ACTIAN™
a division of HCLSoftware

zeenea

EBOOK

What is a ‘Smart’ Data Catalog?

Empower every data user with a data catalog built for modern enterprises.

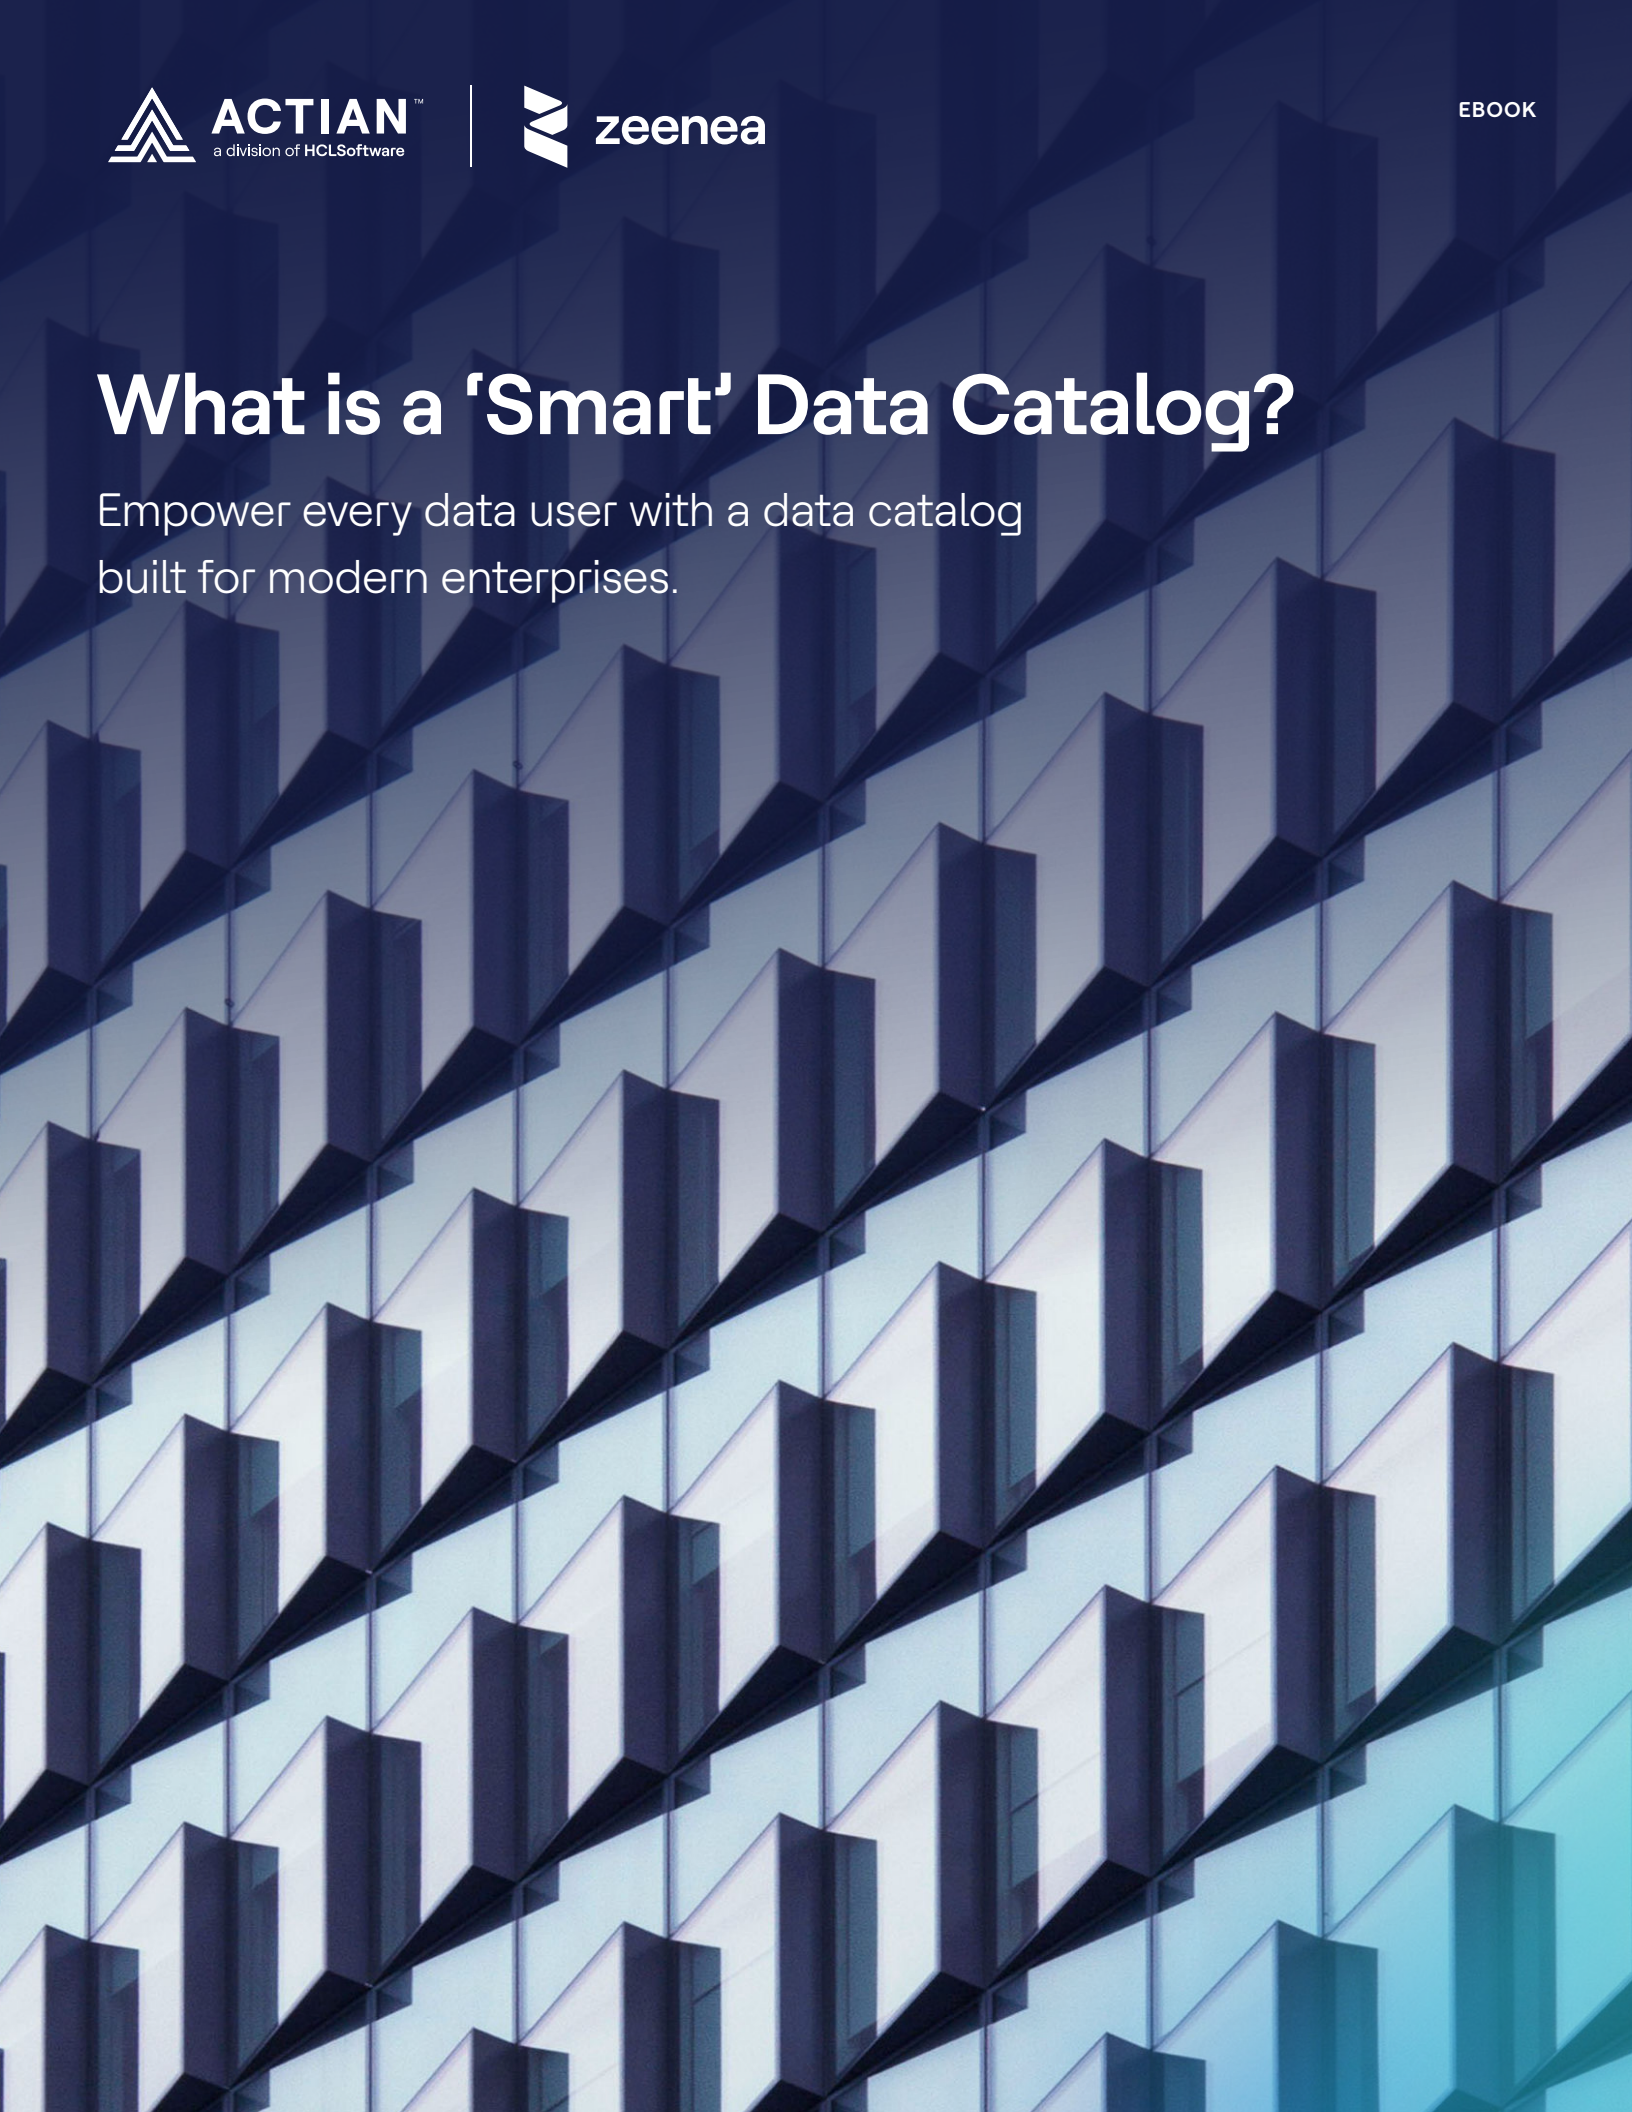

Table of Contents

Executive Summary

The idea of a smart data catalog has been around for several years in metadata management related literature, although it has no official definition. The general consensus is that a modern data catalog must have machine learning and artificial intelligence (AI) to unlock its full potential.

This eBook explains how Actian approaches the idea of a smart data catalog, which cannot be limited to only machine learning capabilities. Before delving into what's meant by "smart," it's important to first describe what a data catalog is. A data catalog is a detailed inventory of all available company data assets, along with the metadata needed to make use of these assets.

A data catalog enables users to efficiently locate the data assets they need, streamlining access for various use cases. It therefore goes beyond a structured data inventory to provide an operational system aimed at accelerating data initiatives. The data catalog should meet the needs of a variety of users, such as analysts, data engineers, risk and conformity professionals, data scientists, product and business managers, and others. In short, the data catalog should cater to the needs of end users.

Data: The Invisible Asset Driving Value

Data fits perfectly with an accounting definition of an asset—a company resource that can be used to produce added value or help the company function properly. The difference is that unlike other assets, data does not show up on the balance sheet.

Enabling Data as a Corporate Asset

