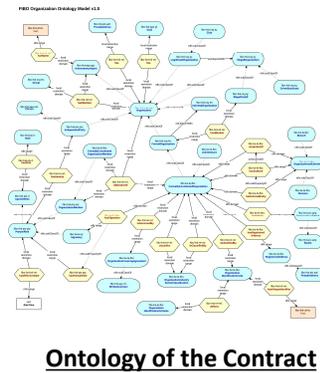
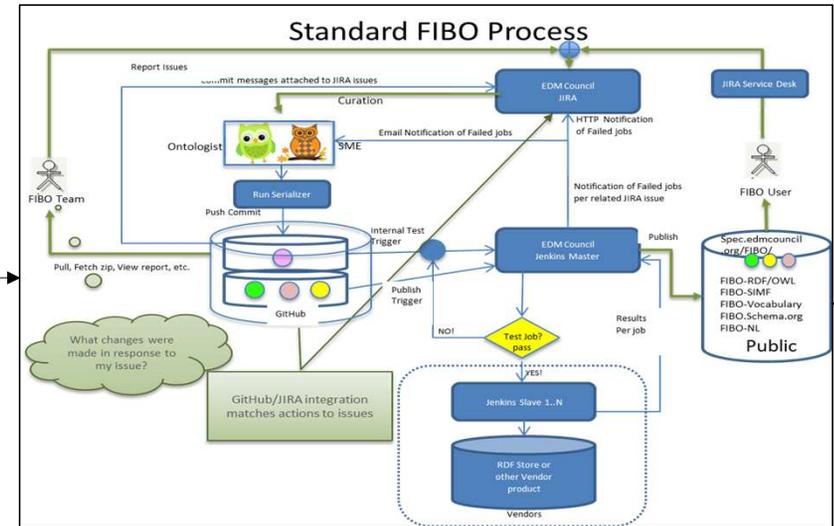
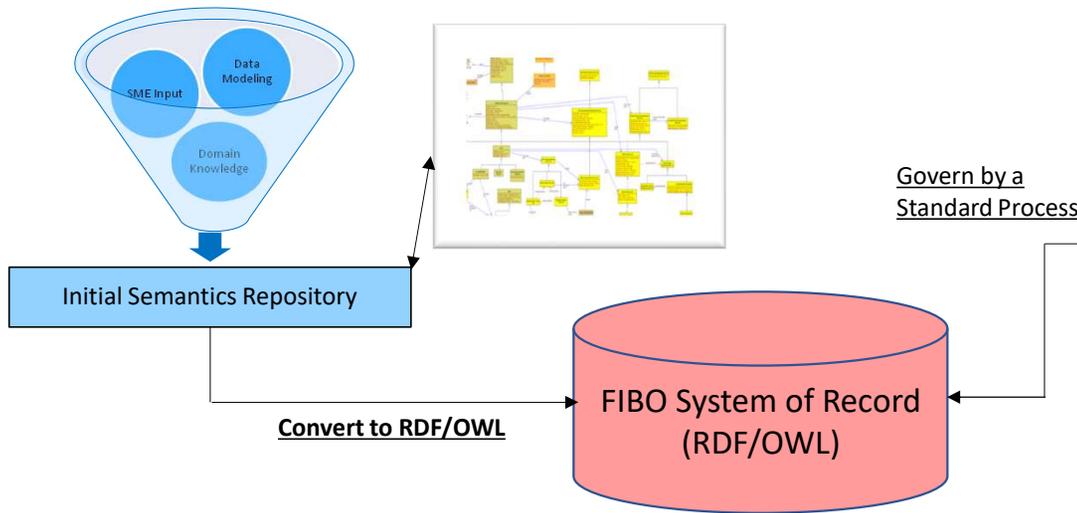


Historical Precedents

Lessons from the Financial Industry Business Ontology (FIBO) Process

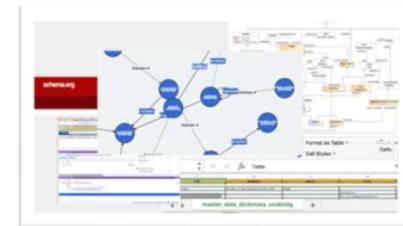
LESSONS FROM FIBO

Financial Industry Business Ontology



ABSOLUTES

- A Proven documented and automated methodology/process (BTDM for Continuous Agile Development)
- Early validation and production of useful products (e.g. FIBO Vocabulary, visual OWL)
- POC's that can be made operational
- Encourage early adopters



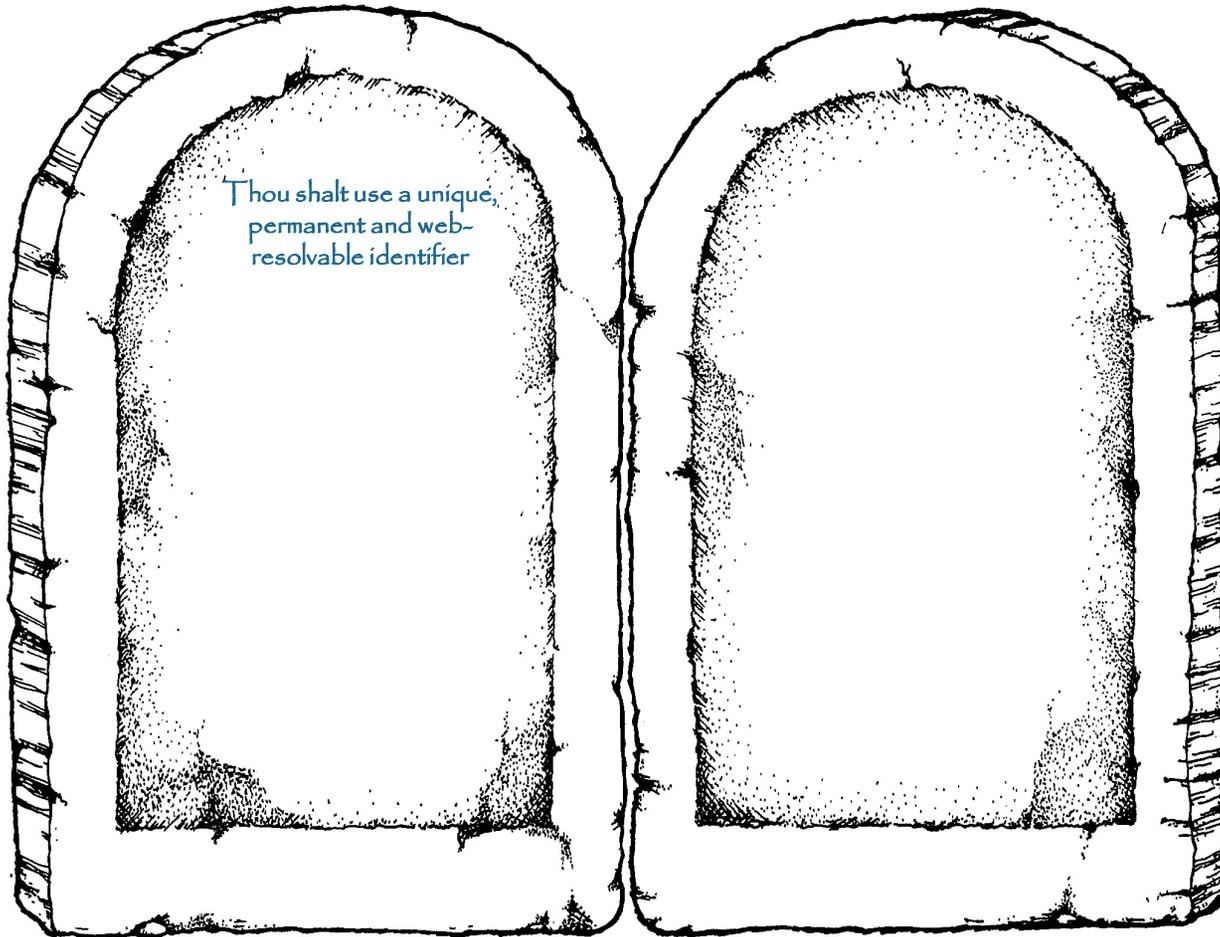
Principles of Enterprise Knowledge Graph

Core Concepts Required for Sustainable EKG Development

CHARACTERISTICS OF AN ENTERPRISE KNOWLEDGE GRAPH

- ★ **Smart Data** (*IRI, ontology, standard data type, open standards, designed for interoperability*)
- ★ **Transparency** (*multiple contextual viewpoints of data to expose inconsistencies and demonstrate provenance*)
- ★ **Automation** (*supports machine learning and reasoning, powers event-driven data processing*)
- ★ **Reuse** (*extensible for a wide range of use cases, shortens time-to-market*)
- ★ **Stability** (*defined access methods, logically separated services, no single points-of-failure*)

