



# Introducing FiORO: The Financial Industry Operational Risk Ontology



Prof. Tom Butler  
Dr Tabbasum Naz  
Dr Marcello Ceci



## Abstract

The Financial Industry Operational Risk Ontology (FiORO) helps address several problems in financial services organizations. Briefly, its primary purpose is to enable the systematic identification, assessment, management, mitigation and regulatory compliance reporting of operational risks (OR) in a financial services organisation. As a knowledge base, it also facilitates knowledge sharing around operational risks and hence may also be used for staff training and regulatory compliance. It will ultimately contribute to the development of 'predictive' operational risk analytics. Operational risk emanates from failed people, processes and technologies, both internal and external (e.g. internal and external fraud, cybercrime etc.). In a typical financial services organisation, operational risks occur across business lines. Hence, enterprise-level operational risk management requires the identification, categorisation and assessment of OR events at a granular level, if such risks are to be managed mitigated and reported. The Financial Industry Operational Risk Ontology (FiORO) is expressed in the Resource Description Framework (RDF)/ Web Ontology Language (OWL). It is, therefore, human readable (using Protégé or a Graph Visualizer) and machine readable, using computer-based software algorithms, reasoners, and other software applications. As a Master Data Model, FiORO will enable operational risks data from heterogeneous silos to be virtualized and/or persisted in a graph data store for querying using SPARQL, analysis using standard reasoners and inferencing engines, and visualization using Open Source and/or proprietary tools. It is also well recognized that financial services organisations also require an OR knowledge base in order to ensure that staff are knowledgeable about operational risks. Thus, the proposed ontology aims at facilitating OR information sharing within and across organizational units and business lines.

## Introduction

This white paper presents a Vocabulary for Operational Risk (VOR), consisting of a taxonomy for Operational Risk Management captured in an OWL ontology (the Financial Industry Operational Risk Ontology or FiORO). The objective of enrichment with axioms is aimed at answering specific competency questions. As a preliminary explanation, the concepts of Operational Risk Management and of Computational Ontology will be introduced.

## Operational Risk

Operational Risk is defined as the risk of negative financial, business and/or reputation impacts resulting from inadequate or failed internal governance and business processes, people, systems, or from external events<sup>1</sup>. Operational Risk Management is the entire process of continuously and systematically identifying, analyzing, responding to, reporting on and monitoring operational risks. Managing operational risks is usually performed through self-assessments and collection by single business areas of event data, which are then consolidated at the corporate level. Because various business areas use different information systems, it is not easy for the operational risk management group to integrate them. This led to growing numbers of books, articles and conferences being devoted to operational risk management. Despite this, very little academic research has been conducted to provide a better understanding of how operational risk management can be improved within organizations.<sup>2</sup>

---

<sup>1</sup> BCBS, Sound Practices for the Management and Supervision of Operational Risk, in: Sound Practices for the Management and Supervision of Operational Risk, Basel Committee on Banking Supervision, 2003.

<sup>2</sup> S.R. Iyer et al., "Academic Research on Enterprise Risk Management", Enterprise Risk Management, J. D. a B. S. (eds.), ed., Wiley, 2010.].

## What is an Ontology?

In Computer Science, an ontology is a semantic tree that “*defines the basic terms and relations that compose the dictionary of the field of interest and the rules that combine the terms and relations so that the dictionary of terms is extended*”. Through the use of Description Logic it is possible to express any object or concept in a formal way. Ontologies contain the specifications of those objects or concepts which are necessary for the understanding of the knowledge domain, its vocabulary, the way concepts and vocabularies are connected, and the way classes, instances and their properties are defined. The Financial Industry Operational Risk Ontology (FiORO) belongs to the category of Domain Ontologies, subclass Enterprise Ontologies<sup>3</sup>. This category includes several interesting experiments such as the Semantic Business Process Management (SBPM) frame<sup>4</sup>, STRATrisk for managing risks in strategy level<sup>5</sup>, as well as others have been delineated<sup>6</sup>.

The research presented in this paper has a progressive approach to the subject: as already explained, the tool initially consists in a comprehensive taxonomy for Operational Risk, extending the official Basel II classification with more than 300 subclasses which specify two more layers of analysis beyond the original two. It also includes a taxonomy of business lines, losses, and product types. In addition, the FiORO includes modules whose focus is strictly driven by proof of concepts and foreseeable practical applications.