The perception of data as a company asset is fairly recent, but it's being increasingly embraced by innovative businesses. In a massively digitalized world, the more successful organizations are those that make the most use of the vast volumes of data available to them. The data can be integrated and analyzed to improve product and service positioning, used to enter new markets, or for myriad other use cases.

Data becomes an asset on one key condition—it must be exploitable. The quest for efficient data exploitation has led to substantial investments by most organizations and touches on many aspects, including company culture, technical architecture, and operations. One aspect at the heart of a data catalog is metadata management.

ACTIAN™ a division of HCLSoftware | zeenea

The Importance of Metadata in Information Exploitation

It is important to keep in mind that company data is essentially a binary magma—zeros and ones—and is incomprehensible to most people. For example, consider frequently used data on computers and mobile devices. On the physical side, there's nothing unusual—the data represents a series of zeros and ones on a hard drive or another storage system. The person who owns the computer or mobile device never actually accesses this raw data.

A series of software solutions gives them the information on the data that enables them or another software to use the data:

- A controller keeps the information relating to the physical location of the bits that constitute an atomic dataset, typically in a file.
- A filing system organizes the datasets in a logical manner and manages crucial information, such as directories, names, or extensions, to make sense of the files. The filing system also manages security, including owners, permissions, date of creation or updates, and more.
- A dedicated file explorer harnesses the information to allow users to consult and understand the content, such as exploring the directory hierarchies, previsualizing files, searching, and linking extensions to applications.