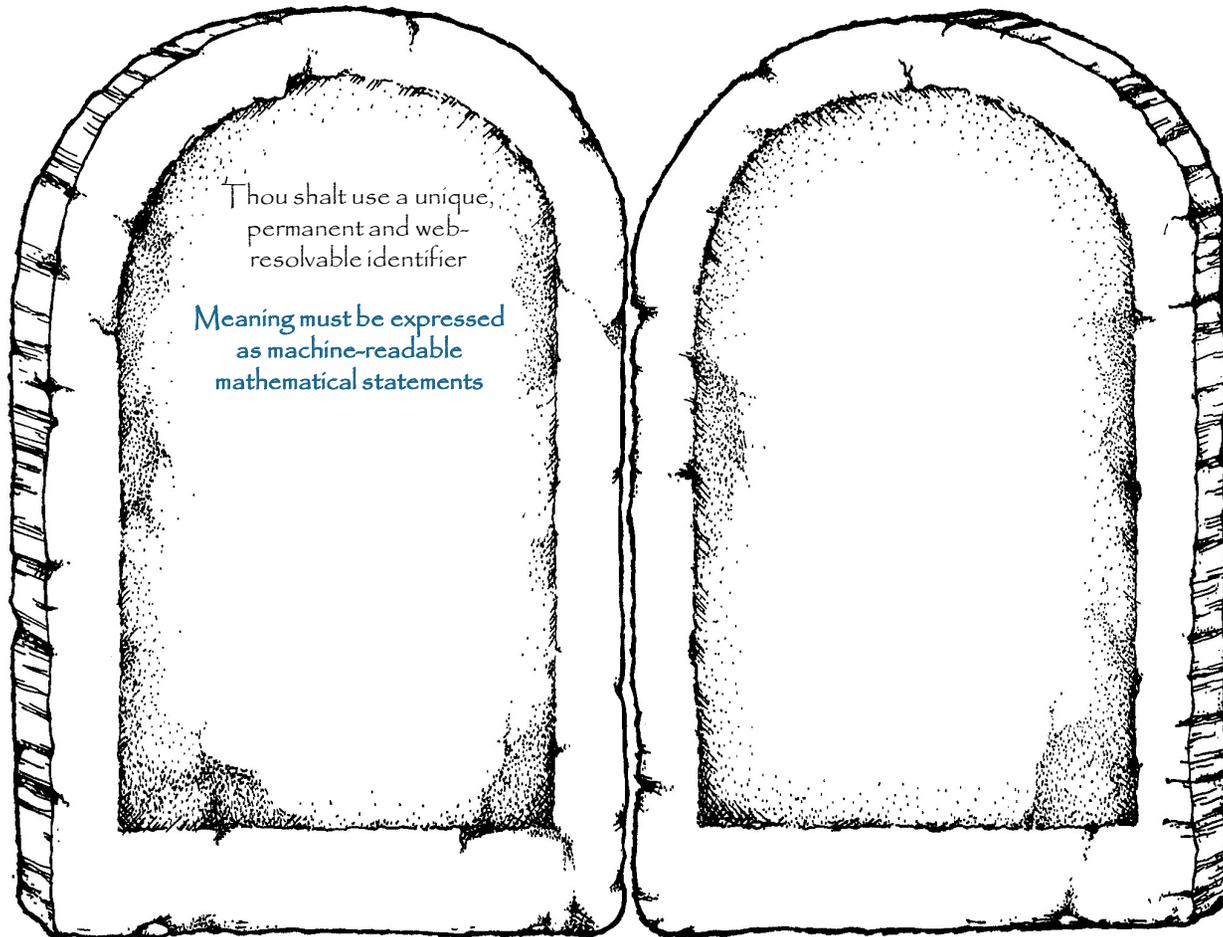
GUIDING PRINCIPLES OF EKG



Principle 1: Identity

Any given object is identified with at least one universally unique, opaque, permanent and web-resolvable identifier.

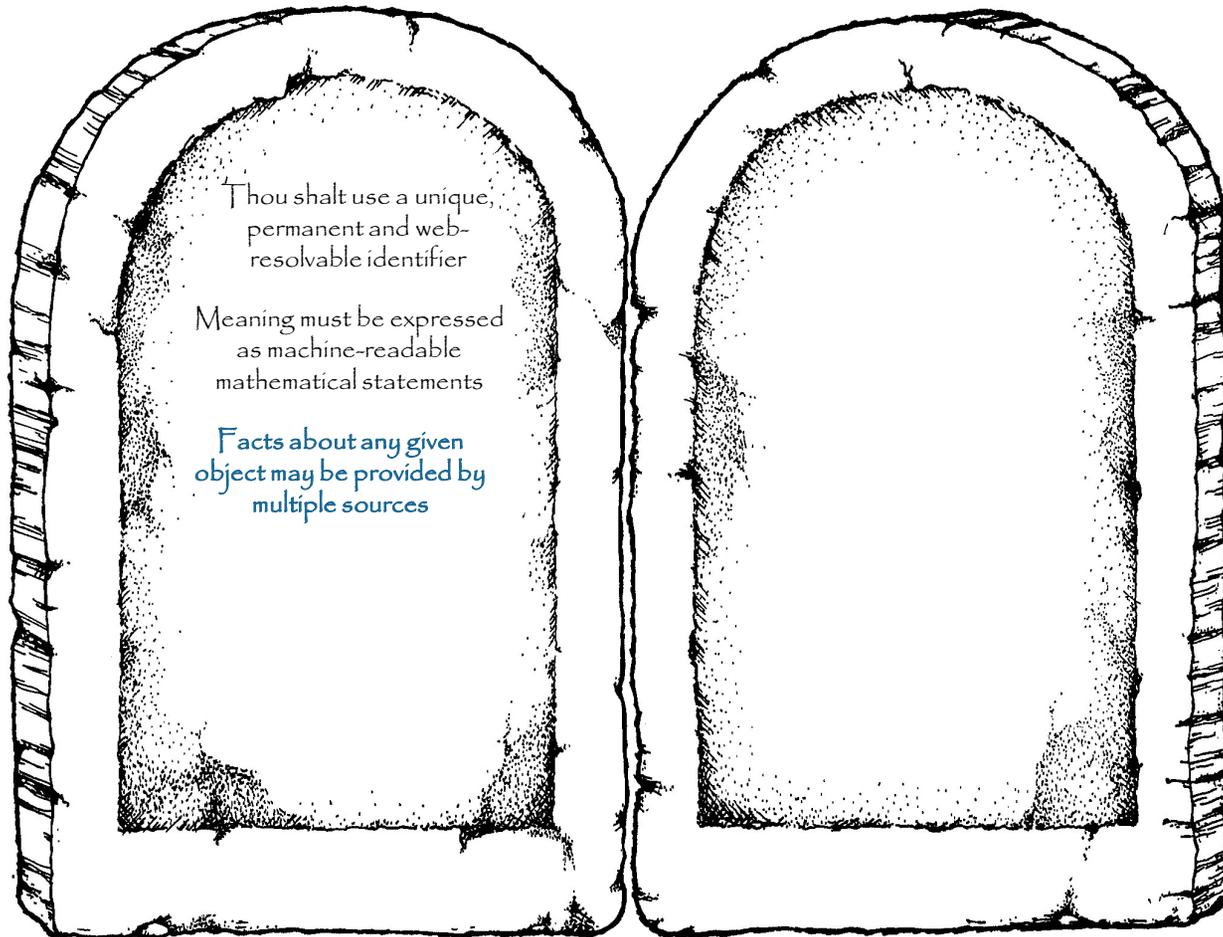
GUIDING PRINCIPLES OF EKG



Principle 2: Meaning

*The meaning of every
data point must be
directly resolvable to
a machine-readable
mathematical
definition.*

