Journal of Innovation
February 2025 | 26th Edition

Object Management Group

Applying Standards to Information Centric Operations

Author:

Mike Bennett
Object Management Group
bennett@omg.org

CONTENTS

FIGURES

We live in a world of information. From the moment we are born we are surrounded by shapes and sounds and color, and we spend the next few years figuring out what it all means. In our adult working lives, almost all of us live in information some of the time, and some of us all the time – whether it's profit and loss figures, sales pipelines, engineering design specifications, software or whatever. Information is our natural habitat.

And yet, to interact with this sea of information in which we live, we must constantly hop from one computer application to another. Or in some cases do all our work within one application and use extra features they introduced to let us handle diagrams in a word processor, paragraph layouts in a spreadsheet and so on. The idea that we could spend our working lives within a single integrated information space ("Cyberspace," to use William Gibson's term for it in the novel Neuromancer[1]) has continued to elude us, despite the writings of visionaries like Gibson and others like Douglas Adams[2] and Vannevar Bush as far back as 1945[1].

Some of the pre-conditions for information-centric working are well into the early adopter stage of the technology evolution cycle. For example, the introduction of semantic models, along with the "Data Centric" Manifesto[3] are starting to provide the thinking, and more importantly the tools, to start the journey from here to there. A more complete inversion of applications and data does not require new technology but a different way of using the technology we already have; a discontinuity in the technology adoption curve.

To complete that journey is going to require a combination of imagination, vision and standards.

This article explores the case for standards to support data-centricity, both for current uses in data centric operations and digital twins, and for future uses that may only become clear under a more ubiquitously information-centric way of working.

1 DATA CENTRICITY

The Data Centric revolution in computing [2] brings promises and challenges, separating content, semantics and presentation. This empowers data to be treated as a resource in its own right, with its own lifecycle, quality measures and ownership independent of the applications which use it and potentially outliving them. Use cases include increased flexibility, risk and regulatory compliance and the provision of datasets for training generative AI solutions. As a simple by-product, firms providing semantically enabled data will also find a faster and smoother route towards FAIR (Findable, Accessible, Interoperable, and Reusable) data operations [3].

---

[1] Gibson, W. (1984). Neuromancer. *Ace*. ISBN: 0-441-56956-0

[2] Adams, D. (1990). Hyperland (video in 5 parts). Available: *http://www.youtube.com/v/rOsPKjbMvxY&hl=en&fs=1&*

[3] The Data Centric Manifesto. Available at *http://datacentricmanifesto.org/*

1.1 THE DATA-CENTRIC MANIFESTO

The Data Centric Manifesto states (emphasis in original):

quote

> “We have uncovered a root cause of the messy state of Information Architecture in large institutions and on the web today. It is the prevailing application-centric mindset that gives applications priority over data. The remedy is to flip this on its head. **Data is the center of the universe; applications are ephemeral**.

The key principles of this are applicable today in any large organization. These principles set out a way of working that can be achieved today by using semantic technologies, and will drive efficiencies in integration, application development and usability of the suite of applications used in large firms such as financial services providers.

The above quoted paragraph states: “The remedy is to flip this on its head.”

However, let’s take a closer look at this. If we were to really flip this on its head, we would not simply be installing webs of linked enterprise data across a firm, we would be changing the way people interact with the tasks of word processing, presentation, calculation and the rest.

That’s not coming any time soon.

In this article we will consider both the current state of data-centric or information-centric working, with some examples and case studies, and take a brief glimpse at what the next steps might be, starting with recent developments in the digital twins area. The provision of data in a common format with common semantics is a pre-requisite for any “flipping on its head” that may happen later, and the benefits of this first phase can be obtained now. We will also look at how standards can contribute, both to the current innovations that are driving this first wave of data-centricity, and in what may come further down the line.

To start then: data that is shared across the organization and among applications needs to be in a common format and have some common ways of expressing the intended meanings of each piece of data. These are not the same thing.

For common data formats, we would look to the linked data or enterprise knowledge graph formats of RDF[4] and RDF Schema[5], along with other graph data formats. These are easy enough to find, but first you need to get the semantics right. That is, the meaning of the data. To be clear, we are not using the word “semantics” here in any special sense other than to talk about what things mean. The technology for this is often referred to as “Semantic Technology” and at the heart of this is a kind of artifact, or more correctly a class of artifacts, called “Ontology.”

[4] RDF 1.1 Concepts and Abstract Syntax. Available: *https://www.w3.org/TR/rdf11-concepts/*

[5] RDF Schema 1.1. Available: *https://www.w3.org/TR/rdf-schema/*

2 ONTOLOGY

The original definition of Ontology is the study of everything; of what is. This has been around since the Ancient Greeks, but now we use the word in a subtly new way, to talk about the ontologies (plural) that our data represent. Each database, each messaging schema, even the variables in computer programs have some intended relationship with reality. They represent something. Whether they represent things completely or consistently, whether the data means the same thing going in as someone else thinks it means coming out the other end, is another matter. Everything in your computational environment has an ontology, whether it knows it or not; the smart next move is to formalize those ontologies so that what people think a given piece of data means is not left to chance or to local interpretations but is set out logically.

