THE FUTURE OF ENTERPRISE SYSTEMS

Exploring the Integration Revolution with the 3D Enterprise System Platform

Abstract

This whitepaper introduces the 3D Enterprise System Platform, designed to tackle pervasive integration and data aggregation challenges. The platform facilitates the creation of enterprise systems with built-in capabilities for seamless data exchange and aggregation, essential for effective reporting and AI-driven analytics. It enables conglomerates to deploy systems tailored to specific business units, enhancing their unique operational needs. Additionally, it supports supply chains in achieving effortless data interchange. This strategic approach not only streamlines operations within conglomerates but also across their extended supply networks.

Blair Kjenner
bkjenner@method1software.com
www.3d-ess.com

1 EXECUTIVE SUMMARY

In today's fast-paced business environment, the challenges of integrating disparate enterprise systems are a common source of frustration among professionals and a barrier to effective decision-making for management. This whitepaper introduces the 3D Enterprise System Platform (3D Platform), designed from the ground up to address these fundamental issues by reimagining enterprise systems architecture. The platform not only eliminates traditional inefficiencies by structuring systems around business units with shared functional modules but also enhances data accuracy, accessibility, organizational agility, and scalability, setting a new standard in enterprise system management.

Key Features of the 3D Platform:

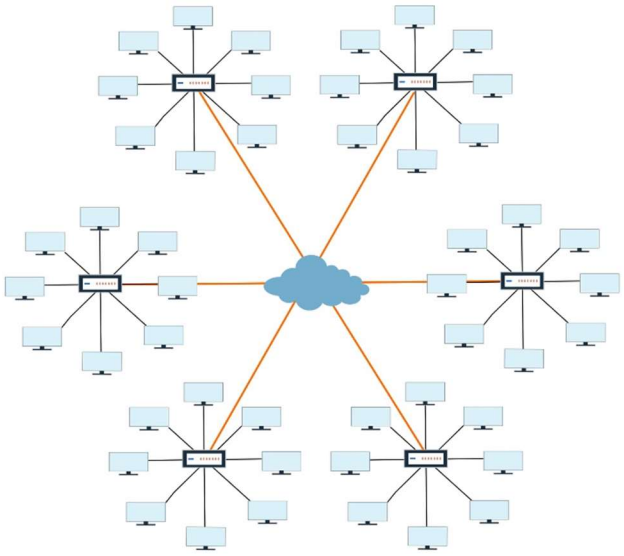

- **Integration by Business Unit:** Prioritizes integration from the initial stages of system implementation to simplify operations across various organizational units.
- **Global Data Exchange:** Employs sophisticated data exchange mechanisms to enable secure, real-time data sharing both within conglomerates and globally, supported by robust governance and compliance frameworks.
- **Advanced Customization and Extensibility:** Supports extensive customization while maintaining core functionalities, allowing organizations to tailor the system to their specific requirements.
- **Ability to Automatically Generate Data Warehouses:** Leverages an integration-first approach, enabling automatic aggregation of data from multiple 3D Enterprise Systems into a continuously updated data warehouse.

Strategic Benefits:

- **Operational Efficiency:** Reduces redundancy and inefficiencies commonly seen in traditional enterprise systems.
- **Cost Reduction:** Lowers total cost of ownership by minimizing repetitive data entry and complex integration efforts.
- **Enhanced Decision-Making:** Provides a consolidated view of enterprise data, facilitating better-informed decisions and quicker response times.

The 3D Platform not only facilitates a seamless transition from legacy systems but also establishes a strategic pathway for data management and operational continuity. It ensures data integrity and connectivity across organizational boundaries, crucial for sectors like healthcare, oil and gas, and government where effective data exchange and collaboration are essential. Tailored for forward-thinkers and innovators, this transformative approach positions the 3D Platform at the forefront of enterprise system evolution. As AI reshapes our future, organizations with deeply integrated systems will fully harness its potential, gaining a competitive edge in the innovation race. Discover how adopting this innovative framework can redefine the landscape of your enterprise solutions.

Contact bkjenner@method1software.com to explore how the 3D Platform can serve as a key catalyst in your strategic transformation.

2 BUSINESS CASE

Creating enterprise systems that can meet the complex needs of today's organizations requires a new approach that is designed to address today's needs for integration and data aggregation. This section discusses the key components of a business plan.

2.1 PROBLEM STATEMENT

2.1.1 How stakeholders experience enterprise systems

Today's business landscape showcases remarkable advancements that have significantly improved productivity, reduced costs, and enhanced customer service. Despite these advancements, there remains substantial potential to further streamline business processes and enhance system integration and connectivity. This section highlights the complexities encountered in environments where numerous systems coexist.

- **Users:** Typically, users interact with multiple systems daily, often facing challenges with poorly integrated tools. This leads to extensive hours spent reconciling data across platforms, frequently resorting to spreadsheets as makeshift solutions. These inefficiencies not only cause data inconsistencies but also lead to significant productivity losses.
- **Management:** Executives regularly confront the challenge of interpreting conflicting data outputs from different systems, complicating strategic decision-making. The time and effort needed to aggregate this data are considerable, often overshadowing potential benefits with increasing system maintenance costs and the risk of investing in new yet potentially ineffective technology solutions.
- **Data Scientists:** Data professionals spend approximately 80% of their time collecting and preparing data from various systems, a task that is continually complicated by evolving data sources. This focus on data preparation detracts from their ability to perform high-value analytical work, thereby limiting the strategic use of AI and data science in organizational decision-making.

The current system landscape not only affects the operational efficiency of these key stakeholders but also limits the potential benefits that streamlined data integration could offer. The 3D Platform is designed to address and resolve these pervasive issues, establishing a new benchmark for efficiency and connectivity in enterprise systems.

2.1.2 Why IT is challenged to address these issues

The traditional approach to enterprise systems, organized around specific business functions—such as HR, payroll, finance, and inventory—works well in isolation but often results in a collection of disconnected systems, each presenting unique integration challenges. These independently developed systems with customized data models meet immediate business needs but create significant barriers to organizational integration, often treated as an afterthought, which amplifies complexity during attempts to connect these disparate systems.