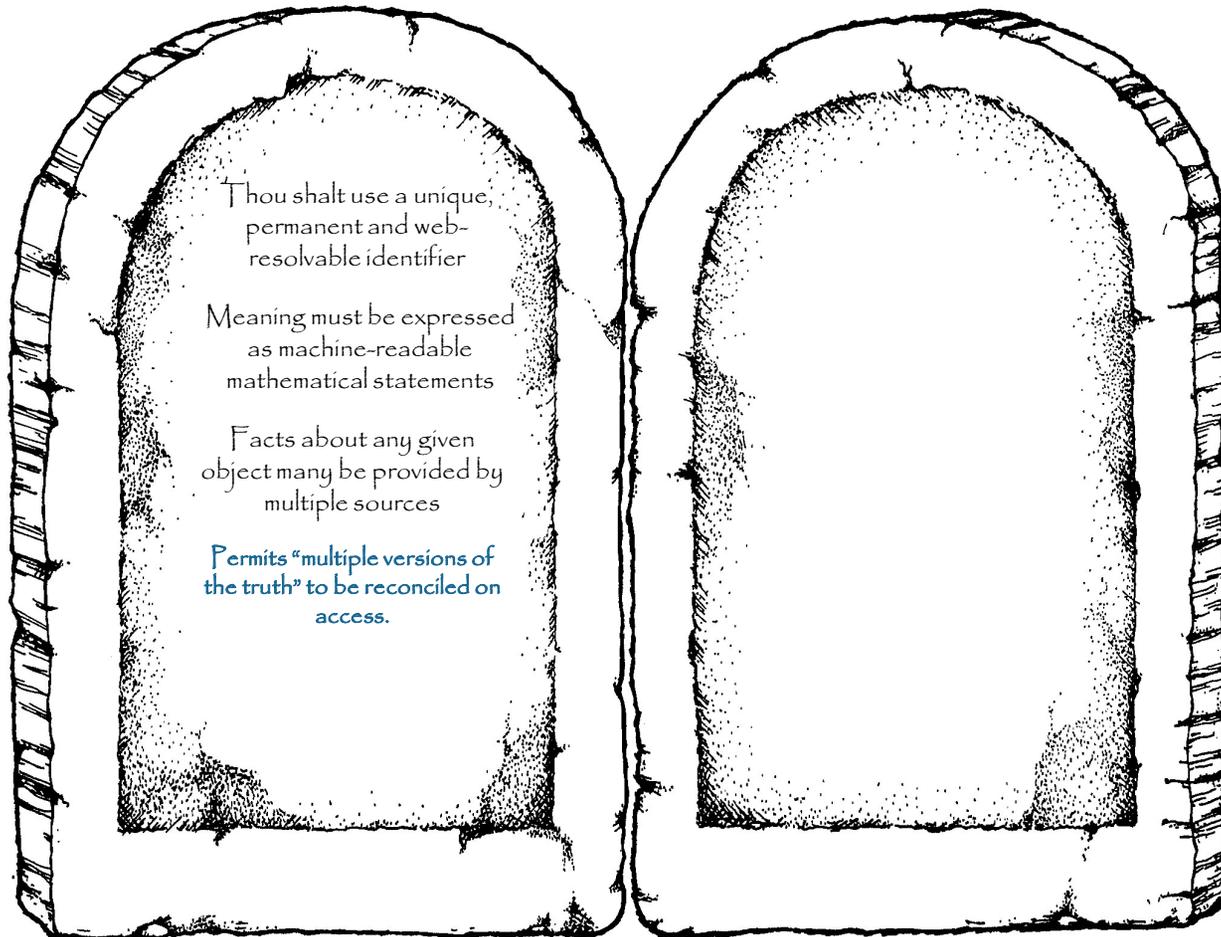
GUIDING PRINCIPLES OF EKG



Principle 3: Distributed

An EKG can incorporate any number of datasets of any number of data sources that can be hosted by any number of independent EKG Platforms. Any given access point provides connectivity to the EKG regardless of where it resides.

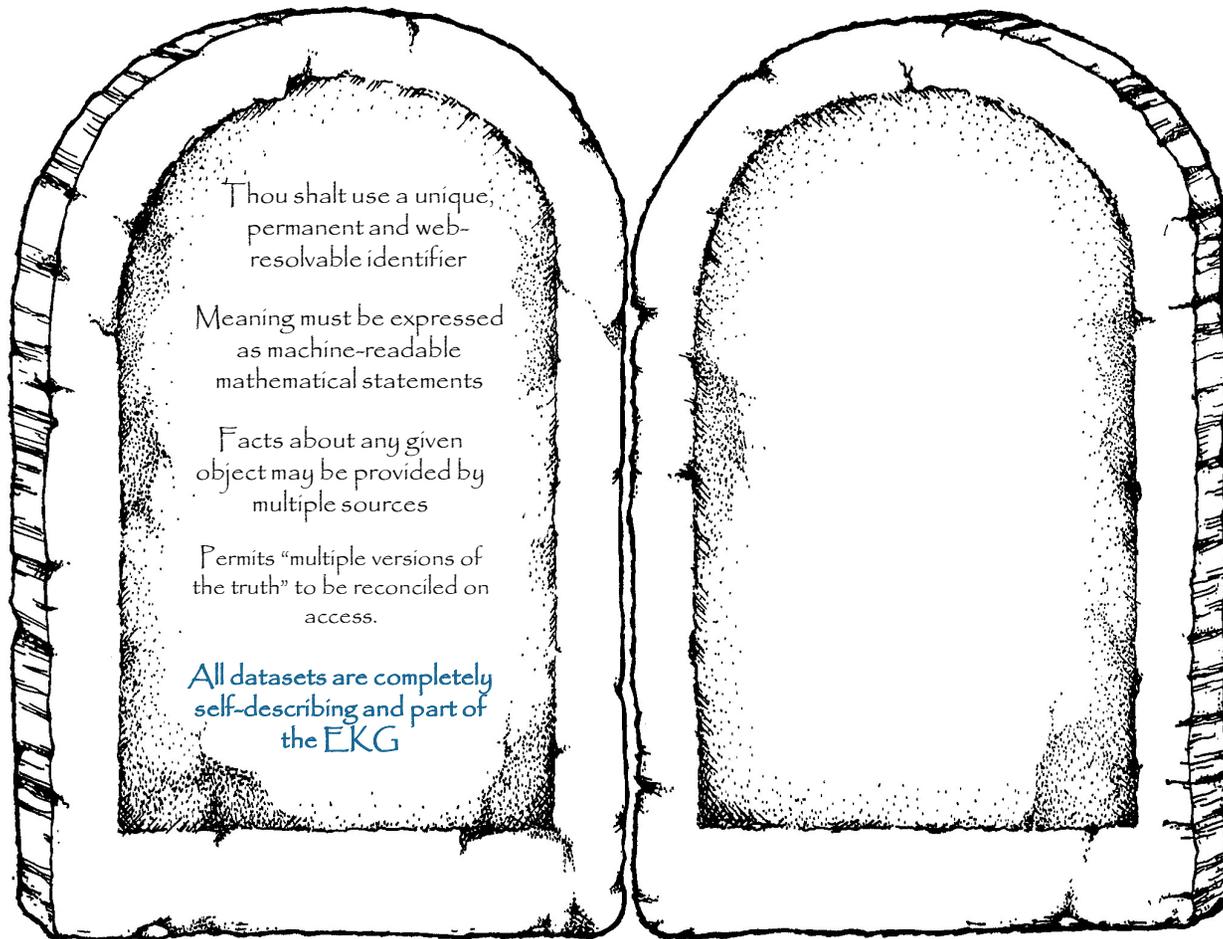
GUIDING PRINCIPLES OF EKG



Principle 4: Open World

Information can vary over time, come from many internal and external sources and be based on different identities and models. These “multiple versions of the truth” need to be reconciled on access by context.