A formal artifact that can be used to represent the meanings of real things, or the meanings that data is referring to, is also called an “ontology.” Ontologies in this sense are models that use formal logic to make consistent, computationally verifiable assertions about things in the world.

An ontology uses logic to make assertions about the world in terms that can be understood in the same unambiguous way by both machines and people. But logic is only part of the story. Logic on its own doesn't make a model meaningful since you can have logically complete and consistent assertions that are utter nonsense, with no bearing in reality. For this you need to define some formal relationship between the model and what it is a model of.

The simplest kind of ontology is represented using a family of logic called Description Logic (DL) [4], a kind of knowledge representation language. This is a simple logic of what there is and what distinguishes things from one another. More complex features of reality would require higher orders of logic to represent them, but we start with DL, which is reflected in the syntax of the Web Ontology Language (OWL) [5].

We need to assert meanings against data to enable a data-centric environment. However, if you start from the bottom up, asserting meanings one by one against pieces of data, there is no guarantee of completeness or consistency and, more importantly, no actual grounding in reality. A more mature engineering approach is to work from the top down: to determine the meanings of things that are of interest to the organization, and only then link this to data, creating the necessary sea of linked data as an enterprise knowledge graph.

Some examples and case studies will help to clarify this relationship between ontologies that reflect reality, and the use of semantic technology such as OWL ontologies. These distinctions will become particularly important in considering digital twins.

3 ONTOLOGY EXAMPLES AND CASE STUDIES

The implications of data centric operations are more far reaching than most of us yet realize. A true inversion of data and operations means a true inversion of how people interact with data and applications. This inversion may remain elusive, but there have been applications of elements of this vision. In this section we will look at some case studies in finance, manufacturing and pharmaceutical industries, where some of these techniques have been tried or used. In these we will see that many elements of data centricity have been already implemented or are starting to fall into place. These examples range from the application of formal semantics through enterprise knowledge graphs to industrial digital twins and data visualization.

We will look at lessons learned from these initiatives. Then we will consider what aspects of these initiatives may benefit from the development of new standards or the use of existing ones.

We start by looking at a core aspect of data centric working as outlined in the Data Centric Manifesto, namely the provision of formal ontologies.

3.1 Ontologies in Finance

In the Global Financial Crisis of 2008 a lot of financial firms were left with exposures to the institutions that went down. They were not short of data; the data was there, and the reality of their contractual exposures was there, yet in many cases it took several weeks to turn that data into the knowledge they needed.

The financial industry realized that it could no longer afford to have data maintained in disparate silos. Siloed data represented the potential of knowledge, but was not readily accessible to business stakeholders in forms they could digest and act upon when they needed it. Integration across data silos has long been recognized as a source of costly inefficiencies, and a number of data standards had been developed to provide some common language at a data level. These standards are well understood and widely adopted. However the feedback from business stakeholders was that what was missing in these data standards was a formal definition of what each data element really meant.

What was needed to supplement financial industry data and messaging standards was a business-facing, computationally independent model of the meanings of the concepts that are to be reflected in data. In response to this industry feedback, the Financial Industry Business Ontology (FIBO)[6] was developed by the Enterprise Data Management Council (EDM Council)[7] as a means to capture "unambiguous shared meaning" across the financial services industry [6].

3.1.1 FIBO Proofs of Concept

Once FIBO was developed and delivered to the industry in its initial draft format, several proofs of concepts were carried out. These included:

[6] *https://spec.edmcouncil.org/fibo/*
[7] *https://edmcouncil.org/*

- Derivatives Messaging[8] – headed up by Wells Fargo, this PoC showed how to map terms from the industry standard FpML messaging standard into the FIBO ontology, with automated classification of derivatives based on their properties.
- Regulation W (Front-running) [8] – headed up by Wells Fargo, this PoC introduced higher levels of logical representation to the ontology, to demonstrate automated compliance with US regulation W against front-running in securities trading.[9]
- Derivatives[10] – headed up by State Street, this PoC demonstrated the use of FIBO in mapping and integrating derivatives data from various sources.
- Financial Reporting [9] – carried out at the Bank of England, demonstrating how use of a common ontology language could aid reporting.

These proofs of concept were highly successful in demonstrating the application of semantic technologies, including automated reasoning and semantic querying, for the target use cases.

However, by demonstrating the use of ontology for a single use case or a limited range of business scenarios, the primary benefit of ontology is not demonstrated in most of these PoCs, namely the ability to use the ontology as a common language across different sets of application data. The technology has been showcased very well, but the benefit of shared semantics for the most part has not.

Some of these initiatives focused unduly on words or data structures rather than real-world meaning. For example, modelers might take a common word, such as “Affiliate” in the Regulation W PoC, and try to account for most uses of the word even if these pointed to distinct concepts, leading to classes with optional properties. Addressing this has been considered out of scope for a proof of concept. There is then no guarantee that the same ontology could be applied in a different business scenario.

There is a distinction between the real-world features that give something its meaning, such as the legal capacities that define a legal person or a bank (what we call “*truth makers*”), and data corresponding to those things, such as company registration numbers or banking licenses. These relationships were not always explicitly reviewed or documented in these proofs of concept.