All of the information managed by these different components is called metadata—literally data about data—and is indispensable to make sense of data content in an organization. Human intervention is very rare. Organizations only need a filename and a place to put the information, and metadata is managed automatically.

The challenge comes with scaling. Moving from a file system to an information system is difficult.

The Purpose of the Smart Data Catalog

A smart data catalog's purpose is to efficiently consolidate vast amounts of information. An information system, regardless of its size, contains several dozen systems and applications that store data from a wide variety of sources. These sources can include relational and nonrelational databases, distributed file systems, APIs, cloud solutions, and more, according to specific protocols, formats, and rules.

Each system manages hundreds or even thousands of datasets—usually tables or files—that are themselves made up of dozens of fields or columns. Each dataset and each field feeds into a metamodel—an ensemble of structured metadata—that makes data exploration possible.

An enterprise metamodel is more sophisticated than an individual system's file. It potentially covers a wide range of aspects:

- **Technical metadata.** Which tools to use to access data based on protocols, authorizations, formats, types, and other criteria.
- **Semantic metadata.** What operational information does the dataset contain, and which business rules govern it.
- **Organizational metadata.** Who owns the data, who produces and controls it, and how is it classified.
- **Usage metadata.** Where and how are the data assets used, what is their quality, and how are they monitored.
- **Compliance metadata.** Which internal rules and regulations must be adhered to in order to use the data.