As demands for integrated operations increase, organizations frequently attempt to merge these fragmented systems into a single, comprehensive platform covering multiple functions. However, expanding functionality within these systems increases complexity exponentially, especially in large organizations, making the integration of a fully functional system for a large government entity requiring seamless data sharing across departments particularly daunting.

This challenge extends beyond transactional data to include reference tables, contacts, equipment, projects, land parcels, and other types of master data common across systems. Aggregating this data into a centralized warehouse for AI or analytics is especially challenging under the current model due to inherent difficulties in system integration.

Even reliance on a single ERP vendor does not fully solve these challenges. Traditional ERP systems, organized around functional silos, lack the flexibility to deploy systems by business unit and then automatically share and aggregate data across all instances. These systems often use a mix of composite and surrogate primary keys, complicating the assignment of a consistent primary key to records across systems, regardless of transfer frequency.

Additionally, traditional ERP solutions fall short in record governance, a critical aspect when aggregating data from multiple sources. Determining which system holds the authoritative version of a record can be problematic. For instance, when aggregating equipment data across different departments, it is essential to identify the "golden" record version.

Existing ERP solutions also lack robust mechanisms for defining groups of related data and allowing systems to subscribe to updates on that data. For example, consider land registry data—a record group might include a parcel and its related details, such as encumbrances. An organization subscribing to that parcel's data should receive automatic updates whenever changes occur. Moreover, organizations often need to manage groups of subscribers, such as creating a subscriber group that includes multiple government departments subscribing to shared reference data.

Existing solutions also lack robust audit logging systems capable of tracking changes at a transactional level across multiple records and tables. This capability is vital for detecting changes in subscribed data and automatically distributing updates to systems that rely on those records.

These are just some of the essential features needed to create truly networked enterprise systems. Attempting to address this level of integration with a traditional ERP solution typically results in an overextended system where functionality is packed into a single solution, often using row-level security to separate departmental data. While this approach may seem comprehensive, it introduces significant challenges when customizing the system to meet diverse departmental needs. The growing complexity drives up costs and increases the risk of system failure, eventually causing the system to become unwieldy.

In today's landscape, organizations need solutions that can handle this level of complexity and go beyond, addressing integration across entire supply chains. Industries within a supply chain share vast amounts of data, and integrating disparate systems across these sectors is a formidable challenge.

The need for better integration is growing rapidly as expectations for AI increase. The true potential of AI can only be realized if organizations can aggregate data from all their systems in a cohesive manner. However, with current systems, this vision remains out of reach unless we adopt a different approach.

Imagine a world where an organization can automatically aggregate data from all systems into a single data warehouse. This capability would unlock the full power of AI, enabling smarter decision-making, real-time insights, and a competitive edge. Businesses operating with fragmented systems and limited data connectivity are increasingly at a disadvantage. The future belongs to organizations that eliminate integration barriers and enable AI to function across all aspects of their operations. Those that adapt will lead their industries, leaving behind those struggling with outdated, disconnected systems.

2.1.3 Solutions We've Tried

Over the years, various solutions have been proposed to tackle the challenges of integrating enterprise systems. Although each presents specific advantages, none fully overcomes the core issues that lead to fragmented data and complex integrations. A closer examination of some commonly suggested approaches reveals their limitations:

1. **Software as a Service (SaaS):** While SaaS offers scalability and ease of integration, its reliance on cloud-based services often leads to reduced control over data. The limited access to underlying databases can hinder comprehensive data integration across systems, posing challenges in thorough data analysis and system customization.
2. **Electronic Data Interchange (EDI):** EDI supports automated data exchange but is limited by varying data formats across systems, leading to inconsistent data structures and incomplete integration. It tends to address symptoms rather than the root causes of system fragmentation.
3. **Powerful Software Development Tools:** New development tools accelerate system creation but typically promote the proliferation of isolated systems, each with its own data models and interfaces. This increases complexity within an organization's IT infrastructure and raises long-term maintenance costs.

4. **Data Warehouses:** Although designed to consolidate data for analysis, data warehouses depend on consistent, clean data inputs. They often replicate existing integration challenges, as fragmented source systems complicate the aggregation process, diminishing the effectiveness of data analysis.
5. **Master Data Management (MDM):** MDM attempts to synchronize core data across systems but struggles with complex, frequently updated multidimensional data across diverse systems. It fails to address the foundational issues of system integration and data governance.
6. **Microservices Architecture:** Offering modularity and scalability, microservices do not inherently address data standardization challenges or ensure consistent record governance across services, especially when different services use disparate data storage solutions.

While these approaches each contribute to a holistic solution, they fall short for organizations aiming to achieve true, seamless integration at scale. To transcend these limitations, a platform that is fundamentally designed with integration, data governance, and scalability at its core is crucial. The 3D Platform introduces a novel approach by targeting the root causes of fragmentation and establishing a unified foundation for the enterprise systems of the future.

2.2 MISSION

To redefine enterprise system management by delivering a platform that ensures effortless data exchange and aggregation, enabling organizations to seamlessly integrate systems and harness the power of AI for superior decision-making.

2.3 KEY OBJECTIVES

To solidify the transformative impact of the 3D Platform, we are committed to achieving the following key objectives:

1. **Core Data Models**: Develop and refine robust core data models that serve as the foundation for all systems. These models are designed to be easily extendable to accommodate the unique needs of various business units while maintaining global consistency.
2. **Extensibility**: The platform allows for the straightforward extension of core models, ensuring that each business unit can tailor the system to fit its specific needs without compromising the integrity or interoperability of the data.
3. **Unique Identification**: Implement a high-performance, universally consistent primary key strategy that ensures each record can be uniquely identified across all systems, facilitating seamless data integration and retrieval.
4. **Secure Data Exchange**: Enable secure and efficient data exchange within conglomerates and across global networks, ensuring that data transfers are both rapid and protected against unauthorized access.