---

[8] Newman, D., & Bennett, M. (2012). Semantic Solutions for Financial Industry Systemic Risk Analysis (presentation). *Next Generation Financial Cyberinfrastructure Workshop, Robert H. Smith School of Business, University of Maryland*. Available: *https://wiki.umiacs.umd.edu/clip/ngfci/images/8/80/BennettNewman.pdf*

[9] U.S. Electronic Code of Federal Regulations (2002). Available: *https://www.govinfo.gov/app/collection/cfr/2002/*

[10] David, E., (3 June 2016). FIBO Marches Forward: A Look Inside State Street's FIBO Proof of Concept. *Waters Technology*. Available: *https://www.waterstechnology.com/data-management/2459451/fibo-marches-forward-a-look-inside-state-streets-fibo-proof-of-concept*

These tendencies – a focus on words rather than concepts, data and data shapes rather than truth makers and so on – will have no impact on the integrity of any stand-alone ontology application. However, there is no guarantee that an ontology developed for one use case will capture concepts reflected in data across the rest of the problem space. This may go some way towards explaining why these have not enabled a seamless connection between data and reality.

3.2 ONTOLOGIES IN WIDER INDUSTRY

A more principled approach to ontology modeling can be seen in the standardization work of the Industrial Ontology Foundry (IOF)[11]. This is an ontology standards project to represent subject matter in industrial applications such as manufacturing and supply chain management. The IOF ontology uses a Top Level Ontology (TLO) called the Basic Formal Ontology (BFO)[12]. BFO is widely used in a range of scientific applications and is distinguished by having a "Realist" stance, meaning that a BFO ontology may only represent things that exist in some real or possible world.

While BFO may not be the best fit for some industrial subject matter, for example in reflecting the distinctions between intended designs and actuality, the IOF experience clearly demonstrates that the use of a well-established TLO brings much-needed rigor and discipline to the set of ontologies that are developed under this standard.

Similarly the Pistoia Alliance[13] has undertaken the first phase of a project to model the processes or *recipes*, by which pharmaceutical products are manufactured[14]. This uses a TLO and makes clear distinctions between for example an intended or planned process activity, and the actuality of carrying out that activity.

Another good use of a TLO has been observed in a project under way for the United Kingdom's Digital Twin project[15]. One case study in this program[16] addresses energy efficiency in the UK housing stock and uses an ontology to integrate concepts such as household energy efficiency appraisals and benchmarks. The TLO used is the Information Exchange Standard (IES4)[17], a 4D realist ontology that follows the BORO Methodology[18]. Again some of the concepts in play in these datasets, such as energy efficiency benchmarks, are not a trivial fit to this specific TLO, but the existence of the TLO and the discipline that it represents [9] means that the resulting

[11] *https://oagi.org/pages/industrial-ontologies*

[12] *https://basic-formal-ontology.org/*

[13] *https://www.pistoiaalliance.org/*

[14] *https://www.pistoiaalliance.org/projects/pharmaceutical-cmc-process-ontology/*

[15] *https://www.gov.uk/government/collections/the-national-digital-twin-programme-ndtp*

[16] Retrofitting in housing: National Digital Twin Programme (NDTP). Available: *https://www.gov.uk/government/case-studies/retrofitting-in-housing-national-digital-twin-programme-ndtp*

[17] *https://telicent.io/ies-ontology/*

[18] *https://www.borosolutions.net/librarysearch?search_api_views_fulltext=*

ontologies are not only a good representation of the subject matter but also are mappable to the available data sources.

3.3 TOP LEVEL ONTOLOGIES

Some important conceptual nuances are best addressed by "partitioning" the ontology to reflect different concerns. These include for example the distinction between a thing in itself versus a thing in some role or context (such as a legal entity versus a client or supplier), and issues regarding occurrences, business processes, plans and actualities, temporality, quantities and values, and many others.

Another outcome of partitioning the ontology is the ability to deal with data as a kind of thing, alongside representations of other kinds thing. Given the kind of ontology in which these concerns are both represented, the resulting ontology framework is able to represent data to world relationships such as data provenance, data quality measures, timeliness and accuracy.

3.4 ONTOLOGY OBSERVATIONS

These case studies reflect the importance of a principled approach to the ontological modeling of real things in the world, and the application of established techniques for this kind of ontology.

To provide the kind of ontology described here (sometimes called a *concept ontology*), the organization needs to ensure that a range of ontology techniques is brought to bear, within a coherent semantics strategy. These include:

- A clear commitment that the ontology represents things in the world
- A principled application of classification theory
- Use of formal logic to represent what kind of thing something is and what distinguishes it from other things
- A focus on concepts rather than words
- Capturing what things mean independently of any specific use case
- The relationship between real-world truth makers and data surrogates for those things
- Clarity about the ontological stance or ontological commitment (e.g. Realism versus Conceptualism, 4D versus 3D and other such distinctions)
- Use of a good Top Level Ontology, such as BFO, IES or UFO[19]
- Use of an ontology in which real world value spaces and data representing measures of those values are both represented and segregated.