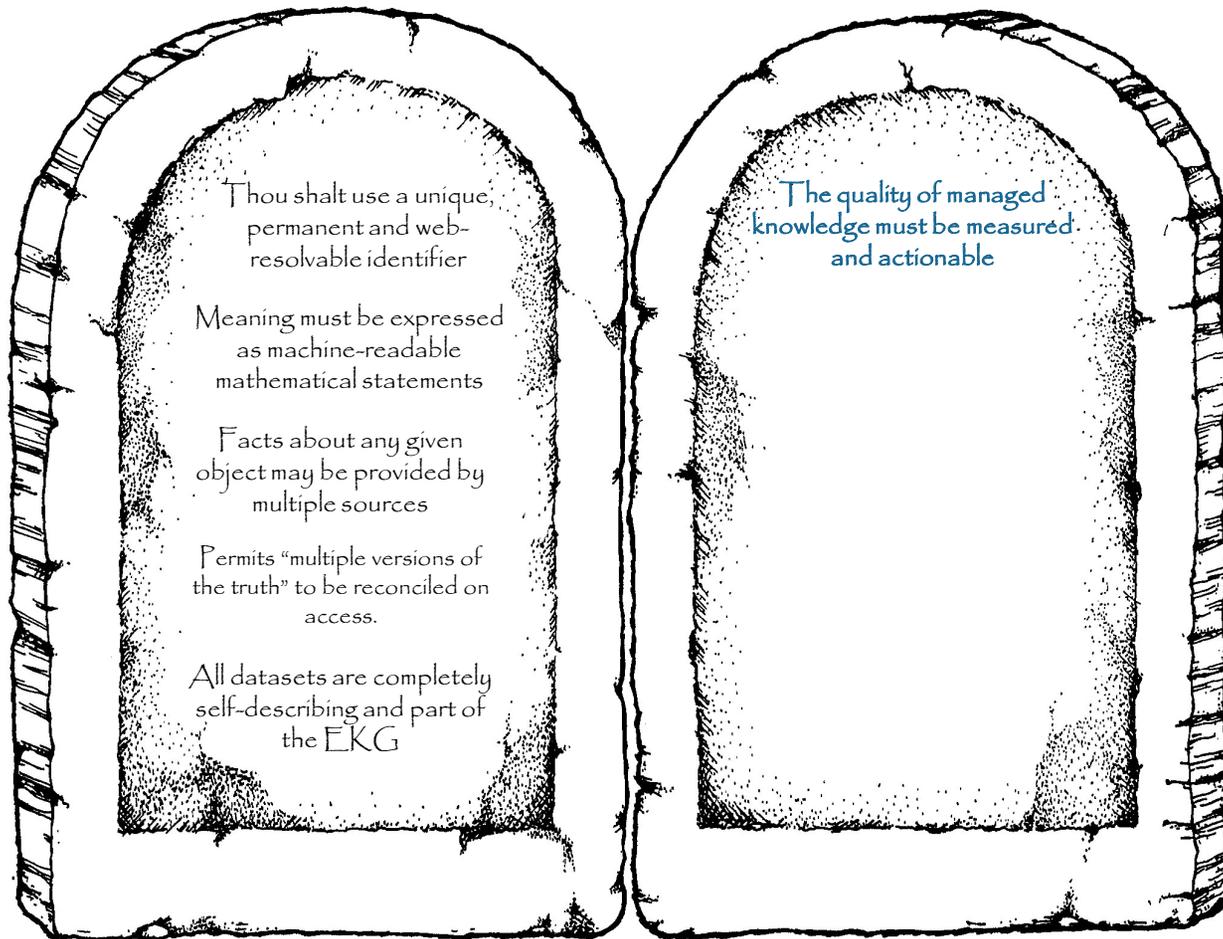
GUIDING PRINCIPLES OF EKG



Principle 5: Self-Describing

An EKG is composed of a set of self-describing datasets that provide information about lineage, provenance, pedigree, maturity, quality and governance. The EKG Platform has policy enforcement services that are driven by these self-describing datasets.

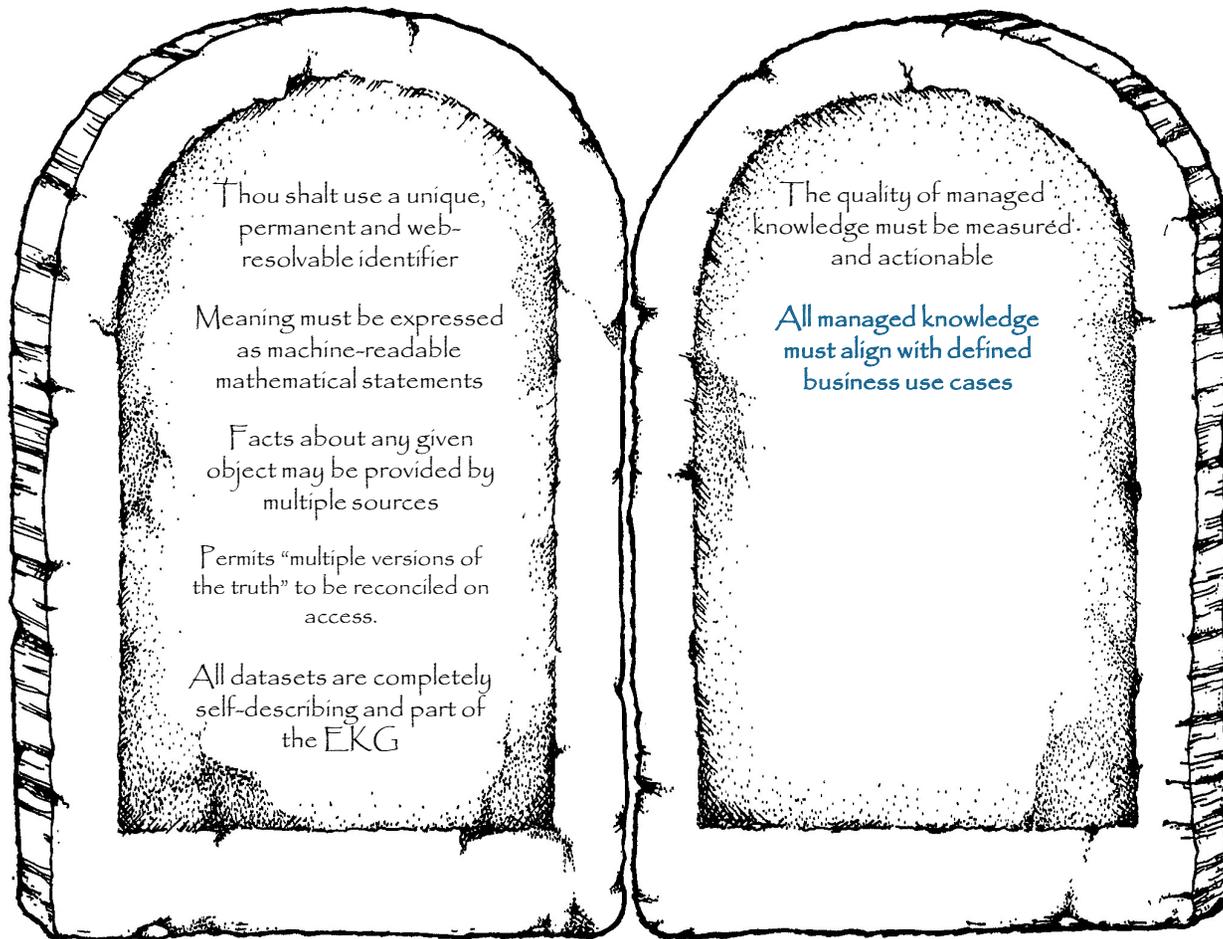
GUIDING PRINCIPLES OF EKG



Principle 6: Measurement

The quality and characteristics of the managed knowledge must be measurable and measured. Measurement criteria are used to designate fitness-for-defined-purpose and must be actionable.

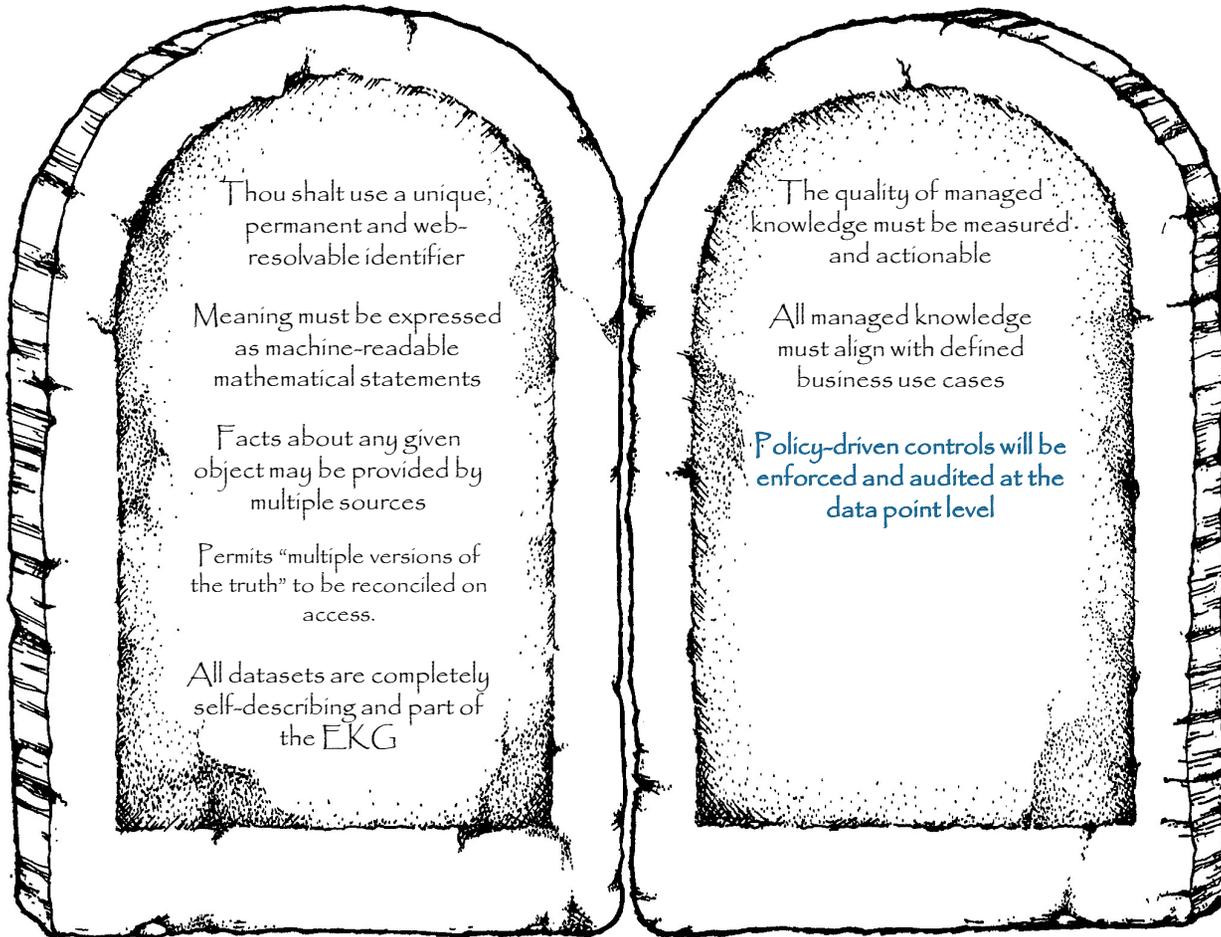
GUIDING PRINCIPLES OF EKG



Principle 7: Use Cases

All artifacts around and information in the EKG are directly linked to defined and prioritized use cases. Nothing in the EKG exists without a known business justification and purpose.