5. **Master Data Exchange**: Facilitate comprehensive master data management that supports the synchronization of essential data across the entire supply chain, enhancing the accuracy and availability of key information.
6. **Subscription Mechanisms**: Provide advanced subscription mechanisms that allow systems to subscribe to updates of specific data groups and tables, ensuring that all entities are working with the most current data available.
7. **Subscriber Management**: Support efficient management of subscriber groups, simplifying the administration of data access and distribution across diverse systems and user bases.
8. **Comprehensive Audit Logging**: Maintain detailed records of all changes to both temporal and non-temporal data, ensuring transparency and facilitating compliance with regulatory requirements.
9. **Automated Data Transfers**: Automate the process of data transfers from the data manager to subscribing systems, streamlining the update process and reducing manual intervention.
10. **Flexible Import Mechanisms**: Offer customizable import mechanisms that allow users to control how updates are applied within their systems, enhancing adaptability to specific operational workflows.
11. **Flipping the Axis**: Implement systems by business unit, allowing for shared functional components across units, which optimizes resource utilization and system coherence.
12. **Versioning Support**: Provide robust versioning capabilities for development, UAT (User Acceptance Testing), and production environments, ensuring that all system modifications are traceable and manageable.
13. **Automating Aggregation**: Establish mechanisms to automatically aggregate data from multiple business unit systems into a single data warehouse, supporting comprehensive reporting and advanced analytical capabilities.

By fulfilling these objectives, the 3D Platform is uniquely positioned to transform how organizations create, manage, and integrate their enterprise systems, providing a competitive edge in the rapidly evolving digital landscape.

2.4 VISION

The vision for 3D Enterprise Systems is to fundamentally transform how organizations create and manage enterprise systems. Instead of building systems by business function, the 3D Platform introduces a paradigm shift where systems are implemented by business unit, with common modules shared across units within an organization. These modules can either be unique to a conglomerate or universally available for any organization leveraging the 3D Platform.

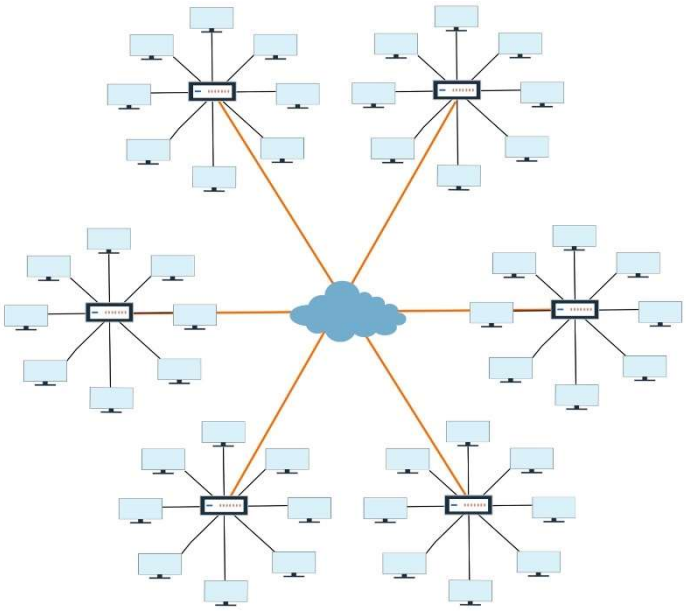

In this new approach, businesses will only ever enter data once, and it will automatically flow across all relevant systems—both within the organization and to external entities. Data exchange will be seamless, eliminating the need for manual reconciliation and data aggregation efforts. The days of reconciling mismatched data between systems using spreadsheets will be a thing of the past.

For conglomerates, the ability to automatically generate comprehensive data warehouses from their business unit systems will revolutionize their ability to leverage AI. They will leap ahead of competitors by gaining insights that drive smarter decisions, all while significantly reducing integration complexities.

The 3D Platform enables us to tackle challenges that are currently insurmountable. For instance, we could deploy systems by business unit across a conglomerate as large as a government, integrating departments from healthcare to education within a unified architecture. This integration could extend even further, connecting government systems with external entities like oil and gas businesses, enabling seamless data exchange wherever significant data sharing is required.

By embracing this vision, organizations will be equipped to gradually overcome the inefficiencies of legacy systems while smoothly integrating them into the 3D Platform as part of a phased transformation. This approach unlocks a future of true integration, innovation, and growth.

2.5 OPPORTUNITY

There is a significant opportunity for an innovative solution like the 3D Platform. Our integration challenges will persist until we reconsider how we create enterprise software for organizations. Currently, ES vendors have not designed their systems from the ground up to be networkable. This

oversight presents a substantial market gap that can be addressed by enhancing existing ES systems or constructing new ones based on our model. Implementing this approach provides organizations with a unique differentiator that no other ESs possess, positioning them ahead in the market.

We can no longer afford to develop systems in isolation and then struggle with integration, nor is it feasible to construct monolithic applications, as they inevitably succumb to their complexity. This is the first solution that adopts a holistic approach to these issues, presenting an opportunity for thought leaders who recognize the problem and are prepared to implement the 3D Platform.

There is a vast amount of funding allocated annually to enterprise software. According to Gartner, IT spending is on the rise, increasing by 5% annually, with organizations worldwide projected to invest over a trillion dollars annually in software alone. This significant market growth underscores the demand for a networkable ES system capable of addressing currently insurmountable challenges.

Table 1. Worldwide IT Spending Forecast (Millions of U.S. Dollars)

	2022 Spending	2022 Growth (%)	2023 Spending	2023 Growth (%)	2024 Spending	2024 Growth (%)
Data Center Systems	216,095	13.7	224,123	3.7	237,790	6.1
Devices	717,048	-10.7	684,342	-4.6	759,331	11.0
Software	793,839	8.8	891,386	12.3	1,007,769	13.1
IT Services	1,250,224	3.5	1,364,106	9.1	1,502,759	10.2
Communications Services	1,424,603	-1.8	1,479,671	3.9	1,536,156	3.8
Overall IT	**4,401,809**	**0.5**	**4,643,628**	**5.5**	**5,043,805**	**8.6**

Source - Gartner (April 2023)

This transformative model extends beyond traditional industries and holds immense potential for sectors like healthcare, oil and gas, and government where data exchange and collaboration are critical for delivering efficient and effective services.