The intended outcome of this research is thus not a single, general-purpose ontology, but rather a set of ontology modules built around a taxonomy for operational risk, each of them with its own task(s), which can be represented, for example, by a decision tree that the module will be able to follow. This will, from the beginning to the end, query/node after query/node, enable queries to be answered through axioms and instances contained in the ABox, and delivering the outcome of the decision tree as an output of this set of queries.

The remainder of this white paper is structured as follows: We present the state-of-the-art and motivation for this research. Then we present the FiORO, its design approach and its main components. The final section concludes by showing the possible uses of the ontology as well as plans for further research in that subject.

## Background and Motivation

The subject of risk management is covered by the Basel II regulatory framework for capital adequacy, which distinguishes between financial, market and operational risks. The category of operational risk is further specified along two levels into 15 subtypes. A slightly different classification is put forward by different authoritative sources such as the “*common risk classification system for the actuarial profession*” (23 subtypes) and the “*operational risk reporting standards*” issued by the ORX Association (19 subtypes). In addition to that, private stakeholders in the financial industry have developed their own instruments

---

<sup>3</sup> R. Rittgen, Handbook of Ontologies for Business Interaction, IGI, 2008.

<sup>4</sup> M. Hepp and D. Roman, “An Ontology Framework for Semantic Business Process Management”, 8<sup>th</sup> international conference wirtschaftsinformatik 2007, pp.423-440.

<sup>5</sup> Y. Jussupova-mariethoz and A. Probst, “Business concepts ontology for an enterprise performance and competences monitoring, comput. Ind., vol. 58, no. 2, 2007, 118-129.

<sup>6</sup> C. Blanco et al., A systematic review and comparison of security ontologies. In proceedings of the 2008 third international conference on availability, reliability and security, 2008, pp. 813-820.

to further describe standards and guidelines for reporting Operational Risk events and losses. These definitions and principles, conceived for the purpose of promoting consistency of reporting and data, constitute a fundamental source for the creation of intelligent regulatory compliance systems and RCMSs: what was created to organize available data on past events can be used to categorize and reason on new data concerning events as they occur.

Concerning the use of external loss data, Ames et al.<sup>7</sup> presents an approach to reduce the uncertainty around effective use of such data. This approach consists in *“Expand[ing] data capture beyond the Basel event type and business line classifications to allow further investigation into key explanatory factors. Examples could include product, flags for litigated events, relational links to insurance policy and coverage, etc.”*. The recommendation of the authors is *“not that the current taxonomy be thrown out, but rather complemented by additional field specifications as product type, litigated vs. not-litigated liability, more detailed data on type and timing of recoveries [...], and others that could be further exploited through modelling the factors that would better explain loss.”*

Building a knowledge base for operational risk management is not an easy task. Being the category of Operational Risk quite heterogeneous, classification of OR-related concepts is not straightforward. The research towards such classification can however rely on interesting previous research in the field of Domain Ontologies, subclass Enterprise Ontologies<sup>8</sup>. Examples in this field include the Semantic Business process Management<sup>9</sup>, attempting the complete corresponding of all layers of processes in an organization (from administrative to functional), BusCo<sup>10</sup>, and STRATrisk, focused on the strategy layer<sup>11</sup>. More recently, ontologies have addressed the operational risk information sharing issue by attempting to go beyond the representation of the main concepts of the operational risk management knowledge domain.<sup>12</sup> Such ontologies are said to aid the operational risk management group in gathering heterogeneous information and using it in support of the decision making process, and to enable computation and inference over information residing in heterogeneous data stores. The approach adopted by the authors to build the ontology is similar to that adopted in the present research: it starts off from the definition of requirements through a set of competency questions, it proceeds by defining a glossary of concepts and relations and it completes by establishing a conceptual network out of those two components.

---

<sup>7</sup> Ames, Mark and Schuermann, Til and Scott, Hal S., Bank Capital for Operational Risk: A Tale of Fragility and Instability (February 10, 2014). Available at SSRN: <http://ssrn.com/abstract=2396046> or <http://dx.doi.org/10.2139/ssrn.2396046>.

<sup>8</sup> R. Rittgen, Handbook of ontologies for Business Interaction, IGI, 2008

<sup>9</sup> M. Hepp and D. Roman, “An Ontology Framework for Semantic Business Process Management”, Proc. 8th international conference Wirtschaftsinformatik 2007, pp. 423-440.