[19] Unified Foundational Ontology (UFO). Available: *https://nemo.inf.ufes.br/en/projetos/ufo/*

4 DIGITAL TWINS CASE STUDIES

In the Digital Twins arena firms like ExxonMobil and Nextspace demonstrate a data-centric separation of concerns, with data being maintained separately from its projection for consumption.

4.1 CASE STUDY: EXXONMOBIL

The oil company ExxonMobil has been developing a system called the "Digital Reality Ecosystem"(DRE)[20]. This is a Digital Twin initiative, addressing the provision of digital twins that unify data from surveys, scans, design artifacts and other sources. The main goal of this system has been to be able to unify a full range of information assets, including 2D assets like diagrams, 3D assets, reality capture and so on (Figure 4-1).

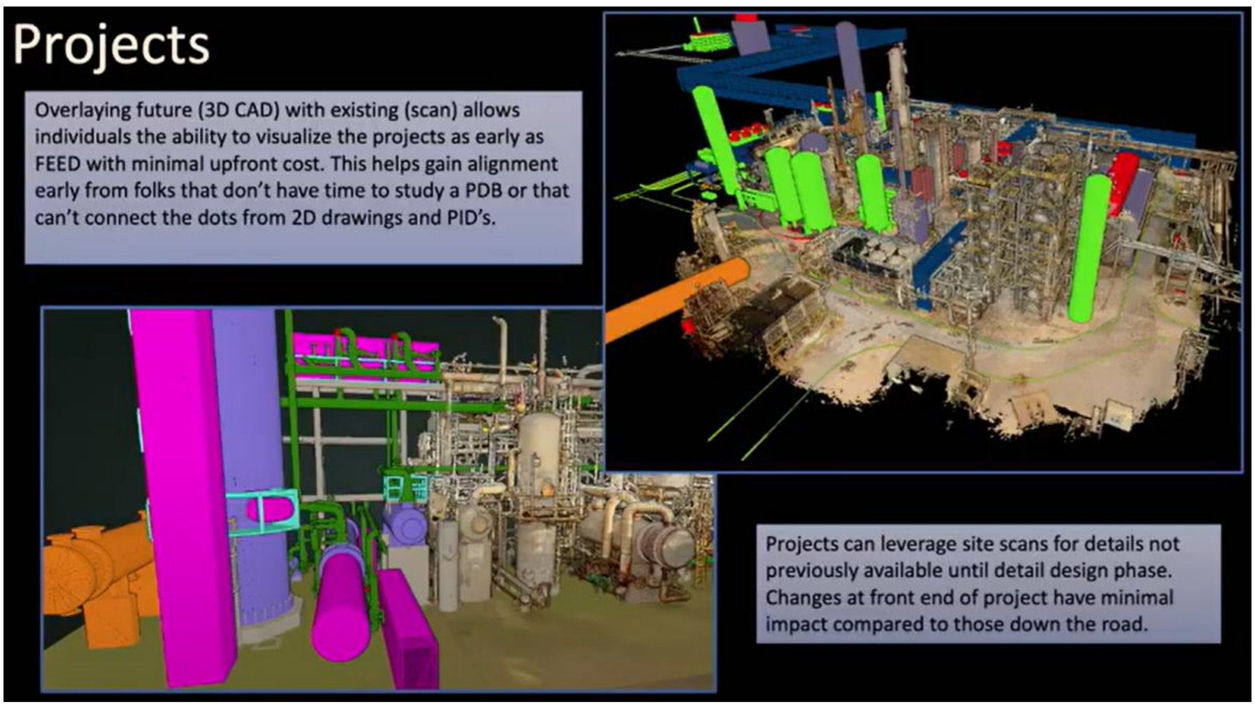

Figure 4-1: Example data formats from ExxonMobil case study.

To achieve this, the approach to digital twins is broken into three aspects: Simulation, Visualization and Integration. Data is maintained independently of its consumption. Data from CAD systems, PDF, design models, specifications and so on is stored as one central resource, from which information is consumed by the end user in one of several virtualization environments, such as Unreal Engine[21]. The key feature of this data-centricity is the separation of form from

[20] Daughtry , K., Dhanormchitphong, A. (2023). Case Study: ExxonMobil's Journey with the Digital Reality Ecosystem. *Augmented Enterprise Summit, 2023*. Available: *https://www.youtube.com/watch?v=iYTzml2gFuM&ab_channel=BrainXchange*

[21] Unreal Engine. Available: *https://www.unrealengine.com/en-US*

content: the presentation of the content in a virtual space is not itself the model; it is an expression of that model.

This initiative is ongoing, and ExxonMobil have made it clear that they seek and welcome engagement from other industry players and standards development organizations.

4.2 CASE STUDY: NEXTSPACE

A company called Nextspace[22] has taken this data-centric thinking to another level [10]. Like ExxonMobil, their platform separates form from content and supports a seamless transition between one virtuality engine and another, showing 3D views, point cloud information and CAD information that can be rendered as the same information across a range of virtualization environments (Figure 4-2).