3 Key Principles

To fully realize the vision of the 3D platform, a fundamental rethinking of how enterprise systems are constructed was necessary. This led to the establishment of five guiding data principles designed to create enterprise systems that have an inherent capability to exchange data and aggregate it for reporting and AI.

- Core Data Models
- Globally Unique Primary Keys
- Record Governance
- Data Transfer
- Data Driven

Each of these principles is discussed in further detail below.

3.1 CORE DATA MODELS

Core data models are foundational components within any organization, encompassing essential data categories like contacts, finances, activities, assets, and contracts. These models adhere to a 'less is more' philosophy, focusing on attributes commonly needed across various systems. For example, typical contact attributes include name, gender, and birth date—standards that remain constant across decades and systems. Financial models similarly revolve around universal elements like journal entries, chart of accounts numbers, and transaction IDs. The 3D Platform integrates these core data models, creating a comprehensive framework that can be customized to meet the specific needs of any business.

Examples of Core Models

- **Contact Registry:** Manages detailed contact information integrating with all core modules. It supports diverse data like postal addresses, email addresses, and phone numbers. The registry accommodates various subtypes (e.g., employees, customers, vendors) and features like alternate names and billing accounts, ensuring data consistency over time.
- **Activity Registry:** Centralizes tracking of activities across systems such as CRM and billing. It supports hierarchical project organization and subtype-specific details, like billing statuses for activities.
- **Equipment Registry:** Facilitates comprehensive equipment management, tracking everything from subtype details (e.g., vehicle attributes like 2WD/4WD) to equipment location history and classification.
- **Human Resources Module:** Manages employee data and organizational hierarchies, linking activities and financial transactions relevant to specific positions.
- **Financial Module:** Ensures seamless connectivity of financial transactions across all data points, supporting detailed financial analysis and efficient management of budgets and transactions.

Features Built into Core Models

Core models come fully equipped with essential features such as temporal data management, audit logging, data exchange, and record governance. These features enable development teams to start with a comprehensive model that not only handles common data types but also includes robust functionalities essential for creating networkable and scalable solutions.

Core Model Data

Core models not only structure essential data but also prepopulate common reference tables (e.g., GL Transaction Status, Address Type) with standard values. Organizations can then tailor this baseline data to specific operational needs.

Benefits

- **Standardization and Static Nature**: The core models serve as a standardized foundation, much like the attributes describing a financial transaction that remain consistent across systems over the years. This stability allows for the development and sharing of source code across different systems.
- **Communication Enablement**: By standardizing data types like contacts and finances, these core models ensure consistency and interoperability between disparate systems, simplifying data sharing.
- **Flexibility and Extensibility**: While the models provide a standard framework, they are designed with the flexibility to be expanded and customized, meeting unique organizational requirements without compromising the integrity of the data.
- **Embedded Features**: Features such as change history (audit logging) and data exchange are integrated into the core models, enhancing data integrity and tracking modifications.
- **Populated with Core Data**: These models come pre-populated with essential data, such as commonly used reference table values, allowing organizations to deploy the models immediately.
- **Inheritance Capability**: The core models are built to be inherited and extended. For instance, the Oil and Gas industry model would inherit the general core model and then integrate specific attributes relevant to its domain.

3.2 GLOBALLY UNIQUE PRIMARY KEYS

The architecture utilizes a globally unique eight-byte key for primary keys, comprising a system ID and a record ID—an incremented integer. This design is applied universally across all data types, including master, transaction, and system data.

Benefits

- **Universality and Uniqueness**: Each record is assigned a primary key that is not only unique within its system but across all systems, greatly simplifying data identification and exchange.
- **Compactness and Efficiency**: The size of the primary key, significantly smaller than a GUID, enhances the efficiency of data storage and retrieval operations.
- **Simplicity for Utilities**: Having a consistent primary/foreign key format across all tables has made it possible to create the 3D platform.
- **Ease of Reference**: The straightforward format (e.g., 1-123) makes these keys easy to use and reference, unlike the cumbersome strings of a GUID.
- **Optimized Data Storage**: The sequential nature of the record IDs allows for efficient, orderly data storage, improving database performance during large-scale data operations.

3.3 RECORD GOVERNANCE

By nature, systems are going to have duplicate copies of the same record. If only one system governed all common data then it wouldn't be a problem. But that is not real world. Take oil well data that is managed by one company but eventually another company buys the well and they manage it. Or a situation where a CPA member resides in one province but then moves to a different province. Or a GL Account structure which is created by a central organization, inherited and augmented by an industry model then inherited and augmented by an oil company. Record governance ensures we always know who owns the gold version of a record. Governance can transfer from one system to the next. By default, when an organization creates a record it automatically governs it. Record governance is critical when it comes to aggregating data from many different systems into a data warehouse for reporting and AI.
Record governance is crucial for maintaining the integrity and trustworthiness of data across systems. It ensures that there is always a clear indication of which system holds the authoritative version of any record, especially vital when aggregating data from multiple sources.

Benefits

- **Assurance of Data Integrity**: Knowing which system's data to trust ensures that the most accurate and current data is used for data aggregation.
- **Flexibility and Control**: Governance can be dynamically managed and transferred between systems, adapting to changes such as business acquisitions or structural reorganizations.
- **Efficiency and Performance**: With the primary keys indicating record origins, governance can be maintained with minimal overhead, preserving system performance while ensuring accurate data management.