<sup>10</sup> Y. Jussupova-Mariethoz and A. Probst, “Business concepts ontology for an enterprise performance and competences monitoring”. Comput. Ind., vol. 58, no. 2, 2007, pp. 118-129

<sup>11</sup> G. M. Allan, et al., “Risk Mining for Strategic Decision Making”, Proc. Adv. In Intel.Web ASC 43, 2007, 21-28.

<sup>12</sup> Kenett, Ron S., and Yossi Raanan, eds. *Operational Risk Management: a practical approach to intelligent data analysis*. Vol. 106. John Wiley & Sons, 2011.

An interesting metamodel for intelligent regulatory compliance, aimed at the development and implementation of risk strategies and adaptive information systems has been developed<sup>13,14</sup>. The model relies on two technologies by the Object Management Group (OMG), namely the Business Motivation Model<sup>15</sup> and the Semantics of Business Vocabularies of Business Rules<sup>16</sup>, which is used also in the GRCTC. An important value of the ontology derived from the metamodel is its focus on concepts such as loss events, loss types, and business lines.

## Towards an Ontology for Operational Risk

Ontologies are widely used in ontology engineering. *“Ontology engineering in computer science and information science is a field which studies the methods and methodologies for building ontologies: formal representations of a set of concepts within a domain and the relationships between those concepts* . According to G. Antonio et al. the main phases involved in the ontology engineering process are determine scope, consider reuse, enumerate terms, define taxonomy, define properties, define facets, define instances and check for anomalies. Below each step was considered carefully while developing FiORO for the operational risk domain.

### 1.1. Determine Scope

At this stage an ontology engineer must be able to answer the questions like what is the domain that the ontology will cover? For what we are going to use the ontology? For what types of questions the information in the ontology should provide answers? Who will use and maintain the ontology? In our case, the domain of the ontology is operational risk. *“Operational risk is defined as the risk of loss resulting from inadequate or failed internal processes, people and systems or from external events”*<sup>17</sup>. Here FiORO will be used for operational risk management.

When going beyond a simple taxonomy, it is necessary to rely on requirements, expressed as a list of competency questions, for example:-

1. What is the number of financial loss events occurring in the year x?
2. What is the average financial loss in the year x?
3. Provide the aggregate financial loss in the year x?
4. Provide the number of loss events in the region x?
5. Provide the aggregate financial loss in the region x?
6. Provide the average financial loss in the region x?
7. Provide the list of financial loss events by business line x and event type y?
8. Provide the frequency of financial loss events by business line x and event type y?

<sup>13</sup> Spies, Marcus. "An ontology modelling perspective on business reporting." *Information Systems* 35, no. 4 (2010): 404-416.

<sup>14</sup> Leibold, Christian, Hans-Ulrich Krieger, and Marcus Spies. "Ontology-Based Modelling and Reasoning in Operational Risks." *Operational Risk Management* 2010, p. 39-59.

<sup>15</sup> BMM, OMG, Business motivation Model, v 1.0, 2008.

<sup>16</sup> SBVR, OMG, Semantics of Business Vocabulary and Business Rules, v1.3, 2015.

<sup>17</sup> Basel II Accord section V.A.S644.

9. Provide the total number of financial loss events falling within financial loss size boundary (i.e. 20k – 50k)?
10. Provide the total aggregate financial losses falling within loss size boundary (i.e. 20k – 50k)?

During the ontology engineering process, we have also identified the following Incident related questions that can be answered using FiORO.

1. Who is the assigned incident manager for the incident x?
2. Who is the business owner accountable for the risk event x?
3. Who is the key contact for the risk event x?
4. What is the reporting date of the incident x?
5. What is the status of the risk event x i.e. open or closed?
6. What is the closure date for the event x?
7. Provide the summary/description of the incident x?
8. What was the cause of the incident x?
9. Provide the latest update for the incident x.
10. Was the incident caused by third party? If yes, provide the name of third party?
11. Is the event a repeat issue?
12. Does the error impact other business units? If yes, which business units?
13. What is the operational risk event type for the incident x?
14. How was the incident identified?
15. Provide the resolution plan in order to resolve the issue x?
16. Is the event x a near miss?
17. What is the potential aggregate financial loss for the event x?
18. How much is the net financial loss i.e. actual gross financial loss minus total recoveries?
19. Provide the Basel event type, product type and process type for the event x.
20. Is it a credit risk related event? Significance?
21. Is it a market risk related event? Significance?
22. What is the name/GL BL2 code/EC2 code/RefName of loss x?
23. Is the loss x in own data (Y/N)?
24. Is the loss related to external credit (Y/N)? Significance?
25. What process/product is the financial loss x related to?
26. In which region has the financial loss occurred?
27. What is the date of occurrence/discovery/recognition for the financial loss x?
28. What is the event type of the financial loss/risk x?