Ultimately, a data catalog harnesses enormous amounts of very diverse information, and its volume will grow exponentially, just as the volume of usable data will grow. This volume of information raises two major problems:

1. How to feed and maintain the volume of information without tripling, or more, the cost of metadata management?
2. How to find the most relevant datasets for any specific use case?

A data catalog must be smart to resolve these issues.
A smart data catalog should have a much wider scope than integrating AI algorithms and should include a range of smart technological and conceptual features.

A data catalog can be smart in these five areas:

1. Metamodeling
2. Data inventory
3. Metadata management
4. The search engine
5. User experience

What is Metadata?

Metadata is what turns binary content into exploitable information. The purpose of the data catalog is to consolidate metadata from all available datasets and present them in the simplest, most straightforward way to data consumers.

Metamodeling in a Smart Data Catalog

At an enterprise scale, the metadata required to harness data assets can be considerable. Besides a narrow sub layer that's mostly technical metadata, the metadata is specific to each organization and sometimes even among different departments within a company. For example, a business analyst won't necessarily seek the same information as a data engineer or a product manager.

It's important to note that a universal and static metamodel cannot be smart. Attempting to create a universal metamodel therefore is not a best practice. Indeed, such a metamodel would have to adapt to a plethora of different situations and will inevitably fall victim to one of these three pitfalls:

1. Excessive simplicity that won't cover all the use cases needed.
2. Excessive levels of abstraction with the potential to adapt to a number of contexts at the cost of arduous and time-consuming training. This is not an ideal situation for an enterprise-wide catalog deployment.
3. Levels of abstraction lacking depth and ultimately leading to multiple concrete concepts borne out of a combination of notions and originating from a variety of different contexts. Many of these concepts will be useless in any specific context, rendering the metamodel needlessly complicated and potentially incomprehensible.

An Organic Approach to a Metamodel

A metamodel is a field of knowledge. The formal structure of a knowledge model is referred to as an ontology. An ontology defines a range of object classes, their attributes, and the relationships between them. In a universal model, the ontology is static—classes, attributes, and relations are predefined, with varying levels of abstraction and complexity (Figure 1).

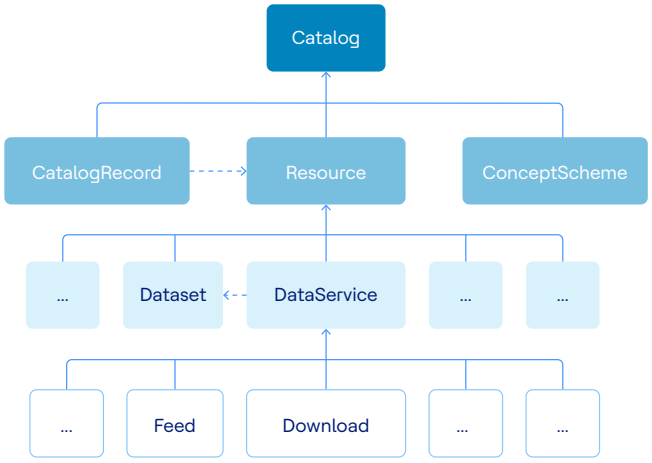

Figure 1: Ontology Example: The Formal Structure of a Knowledge Model

The Actian Zeenea data catalog does not rely on a static ontology but rather on a scalable knowledge graph. The metamodel is therefore simple at the start. There are only a handful of types, representing the different classes of information assets, such as data sources, datasets, fields, and dashboards, each with essential attributes, such as name, description, and contacts (Figure 2).

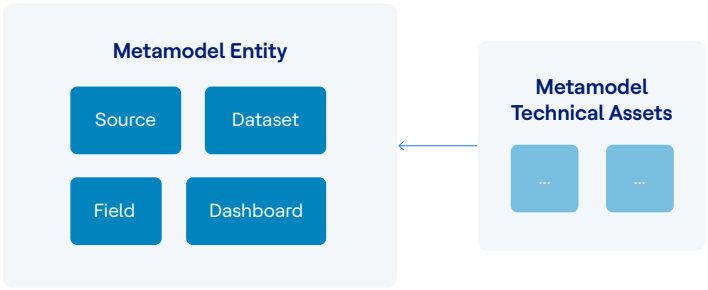

Figure 2: Classes of Information Assets

Adaptability is Key

Smart metamodeling should ensure a metamodel that adapts to any context and can be enriched as use cases or maturity levels develop over time.

ACTIAN™ a division of HCLSoftware | zeenea

This metamodel is fed automatically by the technical metadata extracted from the data sources, which vary depending on the technology being used. For example, the technical metadata of a table in a data warehouse differs from the technical metadata of a file in a data lake.