[22] Nextspace (home page). Available: *https://www.nextspace.com/*

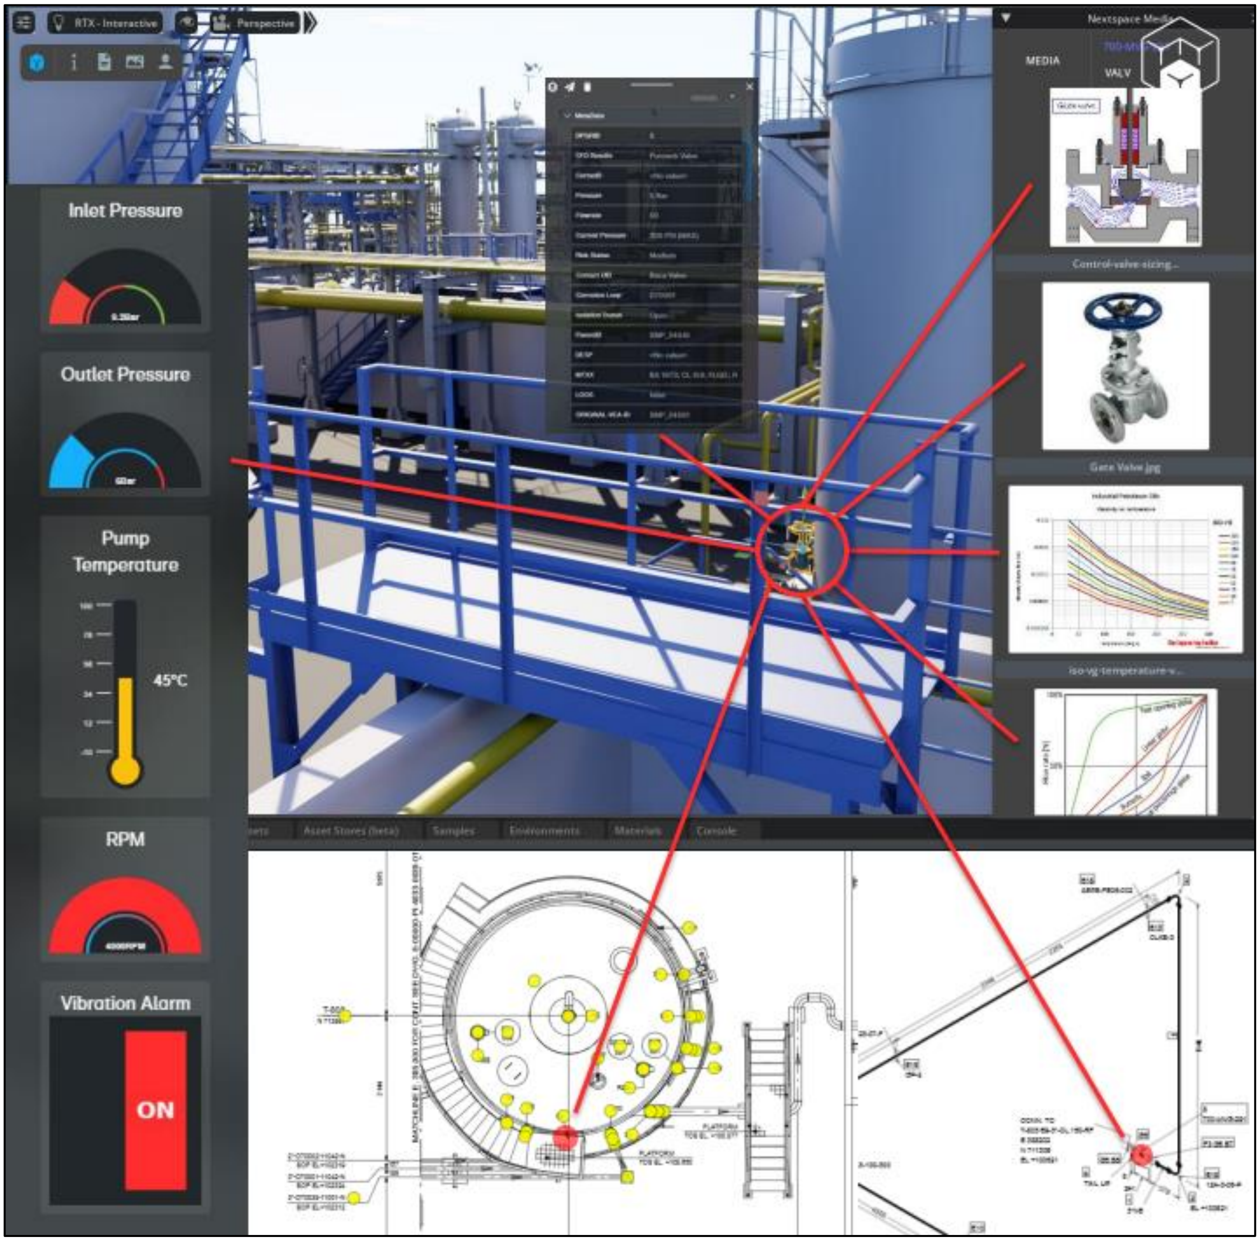

Figure 4-2: Information integration in Nextspace.