This is a long list of competency questions that an operational risk ontology such as FiORO must provide answers for, but only few of them are provided here.

## Reusing Extant Ontologies

During this phase, we needed to check if there already exist ontology that can be reused. According to our knowledge, there does not exist suitable ontologies in the operational risk domain that can answer the competency questions presented above.

## Enumerating Terms

During this phase, we identified all the relevant terms that are expected to appear in ontology in an unstructured list. We have identified categories or classes such as business line, event type, product type, process type, incident, loss, region, operational risk, private banking, retail banking, retail, brokerage, asset management, fraud, internal fraud, external fraud, natural disaster, capital raising, retail brokerage, deposits etc.

- *BusinessLine* class in the ontology represents the profit centers where the revenues are generated from third parties, not allocations from other parts of the firm.
- *OperationalRisk* is subclass of *EventType*. Operational risk event is an event leading to the actual outcome(s) of a business process to differ from the expected outcome(s), due to inadequate or failed process, people and systems, or due to external facts or circumstances.
- *ProductType* class in the ontology represents the products, which also includes services, are the sources of revenue for a bank via direct or indirect fees. The main goal of the *ProductType* class is to improve the understanding of nature of losses, identify the chronic and recurring weaknesses and promote values added dialogue with the businesses and functional areas regarding the impact of their operational risk experience and potential operational risk exposure.
- *ProcessType* class represents that a business process is a set of coordinated tasks and activities the will lead to accomplishing a specific organizational goal i.e. a sequence of interdependent and linked procedures which consume one or more resources to convert input into outputs. The main objective of the process type is to increase the understanding of the nature of losses, identify sources of loss concentrations, identify chronic and recurring weaknesses and promote value added dialogue with the businesses and functional areas regarding the impact of their operational risk experience and potential operational risk exposure.
- *Loss* class in the ontology represents different types of losses including *ExcessClaim*, *Fine*, *ImpairmentProvisions*, *OperationalLoss*, *RiskCrystallization* and *StolenAsset* with further sub-classes in it.
- *Region* class contains different regions of the world.
- *Incident* class represents the failed process.

## The FiORO Taxonomy

The next phase in ontology engineering FiORO was to organize the terms in taxonomic order. Our taxonomy for Operational Risk has been defined and is presented in Figure 1.



Figure 1: Taxonomy for Operational Risk

## Defining Properties

Next, we defined the object, data type and annotation properties of FiORO. The owl:topObjectProperty is the class of all object properties, and hence is a binary predicate relating all individuals in the model to all individuals in the model. Every object property is a subproperty of this property. We have identified following object type properties for the operational risk ontology.

- hasBusinessLine
- hasEventType
- hasLoss
- hasProcessType
- hasProductType
- hasRegion

Every incident has a business line. In operation risk ontology, we can model it as subject (Incident), predicate (hasBusinessLine) and object (BusinessLine). The link between incident and business line is maintained through hasBusinessLine object property as shown in Figure 2.

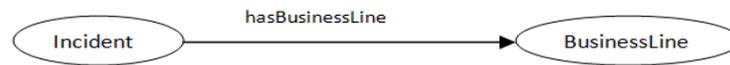


Figure 2 Subject Predicate Object (SPO) in Operational Risk Ontology

The owl:topDataProperty is the class of all data properties, and hence is a binary predicate relating all individuals to all literals. Every data property is a subproperty of this property. We have identified following data type properties for the operational risk ontology:

accrualOrProvision, accrualOrProvisionDate, accrualOrProvisionAmount, actualGrossLoss, assignedIncidentManager, businessLossOwner, businessRiskOwner, closureDate, coveredByInsurance, creditRiskLoss, dateOfIncident, divisionOrBusinessUnit, flagUnofficial, hasDescription, hasID, hasIncidentCause, hasIncidentDescription, hasSynonym, hasWiderImpact, incidentIdentified, incidentName, incidentSummary, isIncidentCausedByThirdParty, isRepeatIssue, keyContact, latestStatusUpdate, lossSizeBoundry, marketRiskLoss, nearMiss, netLoss, potentialGrossLoss, recoveryDirectAmount, recoveryIndirectAmount, reportingDate, resolutionPlan, status, thirdPartyServiceProvider, totalRecoveries, unitFunctionOrBranch, widerGroupImpact, writeOff, writeOffAmount and writeOffDate.