It's possible to define new object classes, or new attributes to existing classes, and the way in which these different objects relate to each other. The ontology isn't coded "in stone" and it's easy to integrate the evolutions of the metamodel in a manner that allows for trial and error. As a result, progressively expanding the data catalog to new use cases is made more straightforward.

The Actian Zeenea Data Catalog Leverages on a Knowledge Graph

Organic metamodeling in the Actian Zeenea data catalog is the smartest way to handle the ontology issue. The data catalog offers several advantages:

- The metamodel can adapt to each context, often relying on a pre-existing model, integrating the in-house nomenclature and terminology without the need for a long and costly learning curve.
- The metamodel does not need to be fully defined before using the data catalog. Organizations will only need to focus on a few classes of objects and the few necessary attributes to cover initial use cases. Companies can then load the model as catalog adoption progresses over time.
- User feedback can be integrated progressively, improving catalog adoption, and as a result ensuring a strong return on investment for metadata management.

Adding functional attributes to the metamodel can facilitate searching for assets. There are considerable advantages to this metamodeling approach. There is also one major inconvenience—because the metamodel is completely dynamic, it's difficult for the engine to understand the structure and therefore difficult for it to help users feed the catalog and use the data, which are the two core components of a smart data catalog.

Part of the solution, however, relates to the metamodel and the ontology attributes. Usually, metamodel attributes are defined by their technical types such as date, number, chain of characters, and list of values.

With the Actian Zeenea data catalog, these library types do include these technical types, and they also include functional types such as quality levels, confidentiality levels, and more. These functional types enable the engine to better understand the ontology, refine the algorithms, and adapt the representation of the information.

How Data Inventory Supports the Data Catalog

Another way to make a data catalog "smart" is through its inventory. A data catalog is essentially a thorough inventory of information assets that include a lot of metadata that help harness data as efficiently as possible. Setting up a data catalog therefore depends, first of all, on an inventory of assets across the organization.

The Challenges of Automating Data Inventory

A declarative approach to building the inventory is not a best practice, however well thought out it may be. It involves a lot of work to launch and maintain the catalog. In a fast-changing digital landscape, the initial effort quickly becomes redundant.

These systems are managed along with all the metadata required for them to work properly. There's no need to recreate this information manually—organizations just need to connect to the different registries and synchronize the catalog content with the source systems. In theory, this should be straightforward, but putting it into practice is actually rather difficult. The fact is, there is no universal standard to which the different technologies conform to as a means of accessing metadata.

Some systems offer simple and well-documented protocols to access the metadata, while others require more in-depth work. Even with relational databases and other common systems, the standard methods are disappointing.

For instance, the Java Database Connectivity (JDBC) API, which provides connectivity to relational databases for Java, offers standard interfaces to access the metadata. It should therefore be easy to gather this information for all systems that have a driver—basically all technologies being used today.

In reality, this standard API provides only basic information. The system tables that come from these solutions possess rich and complete metadata, and it's often better to consult them directly.

The first and obvious step in creating a smart inventory is to automate it. With few exceptions, enterprise datasets are managed by system specialists and can involve distributed filing systems, enterprise resource planning (ERP) systems, relational databases, software packages, and data warehouses, among other solutions.

The Essential Role of Connectivity to System Sources

A smart connectivity layer is a key part of the smart data catalog. That's why the main characteristics of the Actian Zeenea data catalog are:

- **Proprietary.** The data catalog does not rely on third parties in order to maintain a highly specialized metadata extraction.
- **Distributed.** This helps maximize the reach of the data catalog.
- **Open.** Anyone looking to enrich the catalog can easily develop their own connectors.
- **Universal.** The data catalog can synchronize any source of metadata.

This connectivity not only reads and synchronizes the metadata contained in the source registries, but it can also produce metadata. The production of metadata requires more than simple access to the source system registries. It also requires access to the data itself, which will be analyzed by scanners to enrich the catalog automatically.

The Actian Zeenea data catalog produces two types of metadata:

1. **Statistical analysis** to build a profile of the data, such as value distribution, rate of null values, and top values. The nature of the metadata depends on the native type of the data being analyzed.
2. **Structural analysis** to determine the operational type of specific textual data, such as email, postal address, social security number, and client code. The system is scalable and customizable.

The Inventory Mechanism Must Also Be Smart

The inventory mechanism in the Actian Zeenea data catalog is smart in several ways, beyond boasting a connectivity that can automatically feed the data catalog with the assets contained in the different systems. These ways include:

- Dataset detection relies on an extensive knowledge of the storage structures, particularly in a big data context. For example, an internet of things (IoT) dataset made up of thousands of files of time series measures can be identified as a unique dataset, with the number of files and their location being only metadata.
- The inventory is not integrated in the data catalog by default to prevent the import of technical or temporary datasets that would be of little use, either because the data is unexploitable or because it's duplicated data.
- The selection process for data assets that should be imported into the data catalog benefits from some assistance. That's why a best practice is to identify the most appropriate objects for integration in the data catalog.

Leveraging the Benefits of Metadata Management

The concept of the smart data catalog that's most commonly associated with algorithms, machine learning, and AI falls within the realm of metadata management.