3.4 DATA TRANSFER

The Data Transfer principle within the 3D platform ensures robust and efficient exchanges between systems, facilitating seamless management of data flows. This principle states that features need

to be embedded within enterprise systems to ensure that we can define record groups, subscriber groups, and detailed subscriptions to orchestrate and optimize the transfer of data across different systems. We also need mechanisms for automatically detecting changes to data and broadcasting it to subscriber systems.

Benefits

- **Integration Simplification**: Integrations are streamlined through parameter-driven configurations, reducing the need for coding and accelerating deployment times.
- **Master Data Management:** By standardizing master data across systems, the data transfer process simplifies the complexities associated with large-scale data management.
- **Transactional Integrity:** Group updates ensure that changes are processed collectively, maintaining data integrity throughout the transaction lifecycle.
- **Customization and Control:** While the process is designed to be automated, it also offers the flexibility for manual oversight, allowing organizations to tailor the integration to their specific needs.

3.5 SUMMARY

By adhering to these principles, the 3D Enterprise System Platform is meticulously engineered to manage the complexities of contemporary data environments, emphasizing scalability, seamless integration, and stringent data governance. This strategic approach not only meets current operational demands but also positions organizations to capitalize on future technological and data-driven advancements, setting a new benchmark in enterprise system management.

3.6 DATA DRIVEN

3.6.1 Introduction

The central data dictionary is pivotal to the 3D Enterprise System (3D ES), marking a transition from static data model patterns to dynamic, metadata-driven system architectures.

3.6.2 Importance of a Central Data Dictionary

The data dictionary serves as an essential repository for metadata, detailing the purpose, default values, constraints, and relationships of database tables and columns. It underpins various system utilities, such as database creation, data conversions, generation of views, audit logging and data exchange. In the future, front-end systems will connect to the data dictionary to achieve dynamic system configuration.

3.6.3 Reflecting on Industry Practices

While advocating for structured data management, it's crucial that we practice what we preach. Often, IT systems embed crucial data directly in code, contrary to the advised method of dynamic

management through databases. This can lead to less adaptable and maintainable systems, highlighting a gap in industry practices. This observation urges a re-evaluation towards true dynamic, database-driven architectures, enhancing system consistency, maintainability, and scalability.

3.6.4 Extending Core Model Principles to System Data

Just as we use core models and primary keys for business data like contact information, we extend these principles to system data, including the data dictionary itself. We will define core data structures around metadata and employ primary keys for each record within the data dictionary. This approach not only ensures that each piece of metadata is uniquely identifiable and traceable back to its system of origin but also standardizes the handling of this metadata across various systems. By treating metadata with the same rigor as business data, we ensure that our systems are not only robust but also adhere to the highest standards of data governance and integrity.

3.6.5 Summary

This principle advocates for a consistent, standardized approach not only in managing business data but also system data through central data dictionaries. This unified approach is crucial for modern enterprises aiming to maintain robust data governance and seamless adaptability, setting a new benchmark in enterprise system management and aligning our practices with the principles we champion.

4 TECHNICAL DESCRIPTION

This section delves into the technical architecture that underpins our visionary approach. Initially developed with PostgreSQL, the architecture is designed with flexibility in mind. While PostgreSQL serves as the first database environment, the platform is not confined to it alone. The platform could be transformed into other major database environments, such as Oracle and Microsoft SQL Server, to ensure broader compatibility and integration capabilities.

The technical architecture is further validated through an extensive proof of concept (POC) environment that includes 16 different enterprise systems, each populated with scrambled data from live systems. A notable addition to the POC environment is an Oil and Gas company utilizing the PPDM Oil and Gas data model, which features 2000 tables. This inclusion serves as a rigorous test for the 3D Platform, particularly because it challenges the system with complex, temporal data models. Looking ahead, we plan to incorporate additional industry models into the POC environment, such as Healthcare. This expansion will demonstrate how the platform can be extended to support diverse industry models, showcasing its versatility and scalability.

4.1 DATA EXCHANGE MECHANISM

Data exchange is a robust mechanism that facilitates seamless sharing of data between publishing and subscribing systems. This process is centered around record groups, subscriber groups, and subscriptions. Record groups define the specific data to be shared, which may include multiple levels of related child records. Subscriber groups consist of systems with a shared interest in the subscribed data, while subscriptions govern which record groups a subscriber group is entitled to access.

When the export process is initiated, it scans for any changes in subscription data and audit logs, writing the updated information to a data server in JSON format. On the receiving end, each system's import process checks the data server for relevant updates that need to be applied to its database. Users are then able to review and approve the data changes queued for import. The import process intelligently handles differences between the local system's data structure and the incoming data sent by the publisher.

4.1.1 Key Features of the Data Exchange Mechanism

While there are many detailed features and controls involved in automated data exchange, the key high-level considerations include:

- **Structural Flexibility:** Data exchange accommodates scenarios where either the governor or the subscriber has altered the table structures or associated rules. For example, if the governing system adds a new column to a table, subscriber systems can continue to receive updates without interruption, even if they haven't yet updated their structure to include the new column.
- **Support for Temporal and Non-Temporal Data:** The mechanism supports subscriptions for both time-sensitive (temporal) and static (non-temporal) data, ensuring data consistency across varied scenarios.
- **Column Redaction Options:** Redacting sensitive information is straightforward with options to automatically mask, substitute, or null out specific columns as needed.
- **Triggering of Data Exchange:** The platform triggers data exchange whenever a record or a related group of records is updated. For instance, a record group for a contract could include a contract header along with associated terms and participant details.
- **Data Preview During Import:** When importing data, users can preview changes, seeing before-and-after snapshots of all modified values, which helps them assess the impact of updates before approval.
- **Visual Indicators for Pending Updates:** When users browse data in their system, the platform provides front-end cues highlighting records with pending updates from the governing system.
- **Subscriber Group Management:** The platform simplifies subscription management by organizing subscribers into groups, reducing complexity. For instance, a new regional outlet added to a regional subscriber group automatically inherits all relevant subscriptions without additional setup.