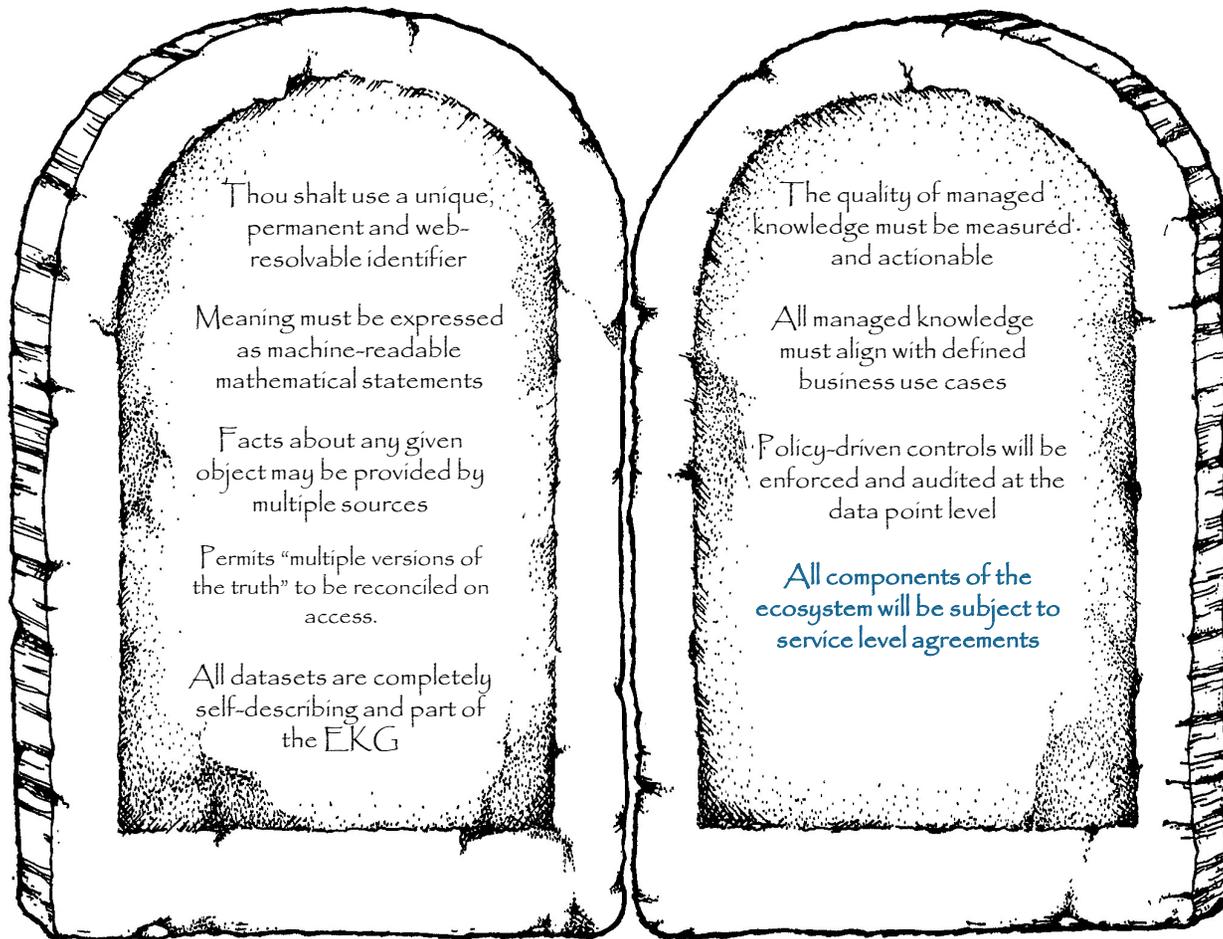
GUIDING PRINCIPLES OF EKG



Principle 8: Control

Entitlement, privacy and business policies will be modeled in the EKG and automatically executed, enforced and audited at the data point level.

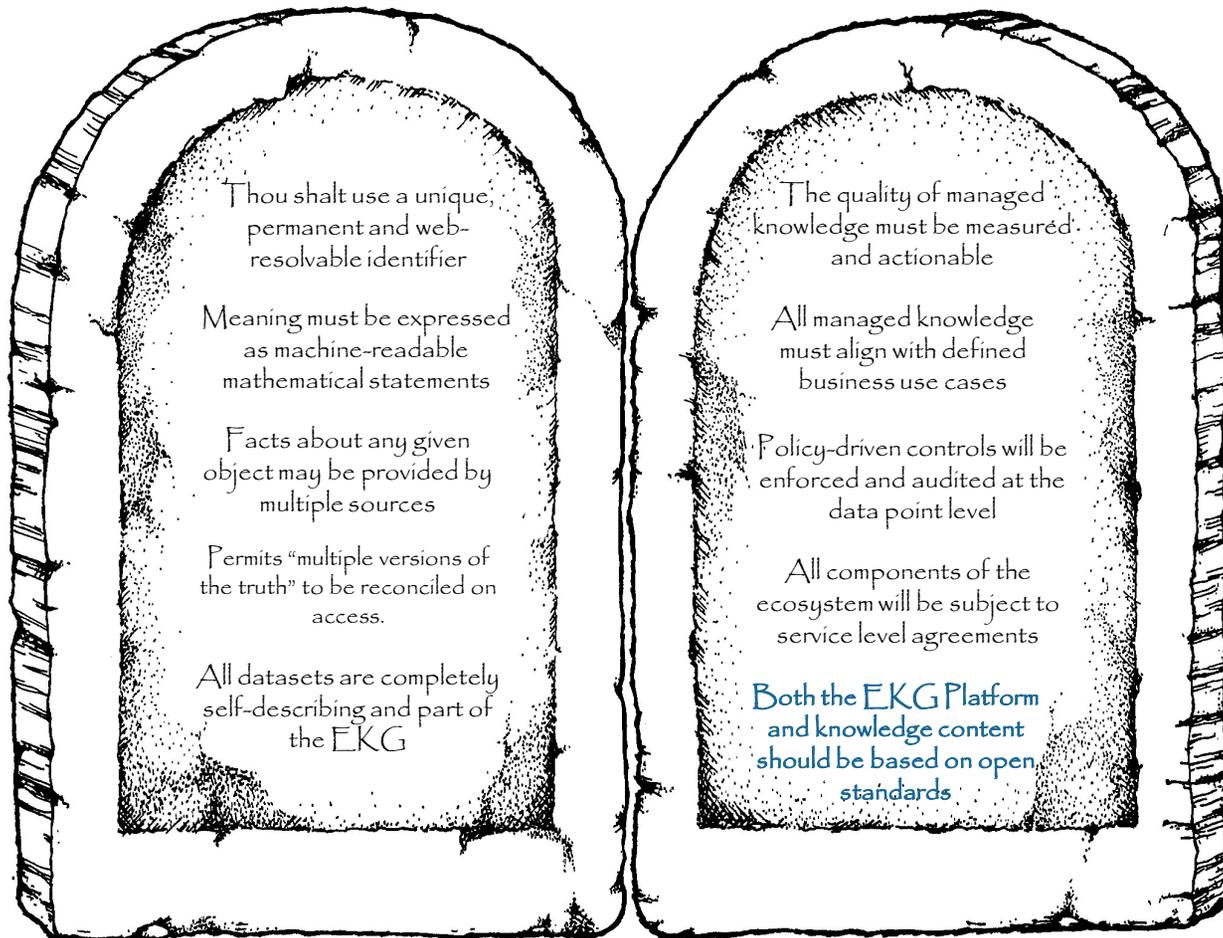
GUIDING PRINCIPLES OF EKG



Principle 9: Ecosystem

An enterprise will use a heterogenous set of technologies and data sources which will be incorporated into the EKG over time. All components of the ecosystem are subject to service level agreements.