How Is Metadata Management Automated?

Metadata management is the discipline that values metamodel attributes for the inventoried data assets. The required workload is usually proportional to the number of attributes in the metamodel and the number of assets in the data catalog—the volume of metadata used to feed and maintain the data catalog can be enormous.

The role of the smart data catalog is to automate processes as much as possible, or at the very least, help data stewards by enhancing productivity and ensuring greater reliability. A smart connectivity layer enables partial metadata automation, although this is typically limited to a subset of the metamodel—primarily technical metadata.

A complete metamodel, or even a modest one, may have dozens of metadata that cannot be extracted from source system registries because the metadata elements were never recorded in the source systems to begin with. For example, metadata such as table structures, column names, or data types can often be extracted automatically. Other types, like business definitions, data quality rules, or usage policies, are not inherently stored in source systems.

Organizations can take several approaches to solve this issue. The most direct approach consists of identifying patterns in the data catalog to suggest metadata values for new assets.

Using Pattern Recognition and Fingerprinting

A pattern will include all the metadata of an asset and the metadata of its relations with other data assets or other catalog entities. Pattern recognition is typically performed with the help of machine learning algorithms.

The difficulty with implementing this approach is precisely qualifying the information assets in a numerical form in order to feed the algorithms and select the relevant patterns. A simple structural analysis is not enough—two datasets can contain identical data but in different structures.

Relying on the identity of the data isn't efficient either. Two datasets can contain identical information, but with different values. For example, "2024 client invoicing" in one dataset, and "2025 client invoicing" in the other.

To solve this problem, the Actian Zeenea data catalog relies on a technology called fingerprinting (Figure 3).

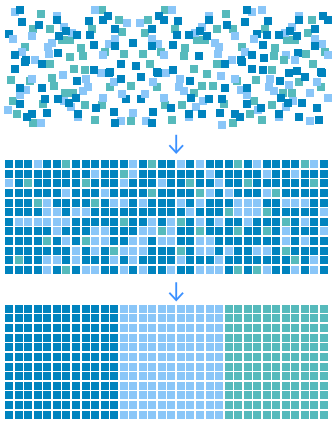

Figure 3: The Actian Zeenea Data Catalog Uses Fingerprinting

Other Embedded Approaches in a Suggestion Engine

While pattern recognition is an efficient approach for suggesting the metadata of a new asset in a data catalog, it rests on an important prerequisite—to recognize a pattern, there must be one to recognize. In other words, this only works if there are a number of assets in the catalog, which is not the case at the start of a project.

It's precisely at these initial phases of a data catalog project that the metadata management load is the highest. It is therefore crucial to include other approaches to help data stewards in these initial phases, when a catalog is more or less empty.

The Zeenea suggestion engine, which provides intelligent algorithms to assist in metadata management, offers other approaches. They include:

- Structural similarity detection. This works when several datasets have identical structures, which is commonplace in layered data lake architectures.
- Fingerprint similarity detection. It does not rely on pattern recognition, but instead on a straightforward Euclidean distance calculation between fingerprints—the fingerprints of two datasets with similar content are likewise very similar.
- Name approximation. This consists of dynamically building a technical names dictionary associated with specific metadata, which works well for specific types of associations. This is the case with the semantic layer, which could, for example, suggest associating a field named "txt_email" with a glossary definition named "Email."

The suggestion engine, which analyzes data catalog content to determine the probable values of the metadata from the assets that have been integrated, is continually updated. New approaches are regularly added. Sometimes they're very simple, and other times much more sophisticated. The Actian Zeenea architecture sees performance improve as the data catalog grows and algorithms are enriched.

The development of the suggestion engine is experimental. Users identify a promising approach, implement it, measure its performance, and then start again. It's a standard approach, but one that raises a major question—how to measure the performance of the data catalog's intelligent algorithms.

A best practice is to use lead time as the main measuring metric for data steward productivity. Lead time is a notion that stems from lean management and measures. In a data catalog context, it's the time that's elapsed between the moment an asset is inventoried and the moment all of its metadata has been valued.

Optimizing a Search Engine to Find Assets Fast

The main objective of a data catalog project is to quickly find relevant data assets. A powerful search engine is needed for efficient exploration.

Given the enormous volumes of data involved in an enterprise catalog, the search engine is the principal mechanism for users to explore the data catalog. The search engine needs to be easy to use, powerful, and most importantly, efficient—results must meet user expectations. Google and Amazon have raised the bar very high in this respect, and the search experiences they offer have become a reference in the field.

A second-to-none search experience should:

- Allow users to write a few words in the search bar, often with the help of a suggestion system that offers frequently associated terms to help narrow down the search.
- Provide near instantaneous results in a specific order, with the most relevant one at the top.
- Enable users to simply add terms to refine the search, or use available filters to cancel out non-relevant results.

Many data cataloging solutions on the market are limited to system indexations, scoring, and filtering. This approach is satisfactory when users have a specific idea of what they're looking for—a high intent search. But the approach can prove disappointing when the search is more exploratory—a low intent search—or when the idea is simply to spontaneously suggest relevant results to a user—no intent.