- **Master Data Exchange Focus:** The primary focus is on master data exchange. For transactional data, a different process is used. For example, service providers might exchange transactional updates with customers through a custom interface.
- **Data Integrity Management:** The import process detects data integrity violations (e.g., foreign key, null constraints) and allows users to either resolve these issues before processing or ignore the update and store comments explaining the exception.
- **Automation vs. Manual Approval:** The import process can be fully automated or require manual review, depending on predefined rules based on the source system and the type of data being imported.
- **Workflow Integration:** The data exchange mechanism can interact with a workflow engine, triggering related processes. For example, a government department subscribing to citizen registry data could automatically initiate workflows when a citizen changes address.

4.1.2 Inheritance

The platform also supports hierarchical inheritance. For example, a GL account structure could be defined by an industry body (e.g., CPA) and then inherited by industry-specific models, such as oil and gas, which could further extend and customize it. This hierarchical setup can extend down multiple levels, with each inheriting and adapting shared data as needed.

4.1.3 Subnets for Enhanced Data Sharing

Each system operates within a subnet server. Small organizations typically run a single system on their subnet, while conglomerates may operate multiple systems. Data within a subnet is shared instantaneously among systems, enhancing performance and enabling secure data sharing within a conglomerate. Subnets can communicate with other subnets via a global master data server.

4.1.4 Stub systems for exchanging data with legacy systems

The platform uses stub systems to bridge gaps between conforming and legacy systems. Once an integration is developed with a legacy, it becomes reusable across multiple organizations, simplifying future integration projects and enhancing the platform's interoperability.

4.1.5 Custom Data Exchange

While the built-in data exchange mechanism covers around 90% of integration needs, custom data exchanges can be programmed to address specialized requirements. Scenarios requiring high-performance data sharing or secure data exchange between systems are easily accommodated through custom integrations.

4.1.6 Accommodating Monolithic Systems

The platform also supports monolithic systems for conglomerates. Sometimes organizations have large enterprise systems with high volumes of data and transactions and are regionalized to further complicate the situation.

The 3D Platform allows these systems to be broken into smaller systems that are all continuously exchanging data. Each of the systems will have its own system ID and will generate data according to that system ID.

A manual process can be set up to instantaneously exchange data between the systems. It can even have flexibility to allow data to be modified to meet unique needs and can limit the data exchange to only tables and data that are of interest to be exchanged.

Furthermore, it will be possible to aggregate all data from these systems into a data warehouse for reporting.

4.1.7 Data Exchange Using QR Codes

Data exchange can also be achieved manually using QR codes (or equivalent technology). QR codes printed on documents contain the core data for a record, including three components: Dictionary Table ID, Row ID, and Header Data. This allows the unique identification of any record within any system, regardless of where it was created.

For example, a QR code on an oil well could be scanned by a service provider to instantly access well details. Any recorded service data would then be linked to the correct well when transmitted to the oil company, streamlining the entire process.

4.1.8 One-Time Secure Subscriptions

In critical situations, such as healthcare emergencies, one-time secure subscriptions allow for rapid access to essential data without copying detailed records to the subscriber's system. For instance, a healthcare provider could authenticate through a patient's device to access vital medical records, receiving both core patient data and a temporary graphical image of more detailed health information.

4.1.9 Benefits of the Data Exchange Mechanism

- **Simplified Integrations:** Parameter-driven configurations replace custom coding, accelerating integration setup and reducing costs.
- **Custom Integrations:** While automated, the system is flexible enough to accommodate custom integration tools for specialized scenarios.
- **Master Data Management:** The platform streamlines master data management by ensuring consistent data across systems.
- **Support for Transactional Changes:** Updates are processed as grouped transactions, reducing the risk of incomplete changes.
- **Controlled Import Processes:** Imports can be tailored for full automation or manual review, depending on organizational needs.
- **Error Reduction and Consistency:** Automated data exchange ensures consistent data across systems, minimizing errors and inconsistencies.

- **Improved Efficiency:** Quick and efficient data transfers reduce manual entry and streamline operations.
- **Support for Monolithic Systems:** Allows monolithic systems to be broken into smaller systems that communicate.
- **Enhanced Collaboration:** Teams across systems can work together more effectively, sharing data seamlessly.
- **Scalable Design:** The platform scales to handle growing data volumes as organizations expand, future-proofing data management.

The 3D Platform's data exchange mechanism is a transformative solution for organizations seeking seamless, secure, and scalable integration across diverse systems.

4.2 AUDIT LOGGING MODULE

The Audit Logging module is a critical component of the 3D Platform, enabling the precise tracking and management of data changes across systems. By logging every change at a transactional level, it supports data integrity, facilitates seamless data exchange, and offers extensive audit capabilities.

4.2.1 Overview

The Audit Logging module is designed to:

- Track changes across all tables, including both temporal (time-sensitive) and non-temporal data.
- Ensure transaction-based logging, capturing every modification in a coordinated manner.
- Centralize audit records for consistent and efficient data management.

4.2.2 Key Features of Audit Logging Module

1. **Invaluable to Users:** It is invaluable for users to view a succinct log of the changes to a record, who made them, and through what mechanism. This helps them quickly assess how a record got to its current state.
2. **Centralized and Efficient:** The module maintains all audit data centrally, minimizing performance impact while ensuring real-time tracking of master data changes.
3. **Configurable Parameters:** Administrators can define which tables, columns, and events trigger audit logs. This flexibility helps avoid unnecessary logging for non-essential changes (e.g., current balance updates).
4. **Transactional:** The system tracks all changes at the transaction level, capturing every modification across related records and tables as part of a single transaction.
5. **Detailed Audit Views:** Users can view audit logs from multiple perspectives:
    - **By Record:** Drill down into the complete history of changes for any record.
    - **By Command/Menu Option:** Track changes triggered by specific user actions.
    - **By Column:** Analyze changes to specific fields across the system.