GUIDING PRINCIPLES OF EKG



Principle 10: Standards

Both the EKG Platform and knowledge content should be based on open standards. Where necessary, the EKG Foundation will work to expand those standards.

Enterprise Knowledge Graph Maturity Model

Standard Measurement Criteria for Sustainable EKG Operations

EKG MATURITY MODEL BACKGROUND

- 1. Established Process** (Information Industry Map/definition; DMM with Carnegie Mellon/process; DCAM with EDM Council/governance; EKG/MM incorporating practical lessons and EKG measurement)
- 2. Measurement Consistency** (definition of components and capabilities; leverage existing frameworks; establish consistent lexicon; introduce standard scoring criteria; obtain insight rather than as a scorecard)
- 3. Synchronized Content** (based on core EKG principles; business friendly definitions of components; provides clear rationale for adoption; starts with core data management requirements; defines the structure of an ideal implementation)
- 4. EKG/MM Usage** (results in dependency maps and supports prioritization decisions; perform gap analysis; leads to operational roadmap; conduct industry benchmarking; easily translatable into Board/executive management metrics)



EKG MATURITY EXPECTATION

2019 2020 2021 2022 2023 2024 2025 2026 2027+

Initiate Journey

1

Implement foundational EKG components for minimum viable product. Establish the baseline architecture using isolated ontologies. Build initial team and processes. Project-based funding and specialist teams.

Demonstrate Capabilities

Extensible Platform

2

Design ontologies for related (multiple) use cases. Create reusable architecture based on expanded design principles. Line of business alignment and funding. Center of Excellence and DataOps environment initiated.

Parallel KG Activities

Enterprise Ready

3

Scalable and resilient EKG platform for mission-critical applications. Datapoint level authentication. Connected inventory linked to governance. Dedicated resources for design, build and support.

EKG as Default Data Hub

Strategic Asset

4

EKG as strategic infrastructure for the organization and authoritative source for all glossaries. Supports application consolidation and cost-reduction. Leverages artificial intelligence and process automation.

Operational EKG Utility

Operational Ecosystem

5

EKG is central to systems and business processes. Full integration into core business operations (including partners). Elimination of data silos. IaC architecture implemented. Inference and reasoning for advanced AI.

Continuous Development

EKG Maturity Model Coverage

Business

Strategy

Data Centric Orientation

Cognition and prioritization of the problems of data fragmentation by line of business management

Business Goals

Prioritization of strategic use cases and clarification financial models for validating the EKG business case

Innovation Maturity

Support for prototyping, experimentation and process for facilitating innovation

Executive Leadership

EKG engagement, commitment and prioritization by CIO, CTO and line-of-business stakeholders

Data

Strategy

Data Architecture

Models, standards and rules for how data is identified, described and integrated into the EKG

Quality Assurance

Overall approach to ensure data is fit for its intended purpose across linked use cases and applications

Entitlement Control

Knowledge graph-based (KG/BAC) access control mechanisms implemented at the datapoint level

Data Governance

Policies, processes and accountability structures to ensure management of the data production/consumption process

Technology

Strategy

Technology Environment

Principles, patterns, processes and tools for storage, integration and deployment of data

DataOps Process

Rules and standards for building, testing, deploying and maintaining the enterprise knowledge graph

End-User Tools

The ability of end users to navigate, search, link and extract value from the knowledge graph



Organization

Strategy

Organizational Culture

Mechanism for addressing internal politics, reducing workplace drama and solving the “culture of competition” across the enterprise

Team Capabilities

Key proficiencies required to design, build, operate and manage the EKG product suite

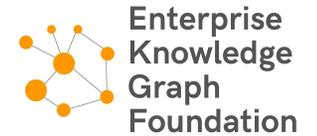
Resourcing

Organizational, financial and human resource plans for implementation of the enterprise knowledge graph

Talent Management

Ability to attract, empower and retain people with the right skills to support the EKG environment

EKG MATURITY ACHIEVEMENT MATRIX



	BUSINESS	DATA	TECHNOLOGY	ORGANIZATION
Level 5 ECOSYSTEM	Fully integrated business (everything is connected)	Elimination of data silos across enterprise	Model-driven architecture broadly implemented	EKG central to systems and business processes
Level 4 STRATEGIC	Applications consolidation and cost-reduction	EKG is serving all core data	EKG is recognized as strategic infrastructure	Multi-year resources fully allocated
Level 3 ENTERPRISE	Business objectives synchronized	Globally unique IRIs for critical data concepts	Scalable EKG platform for mission-critical apps	Technology roadmap and metrics aligned
Level 2 EXTENSIBLE	Strategic business objective	Critical data elements linked to ontology	Cloud architecture for limited use cases	EKG prioritized by executive leadership
Level 1 INITIATE	Tactical alignment of goals with IT	Inventory and systems-of-record curated	Legacy environment limits flexibility	Limited understanding of EKG capability

EKG/MM STRUCTURE

A. Data Architecture

A.1: Data Inventory Management

Description: A data inventory is a repository of data assets for the organization (also known as a data catalogue or metadata repository). The inventory contains information about what data exists, where it resides (“data-at-rest”), responsible parties, upstream/downstream usage, how it moves (“data in motion”) classifications, quality designations, availability and other useful metrics. The contents are catalogued at physical, logical and business levels using defined organizational standards. Consistently documenting data is the backbone of an effective data management program. Policies and controls are required to ensure the inventory remains accurate and relevant.

EKG Rationale: Linking data inventory to the knowledge graph (and business concepts) ensures precision of meaning at the most granular level. Data in the knowledge graph is traceable to all application usage allowing users to find data of interest through assisted and contextual search. The ability of the knowledge graph to connect data with metadata enables users to perform flexible queries in ways that were not previously possible.

Core Questions

- **A.1.1:** Do (one of more) data inventories exist
- **A.1.2:** Is the inventory based on defined standards (for both meaning and format)
- **A.1.3:** Is defined and in-scope data (both breadth and depth) covered in the inventories
- **A.1.4:** Is the data inventory linked to systems of record and authorized data distribution points
- **A.1.5:** Is the inventory linked to the business meaning of the data and expressed using standards
- **A.1.6:** Is the creation and maintenance of the data inventory mandated by policy and incorporated into the data strategy
- **A.1.7:** Is the quality of the content in the inventory measured, reported to involved stakeholders and used for process enhancement

EKG Maturity Scoring (concepts are inherited as the levels progress)

Level 1: Initiate

- Sources, data sets and metadata are onboarded and expressed as formal ontologies
- Authoritative (upstream) data sources and (downstream) consumers are documented and verified by users, data and technology
- The inventory of applications is defined and selected for graph applications
- Requirements and dependencies for each outbound data flow are documented and verified (implementation in the graph is not a requirement)
- Business glossaries for in-scope use cases are defined and verified in the graph (including a list of data sources and datasets)
- Policy implemented mandating inventory maintenance and only authorizing the use of data that has been logged into the inventory
- Access control and entitlement is implemented for in-scope data and related use cases (at the named graph level)

Level 2: Extensible

- Use case trees (taxonomies and dependencies) are defined, standardized and implemented
- All upstream data sources are linked to the authorized systems of record and distribution points
- Policy mandating the use of systems of record and documentation of data flow is implemented
- Entitlements have been defined in the graph (governing access to sources of data in the inventory)
- Classifications (i.e. criticality, security, privacy) are aligned with the use case tree and captured in the knowledge graph
- Governance requirements (i.e. use cases, accountability, data sources, data flows, SLAs) are modeled and registered into the knowledge graph

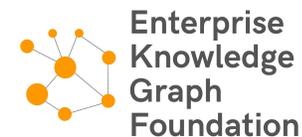
Level 3: Enterprise

- Data inventory is centralized in the graph and linked to governance for defined use cases
- Ontologies and data models (including change history and transformations) are registered in the knowledge graph
- Entitlements are calculated within the inventory and enforced at the datapoint level
- Data quality is automatically calculated (fine-grained with dynamic value resolution) within the inventory for each use case
- Data retention rules are registered in the graph and automatically enforced
- Full audit trail for all upstream and downstream data usage is registered in the graph
- Data elements, calculation methods and CDEs are linked to individual regulatory requirements