In short, simple indexation is great for finding information whose characteristics are well known, but falls short when the search is more exploratory. The results often include false positives and the order in which the search comes out is over-represented with exact matches.

Taking a Multidimensional Search Approach

A simple indexation system has limitations and falls short of providing the most relevant results for users. That's why the Actian Zeenea data catalog isolates the search engine in a dedicated module. This moves away from a search engine and data indexation on flat information processing.

The Actian Zeenea approach is similar to Google PageRank's algorithm. PageRank considers several dozen aspects, called features. These features include the density of the relation between different graph objects, such as hypertext links in internet pages, the linguistic treatment of search terms, or the semantic analysis of the knowledge graph.

Several features are integrated into the Actian Zeenea search engine to provide a high level of relevant results, and those features are always evolving. These core features include:

- Standard, flat, indexation of all the attributes of an object, such as name, description, and properties, weighing them in accordance with the type of property.
- A natural language processing (NLP) layer that takes into account near misses, such as typing or spelling errors.
- A semantic analysis layer that relies on knowledge graph processing.
- A personalization layer that relies on a simple user classification according to uses.

Smart Filtering to Contextualize and Limit Search Results

Completing the search engine is a smart filtering system. Smart filtering is something often found on e-commerce websites and provides contextual filters to limit search results.

These filters work by:

- Only properties that help reduce the number of results are offered in the list of filters. Non-discriminating properties do not show up.
- Each filter shows its impact, meaning the number of residual results is affected once the filter has been applied.
- Applying a filter will refresh the list of results instantaneously.

This combination of multi-dimensional search and smart filtering offers a superior search experience. Actian Zeenea's decoupled architecture—the search engine is an autonomous component—allows for continuous exploration of new approaches, with Actian rapidly integrating those that are efficient.

Delivering a Friendly and Intuitive User Experience

A data catalog should be smart in the experience it delivers to different types of users. One of the main challenges with deploying a data catalog is its level of adoption from the users it is meant for—data consumers. The user experience plays a major role in this adoption.

The User Experience Within the Data Catalog

The user experience revolves around identifying key personas, understanding their behaviors and goals, and then designing a seamless, efficient graphic interface that meets their needs. This is the experience offered by the Actian Zeenea data catalog.

Pinning down personas in a data catalog is challenging That's because it's a universal tool that provides value for any organization, regardless of its size, across all sectors and geographies. Rather than attempting to model personas that are hard to define, it's a best practice to focus on data catalog adoption.

These are two user populations that stand out:

- **Metadata producers**. They feed the data catalog and monitor the quality of its content. They are generally called data stewards.
- **Metadata consumers.** They use the data catalog to meet their business needs. They are often referred to as users.

These two groups can overlap. Some data stewards will also be users.

Challenges of Enterprise-Wide Data Catalog Adoption

The full value of a data catalog is achieved with large-scale adoption by a substantial pool of (meta) data consumers, not just the data stewards and data management specialists. This pool is very diverse. It includes data experts, such as data engineers, architects, data analysts, and data scientists, as well as business people, including project managers, business unit managers, and product managers.

The pool also covers compliance and risk managers. More generally, it also includes all operational managers who are likely to leverage data to improve their performance.

Data stewards use the data catalog regularly, so adoption is typically not an issue. They will accept a long learning curve, as long as the solution is user friendly and helps them in their daily activities. The challenge is usually with users.

Data catalog adoption by users is often slow because:

- **Data catalog usage is sporadic.** Users will log on from time to time to obtain very specific answers to specific queries. They rarely have the time or patience to go through a learning curve for a tool they will only use periodically—weeks can go by between catalog usage.
- **Not everyone has the same stance on metadata.** Some users will focus on technical metadata, others will focus heavily on semantic challenges, and still others might be more interested in the organizational and governance aspects.
- **Not everybody will understand the metamodel.** Likewise, not everyone will understand the internal organization of the information within the catalog. Users can quickly feel put off by an avalanche of concepts that feel irrelevant to their day-to-day needs.

The smart data catalog jumps these hurdles to accelerate adoption. For example, Actian Zeenea facilitates data catalog adoption in several ways, including offering a familiar graphic interface that's inspired by e-commerce platforms.

The users' learning curve needs to be as short as possible. In fact, they should be up and running without the need for any training. To make this possible, Actian Zeenea offers two different interfaces, one for data stewards and one for users (Figure 4).

Zeenea Studio. This is the management and monitoring tool for data catalog content—an expert tool designed for data stewards.

Zeenea Explorer. Designed for users, it provides the simplest search and exploration experience possible.