The hasID data type property has further sub-properties i.e. boiID, crcID and orxID. The hasIncidentDescription property is also further subdivided into sub-properties i.e. controlsInPlace, explanation, failedControls, grossLossAmount, howIssueWasDiscovered, howItOccured, howTheIncidentBeingInvestigated, impactOnCustomers, isLossRecoverable, lessonsLearned, lossRecoverableAmount, preventativeMeasures, timeline, whatHappened and whosResponsible.

The data type boolean, string, dateTime and double etc is assigned to the data type properties as required.

Annotation properties are binary predicates that provide informal documentation annotations about ontologies, statements, or IRIs. For instance, the rdfs:comment annotation property is used to provide a comment. rdfs:comment, rdfs:label, rdfs:seeAlso and rdfs:isDefinedBy are utilized to provide the comments, labels and definitions etc to the operational risk classes and properties.

## Defining Facets

Next step was to enrich the ontology by facets. Facets are the restrictions in the ontology.

If the operational risk event occurs and the potential loss is of €10k or less then it is not reportable to Group Operational risk. If the potential risk loss is €10K to €250K, then it must be reported on monthly loss return. Otherwise, if the potential risk is 250K+, it must be reported to the Group Operational Risk. To handle these types of loss boundaries, the data type lossSizeBoundry property is proposed in OWL by using minExclusive and maxInclusive, as shown in Figure 3.

```

<!-- http://www.grctc.com/oprisk#lossSizeBoundry -->
<DatatypeProperty rdf:about="&oprisk;lossSizeBoundry">
  <rdfs:domain rdf:resource="&oprisk;Incident"/>
  <rdfs:range>
    <rdfs:Datatype>
      <unionOf rdf:parseType="Collection">
        <rdfs:Datatype>
          <onDatatype rdf:resource="&xsd;decimal"/>
          <withRestrictions rdf:parseType="Collection">
            <rdf:Description>
              <xsd:minExclusive rdf:datatype="&xsd;decimal">0</xsd:minExclusive>
            </rdf:Description>
            <rdf:Description>
              <xsd:maxInclusive rdf:datatype="&xsd;decimal">10</xsd:maxInclusive>
            </rdf:Description>
          </withRestrictions>
        </rdfs:Datatype>
        <rdfs:Datatype>
          <onDatatype rdf:resource="&xsd;decimal"/>
          <withRestrictions rdf:parseType="Collection">
            <rdf:Description>
              <xsd:maxInclusive rdf:datatype="&xsd;decimal">250</xsd:maxInclusive>
            </rdf:Description>
            <rdf:Description>
              <xsd:minExclusive rdf:datatype="&xsd;decimal">10</xsd:minExclusive>
            </rdf:Description>
          </withRestrictions>
        </rdfs:Datatype>
        <rdfs:Datatype>
          <onDatatype rdf:resource="&xsd;decimal"/>
          <withRestrictions rdf:parseType="Collection">
            <rdf:Description>
              <xsd:minExclusive rdf:datatype="&xsd;decimal">250</xsd:minExclusive>
            </rdf:Description>
          </withRestrictions>
        </rdfs:Datatype>
      </unionOf>
    </rdfs:Datatype>
  </rdfs:range>
</DatatypeProperty>

```

Figure 3 OWL Restriction in Operational Risk Ontology

Regarding incident, we want to record that reported incident's status is open or closed. In OWL, we have modeled it by using oneOf as shown in Figure 4.

```

<!-- http://www.grctc.com/oprisk#status -->

<DatatypeProperty rdf:about="&oprisk;status">
  <rdfs:domain rdf:resource="&oprisk;Incident"/>
  <rdfs:range>
    <rdfs:Datatype>
      <oneOf>
        <rdf:Description>
          <rdf:type rdf:resource="&rdf;List"/>
          <rdf:first>closed</rdf:first>
          <rdf:rest>
            <rdf:Description>
              <rdf:type rdf:resource="&rdf;List"/>
              <rdf:first>open</rdf:first>
              <rdf:rest rdf:resource="&rdf:nil"/>
            </rdf:Description>
          </rdf:rest>
        </rdf:Description>
      </oneOf>
    </rdfs:Datatype>
  </rdfs:range>
</DatatypeProperty>