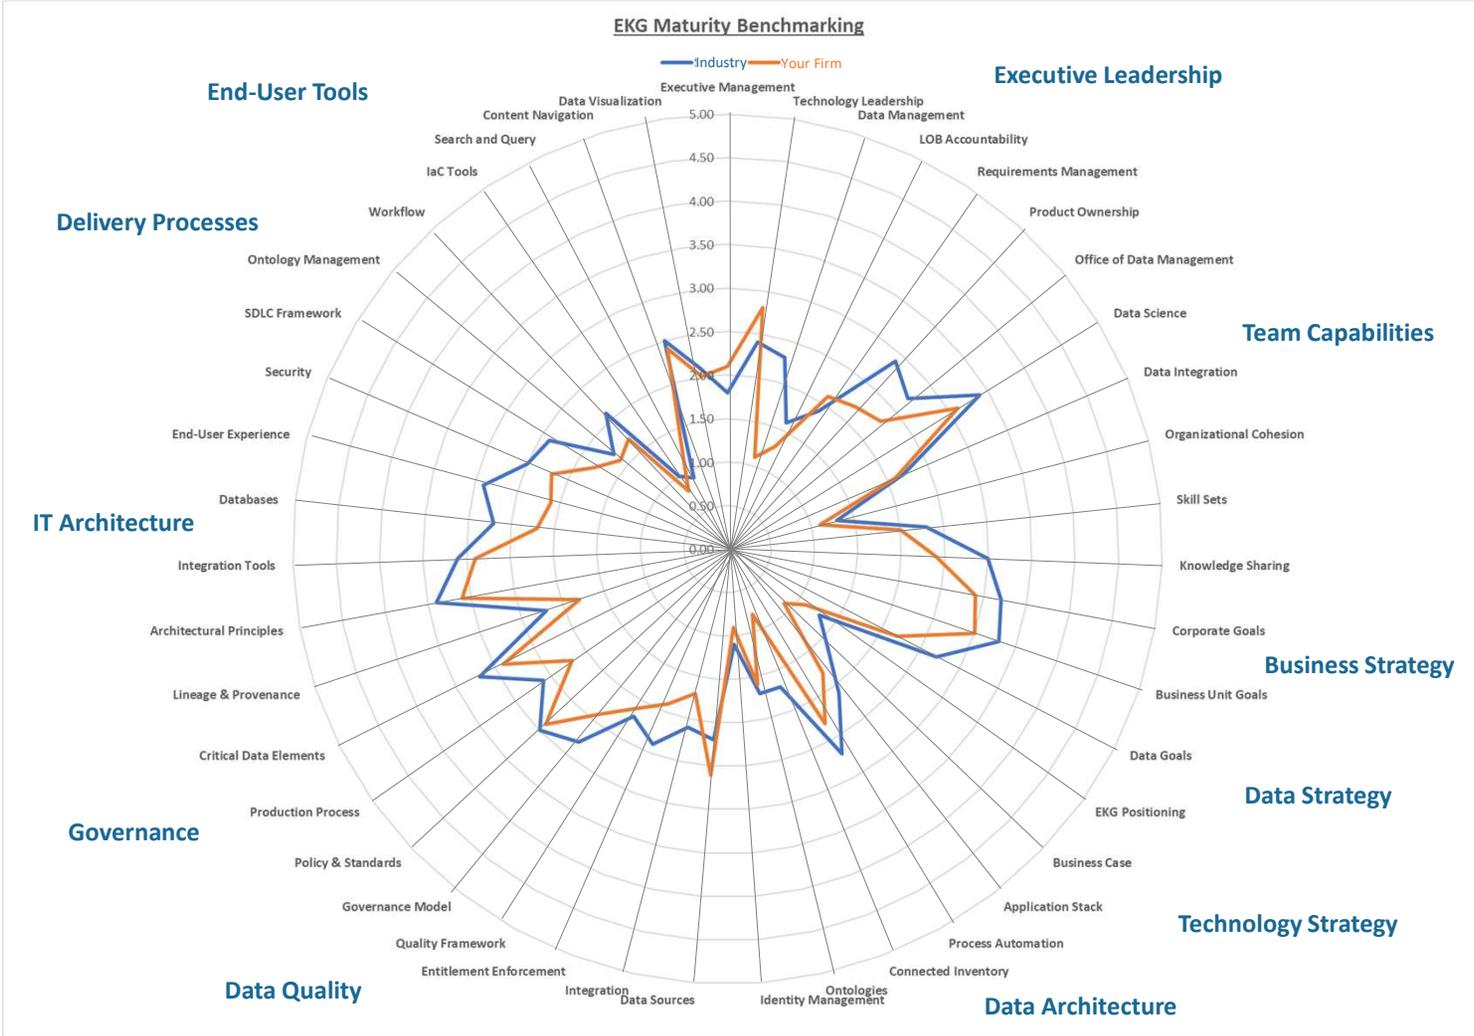
Level 4: Strategic

- Connected inventory has been extended to include real-time (transactional) data
- Inventory is extended to external suppliers and third parties along the supply chain
- The inventory is fully integrated with machine learning to optimize data flow
- The “value of data” is calculated and classified within the organizational inventory

EKG/MM GOVERNANCE (IMPLICATIONS)



Inventory	<ul style="list-style-type: none"> Catalogue of assets created Inventory mapped to consuming applications Data registration processes established 	Quality criteria calculated by EKG, inventory (data sets and sources) onboarded to EKG, ontologies in EKG, field mapping to ontologies, use case	Glossaries	<ul style="list-style-type: none"> Inventory of glossaries with verified business meaning Captured as metadata Mapped to data flow and transformation processes (data reconciliation) Governed by policy (meaning and alignment) 	Identifiers	<ul style="list-style-type: none"> Linked to inventory Mechanism for resolution (cross-referencing) 	IRI principles and policy implemented, host and domain names implemented, onboarded data tagged with IRIs and mapped, SameAs process for resolution	Lineage	<ul style="list-style-type: none"> Lineage and transformation traced to authoritative sources and linked to consuming applications Business provenance defined and linked to use cases and glossaries Governed by policy 	Lineage onboarded to EKG, linked to ontologies, onboarded to EKG, downstream customers automatically tracked <ul style="list-style-type: none"> All aspects of lineage (business flow, logical to physical) need to be identified and converted into the EKG Must track the flow of information into the EKG <ol style="list-style-type: none"> EKG Lineage in EKG itself (native) – lineage will be <u>automatically</u> calculated and tracked for each data set (standard feature of EKG) Legacy Lineage (conversion from legacy) – onboard the metadata from legacy into the EKG (convert to lineage ontology)
Integration & Mapping	<ul style="list-style-type: none"> Cross-referencing and data integration processes implemented Linked to inventory (glossaries and models) 	Data integration standards and business rules onboarded to EKG <ul style="list-style-type: none"> Change Approach: Think about this differently. No longer try to create one dataset from multiple sources. Instead establish one self-describing dataset per data source. IRI links: Cross-reference via IRIs (linked via the identifiers) in the DataOps pipeline (things, not strings). Meaning: Extreme normalization (everything is broken down into its most atomic components) – as a triple. Each dataset represents the opinion of the source but refers to any other object 	Named graph, machine-executable rules and application logic (KG-based access control), EKG linked to organizational management use case, EKG as authoritative source of entitlements, governance	Policy & Standards	Glossaries are onboarded to the EKG (relevant terms in RDF), linked to ontologies, searchable via corporate intranet, EKG is authoritative source <ul style="list-style-type: none"> Glossaries need to be onboarded/formalized inside the EKG (from tools like Colibra to EKG) – need a standard for representing glossaries ANALYSIS - every glossary has a context (lots of glossaries (dozens/hundreds) linked to applications (systems, products, objectives) and is contextual (regulatory, marketing, transaction) – in EKG all glossaries are in the graph – need to define the scope of the glossary meaning of the TERM for each EKG) ONBOARDING/MAPPING - EKG be onboarded (pipeline) – per taxonomy, identify information about governance 	DataOps	Automated testing and validation of ontologies, DataOps environments for ontology lifecycle management implemented	<ul style="list-style-type: none"> Git-Based: Based on Git (as the key governance principles) – thou should NEVER have anything in production that is not managed via the Git repositories. Anything that has an impact on production must go through GIT – auto-log of changes, audit component is always operational. Every change is transparent and recorded Approved: Every change must be approved by the right people; changes must be logged into JIRA (describing reason for the change) Change Management: Changes must be planned and justified (authorized and approved); Changes must be tested before production (automated testing) Mapping files must also be expressed as an ontology and tested. Mappings and ontologies are stored and maintained as a FILE in GIT (key principle) Testing: rigorous and non-human testing – every test must be automated for consistency and precision. No longer rely on humans for testing. No more “big-bang” deployment. Changes are tested as they are introduced into the process (not as a big UAT job) – no “big-banged” deployment. No human intervention. Testing occurs before approval (all changes are validated before entering into production) 		
Entitlements	<ul style="list-style-type: none"> Access and authorization rules defined Linked to SOR, data flow and transformations Provisioning points established Administration controlled Governed by policy 	Assurance that all entitlement policies are implemented (combine all policies, rules, restrictions) <ul style="list-style-type: none"> Stringent entitlement is required for onboarding confidential/PII data; the same security as source system must have assurance of compliance and other privacy regulation Linked to entitlements = vendor contracts and payment requirements 	Use cases defined with stakeholders <ul style="list-style-type: none"> TCO and measurement criteria established Sequencing and prioritization determined Budget defined and aligned with business processes Funding model and LOB budget oversight established Metrics defined for cost and value 	Policies complete, approved and aligned with strategy <ul style="list-style-type: none"> Synchronized with SDLC and business process Linked to operational control functions Mandated and audited 	Core policies (identity (IRI/namespaces), meaning/ontology (RDF/OWL), federation (leave data in place), service layer (quality, access control, queries), self-describing data (linkages), standard-based (RDF, OWL, SKOS, SHACL, Linked Data) <ul style="list-style-type: none"> New (or adjusted) policies that need to be crafted to deal with EKG related processes (inventory registration, glossaries, authoritative sources, DataOps (testing), business rules), metadata, data modeling, lineage/provenance, storage, integration, agile, reuse) 	Business Case	Introduce the concept of foundational dependencies associated with use cases (the concept of reuse) – need to figure out how to pitch KG to business in a way they understand (and buy-into) <ul style="list-style-type: none"> Change the sequencing and prioritization – build foundational ontologies that can be leveraged across business objectives Simplify and better link the business case to the value of data management The KG will generate more accurate information (after populated) about the TCO and value 	NOTE: course-grain that can't easily be components. Regu fine-grained data is different ways (elim producing aggregate model view of the)		