    - **By User:** Review all changes made by a specific user, even tracking commands executed without data modification.
6. **Undo Capability:** Authorized users can reverse complex changes, restoring affected records to their previous state while maintaining an audit trail for the undo action.
7. **Selective Logging for Performance:** Audit logging is primarily used for master data and excludes high-volume transactional data (e.g., journal entries) to avoid performance bottlenecks.
8. **Integrated with Data Exchange:** Audit logs are the foundation for broadcasting updates to subscribing systems, allowing changes to be previewed, applied, or ignored with full traceability.
9. **Data for AI and Machine Learning:** The audit data is a valuable resource for analytics, offering insights into change patterns, user behaviors, and data quality over time.

4.2.3 Benefits

- **Enhanced Data Governance:** Empower more users to manage master data with confidence, knowing that every change is tracked and accountable.
- **Streamlined Troubleshooting:** Support teams can quickly diagnose issues by retracing user actions and reviewing detailed audit trails.
- **Reduced Development Costs:** By standardizing audit logging across the platform, development teams avoid the overhead of building custom solutions for each system.
- **Optimized User Experience:** The module's user-friendly interfaces and robust query options make it easy to access and understand the full history of any record.

4.2.4 Conclusion

The Audit Logging module is a cornerstone of the 3D Platform. It provides the transparency, accountability, and efficiency needed to maintain data integrity across interconnected systems, supporting the seamless data exchange capabilities that define the 3D Platform vision.

4.3 RECORD GOVERNANCE

In order to achieve the record governance principle, the record governance module was developed. This module ensures that users always know which system holds the authoritative version of a given master record. The Record Governance module provides the mechanisms necessary to manage and track which system is responsible for maintaining the "gold standard" version of a record.

4.3.1 Overview

By default, the system that creates a record automatically governs it. Conversely, when a system imports a record from another system, it does not assume governance over that record. Governance is typically implicit, relying on the System ID portion of the 3D primary key, so explicit

tracking is only necessary for exceptions, such as when record governance is transferred between systems.

The Record Governance module focuses on tracking these exceptions by generating governance records when the default governance rules are overridden. For instance, if a system transfers governance of a master record to another system, governance records are created in both the old and new governing systems, as well as on the master data server.

Governance can be applied to individual records or groups of related records (e.g., a contact and its associated address).

4.3.2 Functional Components

1. **Record Governance Transfer:** This function allows record governance to be transferred from one system to another. When invoked, it:
    - Creates a pending governance override record in the originating system.
    - Creates a pending transfer update on the master data server.
    - Exports the governance transfer record to the receiving system.
    - Finalizes the transfer once the receiving system accepts it, updating the status on both the originating system and the master data server.
2. **Record Governance Revoke:** If the recipient system has not yet accepted the transfer, the originating system can revoke it. This triggers a governance revoke transaction on the master data server and updates the override record on the recipient system.

4.3.3 User Interface Considerations

The platform allows users to:

- Inquire about the governance status of any record, identifying which system currently holds governance.
- View the governance transfer history for any record, including details on transfers initiated, revoked, or approved.

4.3.4 Summary

The Record Governance module ensures that there is always a clear and accurate understanding of which system governs any given record, preventing conflicts and supporting seamless data exchange. The module also facilitates controlled transfers of governance, allowing systems to maintain data integrity while supporting complex data-sharing scenarios across different environments. By effectively managing record governance, the 3D Platform provides the foundation needed for reliable and efficient data management at scale.

4.4 TEMPORAL DATA MANAGEMENT

Most enterprise systems will only show you how your data looks today. With the 3D Platform, you can specify a query date, and all data will be returned as if it were that day. This adds a new dimension to your data that is critical for queries but also critical for AI to analyze how your data is changing over time. The temporal features embedded within the platform specifically address the time dimension, enabling organizations to manage data that evolves over time with precision and consistency.

4.4.1 Examples of Temporal Data

Temporal data is essential for various business operations. Here are some key examples:

- **Payroll Data:** Compensation parameters for staff change over time, such as roles, rates, and rules.
- **Property Tax:** Tax assessments depend on property values and applicable rates that vary annually.
- **Health Benefits:** Benefits are tied to enrollment plans and the limits within those plans, which change over time.
- **Oil and Gas Data:** Temporal data is crucial in tracking measurement parameters like plate orifice sizes that impact volume calculations.

4.4.2 Key Features of Temporal Data

A system without temporal capabilities can only capture the present state of data. In contrast, the 3D Platform's temporal capabilities allow for historical data to be maintained and updated based on effective dates. For example, corrections to past payroll periods can be applied retroactively without affecting the current data.

4.4.3 Temporal Resolutions

Temporal data in the 3D Platform is captured at various resolutions—year, month, day, or even second—depending on the system's requirements. The resolution determines how temporal segments are created and managed. For instance, in a property tax system with yearly resolution, changes within the same year are consolidated into a single segment.

4.4.4 Conclusion

The 3D Platform's temporal data management is essential for creating systems that can adapt to changes over time. By distinguishing between temporal and non-temporal data and standardizing temporal methods across all systems, the platform accelerates development and enables seamless data exchange—even for temporal data. This consistent approach ensures that organizations can manage their evolving data needs while maintaining high data integrity across interconnected systems.

4.5 DATA WAREHOUSE GENERATION (DWG)

80%[1] of the effort related to artificial intelligence involves aggregating and cleansing data. What if data scientists didn't have to spend time on these tasks and could focus entirely on artificial intelligence?

The 3D Platform addresses this by including a data warehouse generation capability that automatically aggregates data from various business unit systems into a single data warehouse for reporting.

Key Features:

1. **Aggregated Schema Creation:** The DWG utility begins by merging data dictionaries from all listed systems, generating a consolidated database schema.
2. **Handling System Customizations:** Specialized columns and tables added by individual systems are recognized and seamlessly integrated into the data warehouse.
3. **Merging Transactional and Master Data:** Transactional data, such as journal entries, and master data, such as contacts, are merged across all systems, resulting in a comprehensive and consistent dataset in the warehouse.
4. **Governance-Based Precedence:** When conflicts arise, the system governing a particular master record takes precedence, ensuring data integrity and consistency.
5. **Automatic updates:** The utility will update the data warehouse when data from the source systems change.