The results for the information consumer are comparable to ExxonMobil's Digital Realty Ecosystem. In this case however, the single source of data that lies at the heart of these models is semantic graph data, represented in a triple store and given meaning by the use of formal ontologies.

The ontology-enabled data resource used by Nextspace integrates 1D and 2D assets such as parts lists and specifications alongside different 3D formats in the same set of ontology-defined data. They use LLMs to aid with the mapping of incoming data formats to the ontology.

The use of internationally recognized standards for the data and its semantics (RDF and OWL respectively) makes this data more portable.

5 Additional Observations

5.1 Form Versus Content

In both ExxonMobil and Nextspace, if we consider both the model and the model content, that content has a life independent of the tools in which it is presented.

This is true information-centricity. By decoupling the data from the application, the organization is able to manage the data as a resource in its own right. Data is an asset, with its own lifecycle management covering data provenance, data ownership and data quality measures such as timeliness, accuracy and completeness.

When maintained as an asset in its own right, data outlives individual applications. Data may be updated and ingested by different applications at different points in the data lifecycle and these applications may be updated or replaced, while the data continues to evolve at its own pace.

Segregating data from applications not only enables improved user experiences and data management but also opens the way to smoother incorporation of artificial intelligence applications.

5.2 Artificial Intelligence

The potential of Artificial Intelligence (AI) is enhanced by the provision of an information-centric set of data assets, both for symbolic and generative classes of AI. Generative AI such as LLMs can be trained on user data that is in a common format, while symbolic AI may be supported directly by formal ontologies. We are already seeing a range of ways emerging in which LLMs and ontologies can be used together, for example the Fall 2023 season of the 2024 Ontology Summit [11].

5.3 Semantics, Data and Presentation

Larger institutions are already seeing the benefits of using formal ontologies along with enterprise knowledge graphs to provide a data-centric ecosystem. The next phase in truly information-centric working requires the separation of concerns between the data and its presentation. These kinds of implementation are already emerging in the area of industrial digital twins. Future developments may take things beyond that to something more ubiquitous, moving towards a more inclusive metaverse.

Data centric implementation along with visualization points to a three-way separation between semantics, data and presentation, as shown in Figure 5-1.

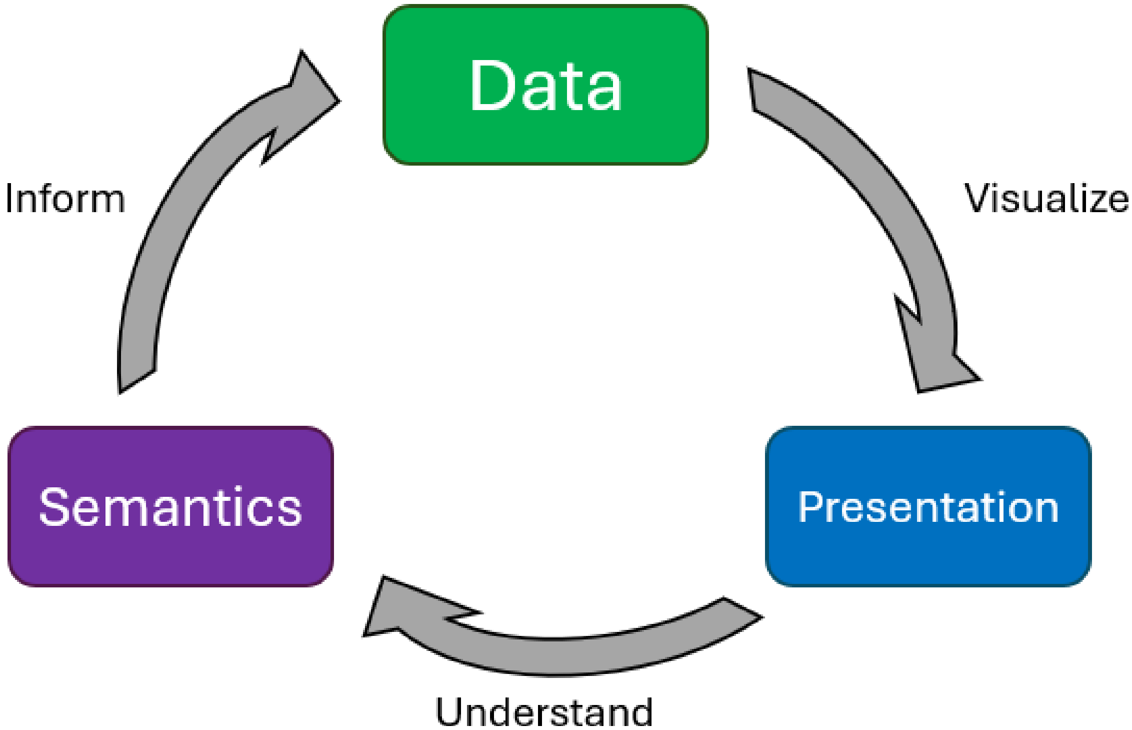

Figure 5-1: Semantics, data, and presentation.

6 STANDARDS

There may be standards applicable both for things that people are doing now in the world of data centric operations and digital twins, and in future possibilities that will likely only become clear once things move to a more ubiquitously information-centric way of thinking and working.