```

Figure 4 Incident Status in OWL

The hasID data type property, has further sub-property orxID with a restriction on it. The restriction on orxID is that it is functional as shown in figure 5. A functional property is a property that can have only one (unique) value y for each instance x, i.e. there cannot be two distinct values y1 and y2 such that the pairs (x,y1) and (x,y2) are both instances of this property. [<https://www.w3.org/TR/owl-ref/#FunctionalProperty-def>]

```

<!-- http://www.grctc.com/oprisk#orxID -->

<DatatypeProperty rdf:about="&oprisk;orxID">
  <rdf:type rdf:resource="&owl;FunctionalProperty"/>
  <rdfs:subPropertyOf rdf:resource="&oprisk;hasID"/>
  <rdfs:range rdf:resource="&xsd:string"/>
</DatatypeProperty>

```

Figure 5 orxID as Functional Property in OWL

## Defining Instances and Checking for Anomalies

In general scheme of things Instances (sample events) are loaded as individuals of incident class in the ontology. The incident's region, process type, product type, loss size, incident date etc is loaded and then ontology is checked for anomalies.

## SPARQL Queries and Results

Based on the requirements and competency questions presented above, we developed SPARQL queries for the operational risk domain and results were extracted. Below we can find a selection of SPARQL queries and their results. Figure 6 shows the SPARQL query and the result for the competency question aimed at identifying the average amount of loss in the Western Europe region.

```

SPARQL query:
PREFIX rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
PREFIX owl: <http://www.w3.org/2002/07/owl#>
PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
PREFIX oprisk: <http://www.grctc.com/oprisk#>

SELECT (AVG(?object) as ?avgLossAmount)
WHERE { ?x oprisk:grossLossAmount ?object . ?x oprisk:hasRegion
oprisk:WesternEurope.
}

avgLossAmount
"100150.0"^^<http://www.w3.org/2001/XMLSchema#double>
    
```

Figure 6 SPARQL Query and Result to Calculate Average Loss Amount in Western Europe

Figure 7 shows the SPARQL query and the result for the competency question aimed at identifying the total number of loss events falling within loss size boundary (i.e. 20k – 50k).

```

SPARQL query:
PREFIX rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
PREFIX owl: <http://www.w3.org/2002/07/owl#>
PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
PREFIX oprisk: <http://www.grctc.com/oprisk#>

SELECT (count (?amount) As ?noOfIncidentsInBoundry)
WHERE { ?incident oprisk:grossLossAmount ?amount .
FILTER (?amount >=20000)
FILTER (?amount <=50000)
}

noOfIncidentsInBoundry
"2"^^<http://www.w3.org/2001/XMLSchema#integer>
    
```

Figure 7 SPARQL Query and Result to Calculate Total Number of Loss Events within Loss Size 20-50K

Other SPARQL queries are posed in the similar way and results are obtained. Figure 8 provides the ontology metrics of operational risk ontology.

Metrics	
Axiom	2924
Logical axiom count	2119
Class count	693
Object property count	6
Data property count	62
Individual count	44

Figure 8 Ontology Metrics

## Conclusion and future work

This short working paper presents the Financial Industry Operational Risk Ontology (FiORO). This ontology will address several problems in regulatory compliance and reporting for financial services organizations. Briefly, its primary purpose is to enable the systematic identification, assessment, management, mitigation and regulatory compliance reporting of operational risks (OR). As a Master Data Model, FiORO will enable operational risks data from heterogeneous silos to be virtualized and/or persisted in a graph data store for querying using SPARQL, analysis using standard reasoners, and inferencing engines, and visualization using Open Source and/or proprietary tools. It is also well recognized that financial services organisations also require an OR knowledge base in order to ensure that staff are knowledgeable about operational risks. Thus, the proposed ontology aims at facilitating OR information sharing within and across organizational units and business lines.

We have described in brief how this ontology can be used. Despite its early stage of development, it provides a more accurate knowledge base for operational risk reporting—internally and externally. We expressly designed it using a modular approach to serve multiple purposes. It provides the basis for deeper, more complex semantic representation of operational risk knowledge.