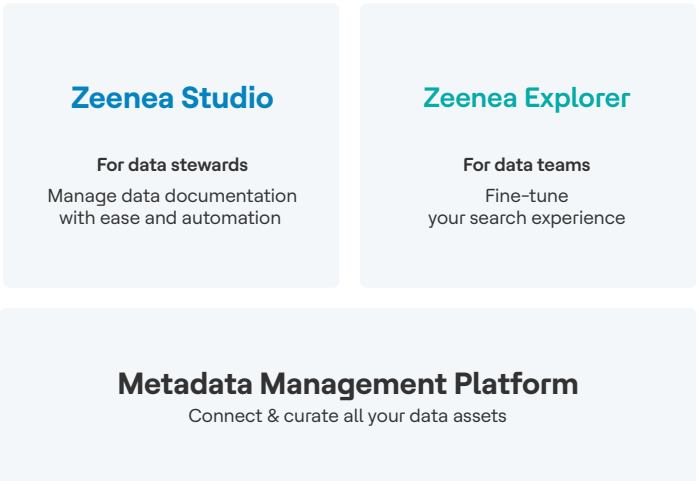

Figure 4: Ontology Example: The Formal Structure of a Knowledge Model

This approach aligns with the user-friendly principles of marketplace solutions. These solutions usually have two applications. The first, a "back office" solution, enables marketplace staff or partners to feed the catalog in the most automated manner possible and control its content to ensure quality.

ACTIAN™ a division of HCLSoftware | zeenea

The second application, for data consumers, usually takes the form of an e-commerce website and enables end users to find data articles or explore the catalog. Zeenea Studio and Zeenea Explorer reflect these two roles.

Zeenea Explorer is the key to large scale data catalog adoption and emulates e-commerce websites for its design:

- The user-friendly design follows in the footsteps of large marketplace sites by offering an efficient search engine, robust exploration capabilities, and a recommendation system that can push a specific set of objects adapted to each user profile.
- As with marketplaces, searching is the main means to access the information. The interface is heavily inspired by mainstream and e-commerce search engines.

A Condensed Learning Curve

With the Actian Zeenea data catalog, the learning curve is considerably shortened as a result of giving users a familiar looking system. An interactive guide feature for new users makes the learning curve even shorter—there's no need for any training.

The Actian Zeenea data catalog ranks information in accordance with the role of the user within the organization. It dynamically adapts the information hierarchy in the data catalog according to user profiles.

This information hierarchy is what differentiates a data catalog from a marketplace type of catalog. The information hierarchy in a marketplace is the same for all users. Photos, descriptions, prices, and delivery information will always be the most important information, followed by technical details, opinions, and delivery conditions.

In a data catalog, on the other hand, the information hierarchy depends on the operational role of the user. For some, the most relevant information in a dataset will be technical, such as location, security, formats, and types. Others will want to know the data semantics and their business lineage. Still others will want to know the processes and controls that drive data production for compliance or operational considerations. The smart data catalog should be able to dynamically adjust the structure of the information to adapt to its different users.

The last challenge is the manner in which information is organized in the data catalog. It's organized in the form of exploration paths by theme, similar to shelving in a marketplace. It is difficult to find a structure that agrees with everyone. Some will explore the catalog along technical lines, such as systems, applications, and technologies. Others will explore it from a more functional perspective, like business domains, and still others from a semantic angle, such as through business glossaries.

The challenge of having everyone agree on a sole universal classification seems insurmountable for most organizations. That's why the smart data catalog should be adaptable and not ask users to understand a classification that doesn't make sense to them.

Ultimately, the user experience is one of the most important success factors for a data catalog. This experience depends on the application designs offered by the catalog, and also by the efficiency and simplicity of the metamodel.

Key Takeaways to Optimize a Smart Data Catalog

One cornerstone of a successful data strategy is effective metadata management across the organization. A smart data catalog is the most effective tool for managing this metadata.

The smart side of the data catalog is not limited to integrating intelligent algorithms. The "smart" must be featured in all aspects of the data catalog:

- The way the metamodel can be designed and added to, altered, or improved as data catalog adoption grows.
- Advanced automation of data asset inventory and the collection of metadata in the systems that host them.
- The ability to assist data stewards in feeding and controlling the content of the data catalog.
- The search engine, which is the simplest and most direct access point for data consumers.
- The user experience, which must take into account the wide variety of personas who will use the catalog.

A truly smart data catalog is more than a repository of metadata—it's a dynamic, adaptable tool that aligns with the needs of diverse users. By prioritizing the user experience, flexible information hierarchies, and intelligent automation, the data catalog ensures that every user, from technical teams to business leaders, can access the information they need in a way that makes sense to them.

In an era where data is a strategic asset, organizations that embrace a smart, user-centric data catalog will be best positioned to drive efficiency, collaboration, and innovation.

About Actian

Actian empowers enterprises to confidently manage and govern data at scale. Organizations trust Actian data management and data intelligence solutions to streamline complex data environments and accelerate the delivery of AI-ready data. Designed to be flexible, Actian solutions integrate seamlessly and perform reliably across on-premises, cloud and hybrid environments. Learn more about Actian, the data division of **HCLSoftware**, at **actian.com**.

Tel +1 512 231 6000 Fax +1.512.231.6010 710 Hesters Crossing Road, Suite 250, Round Rock, TX 78681 **Actian.com**

ACTIAN™ a division of HCLSoftware | zeenea