Considering the two kinds of case studies described above, the move towards data-centricity can perhaps be benefitted by the use or development of standards in two related areas:

1. The creation and curation of real-world oriented or concept ontologies
2. Data-centric digital twins and separation of concerns

Combining these would bring the potential of formal ontology to the provision of data-centric digital twinning implementations in industrial applications and elsewhere.

A pre-requisite of data centricity is the separation of form from content, something which is baked into many OMG specifications, such as UML[23] and SysML[24]. At the heart of the OMG's thinking is the separation of concerns, as documented in Model Driven Architecture (MDA)[25].

[23] *https://www.omg.org/spec/UML*
[24] *https://www.omg.org/spec/SysML/1.7*
[25] *https://www.omg.org/mda/*

6.1 SEMANTICS IN STANDARDS

A necessary feature of information-centric architectures is the provision of common meaning. In industrial digital twins this may use a semantically-enabled systems modeling language like SysML V2[26] with its KerML[27] core semantics. For less tangible things such as financial instruments or contracts W3C based ontology standards have been used to define a formal ontology.

"Semantics" can refer to model semantics or real-world or domain semantics. Advances in KerML and SysML V2 bring these languages closer to being able to reflect real-world meanings of things represented in the models. These take a different approach to conventional ontology languages such as OWL.

Industrial digital twins may make use of system design models in SysML for data-centric working. The examples from ExxonMobil and Nextspace show the semantic unification of 2D and 3D assets and reality capture point clouds. The Nextspace example shows that it is possible to find a way to unify these different dimensions under a common ontology. Standardizing these approaches would be a benefit to the industry.

An ontology that is predicated on real things can also include among those real things, that which is data, thereby incorporating the relationships between the data and the things it is about. The ability to model this distinction is particularly relevant in the case of industrial digital twins, as it enables the enterprise to manage data concerns such as measurement tolerances, instrument calibration history and the like.

For ontologies that represent the real world, there is a need to standardize the ontological commitment. There is an ISO standard, ISO 21838, in which Part 1 [12] sets out the definition of what it is to be a Top Level Ontology. In time there may be a need to further characterize the details of ontological commitments. This includes for example whether it is a Realist ontology, whether it reflects a 3D or 4D view of the world, and other similar considerations.

Every ontology used across the enterprise or project should then be able to declare these features in an unambiguous and standardized way.

6.2 METAVERSE AND DIGITAL TWINS STANDARDS

Standards cover more kinds of subject matter than data formats or modeling languages. Anything that can be stated in normative terms ("you shall do this") can be proposed and formally adopted as a standard. Standards may define formats, business processes, architectures, interfaces and many other concerns. For data-centric digital twins, these may range from being able to formally state what kind of model is used and how this represents reality, through to standardizing the way data capture, storage, distribution and consumption are done.

---

[26] Systems Modeling Language (SysML) V2 (Beta). Available: *https://www.omg.org/spec/SysML*

[27] Kernal Modeling Language (KerML) (Beta). Available: *https://www.omg.org/spec/KerML*

The combination of digital twin data and extended reality (XR and VR) representation of that data effectively defines a key part of what's become known as the Metaverse. The relationships between these various components of the metaverse can in principle be standardized, and many already are, as documented by the Metaverse Standards Forum[28].

Some of the same tools, techniques, environments and potential standards apply across both the industrial metaverse and what we might call the retail or entertainment Metaverse. Others will not. You may not be able to pick up a magic sword in a virtual oil platform, but you may need to pick up a gas analyzer at some point (Figure 6-1). There will be many points of commonality and convergence between these worlds and their standards.

Figure 6-1: A virtual gas analyzer in VR training session[29].

The use of ontologies looks set to be the next evolution in digital twin and metaverse standards. As well as standards for identifying styles or type of ontology and use of top level ontologies, we can expect to see further standardization use cases in the ways in which the ontologically defined data is integrated, mapped and rendered into 3D representations in a range of target 3D rendering environments including virtual reality engines.

6.3 STANDARDS BENEFITS

We have seen how a truly information-centric way of working requires a radical inversion of the relationship between data and applications, and how the use of formal ontologies provides an interim step towards that, which can be and is being put into play right now. Businesses are already seeing benefits in cost and risk reduction from the use of formal ontologies, both in data

[28] *https://metaverse-standards.org/*

[29] Dräger Offshore Rig VR Experience. Available: *https://demodern.com/projects/offshore-rig-virtual-reality-experience*

integration and as a means to seed a more enterprise-wide graph of reusable data in a common, application-neutral format, known as a *knowledge graph*[30].

Standards present a mutual benefit arrangement, a resolution of the tragedy of the commons[31], whereby entities can benefit by collective action in ways which could disadvantage each if carried out in isolation. By participating in standards development, organizations can gain commercial benefit, both from the efficiencies of more consistent interactions between data, visualizations and semantics, and in the new business opportunities that would emerge from a more information centric way of working.

The next steps in standards use for the metaverse and digital twins will speed the adoption and interoperability of 3D asset data and digital engineering, digital threads, extended reality and data visualization. We may see these initially in smart cities, energy and oil and gas. With standards driven by these industries, we should expect to see further benefits from applying these same standards in other industries.