FOUNDRY

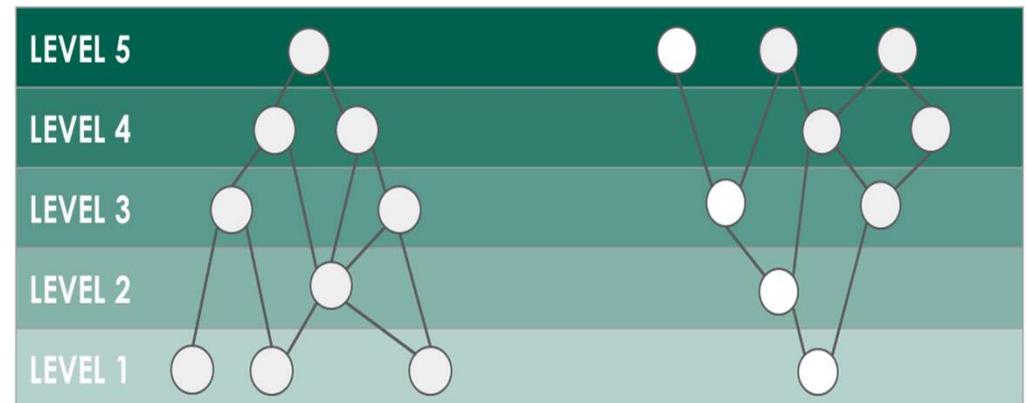
for the EKG Foundation

Note: the Foundation is starting with four initial portals. Others are under consideration and will be created based on member priorities.

The Foundation serves as a master repository of artifacts for people and organizations. The associated portals have been designed to offer curated access to published ontologies, tools, use cases and other best practice recommendations needed for sustainable knowledge graph operations.

USE CASE PORTAL

- Repository of use cases, use case tree structures and test results to help accelerate the adoption of enterprise knowledge graph technology
- Use cases are created to ensure appropriate expectation management with ontology consumers. They are designed to be devoid of any technology assumptions.
- Use cases are structured as a “contract” to create a testable relationship between user requirements and the ontology structure.
- Strategic use cases can be sub-divided into modular (reusable) components to create a “use case tree” that captures the functional requirements, user personas and dependencies for each strategic use case.



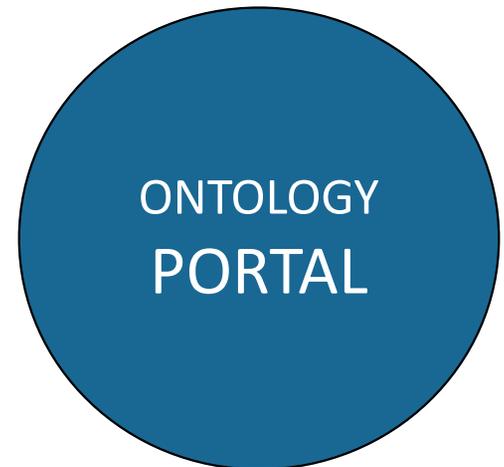
BEST PRACTICE PORTAL

- Establishment of a structured methodology to ensure the performance of the knowledge graph in integrated environments (clarify dependencies, understand limitations and define the scope of work)
- Contributions should comport to the structured method based on Agile principles and the implementation of a DataOps environment for hygiene and policy testing
- Ontologies must be aligned with verified use cases/user stories and standardized governance control processes linked to the EKG maturity model

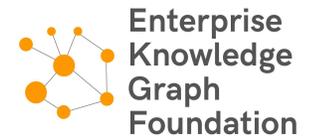


ONTOLOGY PORTAL

- Access published ontologies covering a broad variety of subjects linked to defined use cases, user personas and test data
- Published ontologies include a detailed description of the classes and concepts as well as information on methods (i.e. Git URL) for user access
- Assurance of “rights to publish” as well as intended update frequencies before are required
- Ontologies listed in the portal will be tested for both logic and hygiene



TOOLS PORTAL



- The Portal is designed to provide a factual evaluation of capabilities and product functionality (*based on quality criteria/maturity model capabilities*) in support of EKG user requirements.
- The Tools Portal is designed to cover taxonomy management, ontology editing, ETL, triple stores, reasoners, knowledge-based ML/NLP products, document management and semantic search.
- The Tools Portal contains a catalog of products, tools and open source utilities for the enterprise knowledge graph - as well as the results of Foundation testing processes.



NEXT STEPS: STAGE ONE

1. Organizational – formal launch of EKG Foundation

- Publish Terms of Reference and Strategic Plan (sponsor/member recruitment and operational infrastructure)
- Finalize EKG bylaws and governance structure (operating model)
- Recruit key participants to refine and prioritize the agenda

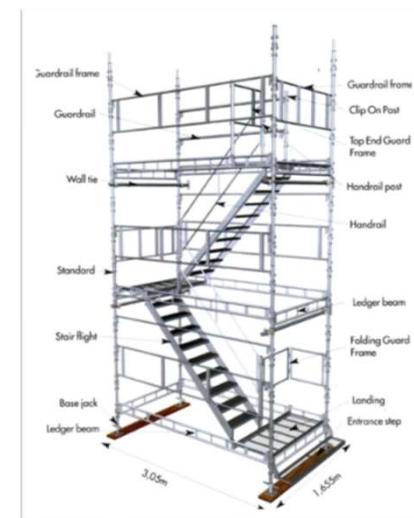
2. Working Content – working group on EKG maturity model

- Populate repository with EKG/MM working documents
- Working meetings to review content and agree on scoring criteria
- Implement initial benchmarking survey (baseline)

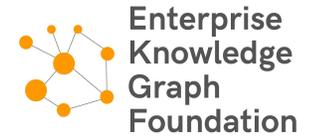
3. Ontology Portal – fork Ontology Publisher with enhancements

- Processing service for quality evaluation, formatting and enrichment
- Mechanism for ontology submission via website and REST API (supporting any RDF syntax)
- Publish results referencing element, datetime, check applied, resultClass, resultText, severity

4. Communications – launch website, LinkedIn Group and EKG Directory



Join the EKG Foundation



A screenshot of the EKG Foundation website. The top navigation bar includes the EKG Foundation logo, the URL "EKGf.org", and links for "Home", "Mission", "Organizing Team", and "Contact/Join us". The main content area features the title "ENTERPRISE KNOWLEDGE GRAPH FOUNDATION" and a mission statement: "Promoting semantic standards and best practices to support the growing enterprise graph marketplace". A quote by Dennis E. Wisnosky is displayed: "The development of a collaborative environment is one of the principal operational goals of the EKG Foundation". Below this, three sections are listed: "MISSION" (To be the focal point for coordination and engagement across the knowledge graph community), "CONTENT" (The Foundation will promote, share, and advocate for standards, best practices, products, services and open source components), and "COMMUNITY" (A collaborative community and marketplace for all those involved with using and creating knowledge graphs). The contact email "info@ekgf.org" is provided. On the right side, a "CONTACT US" form is visible, containing fields for email, first and last name, organization, job title, country/region, phone number, and mobile phone number, along with checkboxes for contribution interests and a "Submit" button.

www.ekgf.org