7 Conclusions

For industry and consumers to move towards information-centric working, standards will be needed at several levels, ranging from standardizing how real-world concepts are modeled in ontologies, through to standards for shared data spaces such as the International Data Spaces initiative[32] and the European Gaia-X project[33]. A more complete inversion of data and applications will also require standards connecting data visualization techniques and datasets.

These standards will need broad participation by industry and academia, with benefits accruing to industry such as having greater trust in data and enabling a marketplace of applications and data visualizations within an institution-wide graph of knowledge, offering flexibility and resilience to end users.

With standards-enabled architectures in place, firms will gain greater value from existing data assets, be resilient against data loss and quality issues, implement FAIR principles and provide provable, responsible deployments of generative AI based solutions, deploying a common sea of knowledge across the enterprise.

[30] See for example the Enterprise Knowledge Graph Forum (an OMG Managed Community). Available: *https://www.ekgf.org/*

[31] The notion of the Tragedy of the Commons was introduced by Garret Hardin [13] in the context of shared resources, but the principle whereby groups can act for mutual benefit in ways which would not benefit them if they acted alone has broader applicability, for example in industry regulation, standards, security and so on.

[32] *https://internationaldataspaces.org/*

[33] *https://gaia-x.eu/*

8 REFERENCES

[1] Bush, Vannevar (Jul 1945). As We May Think. The Atlantic Monthly. Vol. 176, no. 1. pp. 101–8. Available: *https://www.theatlantic.com/magazine/archive/1945/07/as-we-may-think/3881/*

[2] McComb, D. (2024). Data-Centric: How Big Things Get Done (in IT). *Semantic Arts.* Available: *https://www.semanticarts.com/data-centric-how-big-things-get-done-in-it/*

[3] Mark D. Wilkinson; Michel Dumontier; IJsbrand Jan Aalbersberg; et al. (15 March 2016). The FAIR Guiding Principles for scientific data management and stewardship. *Scientific Data*. *3 (1): 160018*. doi:10.1038/SDATA.2016.18.

[4] Franz Baader, Ian Horrocks, and Ulrike Sattler (2007). Chapter 3 Description Logics. *Frank van Harmelen, Vladimir Lifschitz, and Bruce Porter, editors, Handbook of Knowledge Representation. Elsevier*. doi:10.1016/S1574-6526(07)03028-3

[5] Horrocks, I.; Patel-Schneider, Peter; van Harmelen, Frank (2003). From SHIQ and RDF to OWL: The making of a Web Ontology Language. *Web Semantics: Science, Services and Agents on the World Wide Web. 1: 7–26.* doi:10.1016/j.websem.2003.07.001

[6] Bennett, M. G. (2010). The EDM Council Semantics Repository – Considerations in Ontology Alignment. *Proceedings of the ISWC 2010 Workshops, Volume I, Shanghai, China, November 7-8, 2010. Philippe Cudre-Mauroux, Bijan Parsia, Eds.* ISSN: 1613-0073. Available: *http://ceur-ws.org/Vol-687/*

[7] Grosof, B., Bloomfield, J., Fodor, P., Kifer, M., Grosof, I., Calejo, M., Swift, T. (2015). Automated Decision Support for Financial Regulatory/Policy Compliance, using Textual Rulelog. *Proceedings of the RuleML 2015 Challenge, the Special Track on Rule-based Recommender Systems for the Web of Data, the Special Industry Track and the RuleML 2015 Doctoral Consortium. 9th International Web Rule Symposium (RuleML 2015)*. Available: *http://ceur-ws.org/Vol-1417/paper8.pdf*

[8] Bholat, D. (2016). Modelling metadata in central banks*. ECB Statistics Paper Series. No 13/April 2016*. Available: *https://www.ecb.europa.eu/pub/pdf/scpsps/ecbsp13.en.pdf?a66e0b0448ff62df28d39643573b2b3d*

[9] Partridge, C. (2002). Note: A Couple of Meta-Ontological Choices for Ontological Architectures. *LADSEB-CNR – Technical report 06/02 – Padova, The BORO Program, LADSEB CNR, Italy.* Available: *https://www.academia.edu/483241/Note_A_couple_of_meta_ontological_choices_for_ontological_architectures*

[10] Nextspace Technical Specifications. Available: *https://www.nextspace.com/products/nextspace/technical-specifications*

[11] Ontology Summit (2020): Knowledge Graphs. Available: *https://ontologforum.org/index.php/OntologySummit2020*

[12] ISO (2021). Information technology — Top-level ontologies (TLO) — Part 1: Requirements. Available: *https://www.iso.org/standard/71954.html*

[13] Hardin, G (1968). The Tragedy of the Commons. *Science*. *162 (3859): 1243–1248.* doi:10.1126/science.162.3859.1243.

ACKNOWLEDGEMENTS

The views expressed in the *OMG Journal of Innovation* are the author's views and do not necessarily represent the views of their respective employers nor those of the Object Management Group® (OMG®).

➢ Return to *OMG Journal of Innovation landing page* for more articles and past editions.