The DWG utility addresses a critical issue in the IT industry by providing organizations with a powerful tool for automatically aggregating data for reporting and AI purposes.

4.6 MASTER DATA QUERY ENGINE

The Master Data Query Engine is a powerful tool that provides comprehensive insights into any master record by retrieving not only the record itself but also all related child, grandchild, and deeper levels of data. Designed for efficient and in-depth data analysis, the engine allows organizations to gain a complete picture of any master data entity by specifying parameters like query date, data levels, and filters.

4.6.1 Key Features

1. **Multi-Level Data Retrieval:** The engine supports drilling down through multiple levels of related data, offering a complete snapshot of all relevant records connected to the master record.
2. **Date-Based Filtering:** Users can query data as of a specific date, allowing them to see historical data states. For example, in a government system, users could retrieve

[1] According to dataversity.net and other sources (https://www.dataversity.net/survey-shows-data-scientists-spend-time-cleaning-data/)

information about a land parcel and its associated entities (e.g., title holders, wells, encumbrances) as of a particular date.
3. **Automatic View Generation:** The engine leverages the data dictionary to automatically generate views that present aggregated and structured data.
4. **Flexible Query Parameters:** Users can filter results based on table definitions, query conditions, foreign key relationships, and levels of data depth. Common filtering options include system of origin, record status, and more.

4.6.2 Benefits

- **Comprehensive Data Aggregation:** Like a "Google search" for master data, this engine quickly aggregates and presents a complete picture of any entity.
- **Enhanced Decision-Making:** With instant access to in-depth data insights, users can make well-informed decisions more effectively.

The Master Data Query Engine is a transformative feature that streamlines access to critical data and helps organizations leverage their information for better decision-making.

4.7 Financial Module

The 3D Platform's financial module is engineered to meet the evolving needs of modern enterprises, offering an integrated and adaptable solution that addresses the limitations of traditional accounting systems. Here are four key strengths of the 3D financial module:

4.7.1 Integrated General Ledger and Sub-Ledgers for Seamless Data Consistency

The 3D financial system integrates sub-ledgers directly with the general ledger, ensuring that financial data remains consistent and up to date. This approach eliminates data duplication and reduces the effort required to keep financial information in sync across systems.

- **Real-Time Updates:** Integrated sub-ledgers and general ledgers ensure financial data is always current, giving executives and stakeholders access to accurate and consistent figures.
- **Streamlined Reconciliation:** With no separate sub-ledgers, reconciliation tasks are reduced, enhancing the efficiency of financial operations.
- **Improved Data Integrity:** Consistent and up-to-date information across all financial records prevents discrepancies and builds trust in financial reporting.

4.7.2 Adaptive and Scalable Hierarchical Account Structures

The 3D Platform's financial system employs a truly hierarchical structure for both General Ledger (GL) Account Codes and Cost Centers, offering flexibility and scalability beyond traditional systems. Instead of relying on rigid, text-based account codes with embedded data, the platform's

hierarchical structures allow organizations to tailor their financial setup to their unique needs, regardless of size or complexity.

- **Scalable for Any Organization Size:** Whether for a small organization with a few levels or a large enterprise with complex, multi-tiered hierarchies, the system adapts to fit the need. The hierarchy can be shallow or deep as necessary.
- **Real-Time Flexibility:** Live queries at any level of the hierarchy enable users to analyze breakdowns by account or cost center, view detailed journal entries, or see summary reports by month. The system allows for seamless expansion or collapse of levels based on evolving needs.
- **Seamless Adaptation to Change:** The platform's flexible structure makes it easy to add new levels into the Account and Cost Center structures without disrupting data integrity. Unlike traditional systems, where hierarchical data is embedded in account codes and tied to foreign keys, the 3D Platform decouples these layers, preventing the need for extensive updates to historical data when changes occur.

4.7.3 Multi-Dimensional Financial Analysis Linked to Enterprise Data

The 3D financial system provides the ability to link financial data directly to enterprise entities, enabling comprehensive analysis beyond traditional Account Codes and Cost Centers. Financial transactions can be associated with assets, contracts, projects, land parcels, and other relevant entities.

- **Detailed Financial Reporting:** Generate financial statements for specific entities, such as assets, project types, or acquisition methods, for deeper insights and targeted analysis.
- **Dynamic Entity Linkages:** The system ensures journal entries are automatically linked to relevant enterprise entities, reducing manual workarounds and streamlining processes.
- **Scalable and Contextual Reporting:** Directly connecting financial data to enterprise entities supports robust reporting capabilities that grow with the organization.

4.7.4 Flexible Support for New Transaction Types

The 3D financial module is built to accommodate new transaction types as organizations automate additional processes, such as invoicing for goods and services. The system's flexibility eliminates the need for separate sub-ledgers, allowing for seamless integration of new transactions within the core financial system.

- **Efficient Process Automation:** New transaction types can be easily incorporated, reducing the need for costly and redundant sub-ledgers.
- **Streamlined Data Management:** By supporting a wide range of transaction types natively, the system minimizes duplication and enhances overall efficiency.
- **Customizable Transaction Workflows:** Organizations can tailor transaction processes to align with unique requirements, ensuring the financial system adapts as business needs evolve.

4.7.5 Conclusion

The 3D Platform's financial module offers a robust and adaptable solution that meets the evolving needs of today's enterprises. With integrated ledgers, flexible account structures, multi-dimensional analysis, and seamless support for new transaction types, the platform enhances operational efficiency, scalability, and financial clarity across all levels of the organization.

5 PROOF OF CONCEPT ENVIRONMENT

The 3D Platform includes a comprehensive test environment designed to validate every aspect of the platform's functionality, scalability, and performance. The Proof of Concept (POC) environment consists of various enterprise systems operating within a government subnet and external systems that interact with them. Each system has been designed to illustrate the seamless exchange of data and the ability to aggregate information across business units.

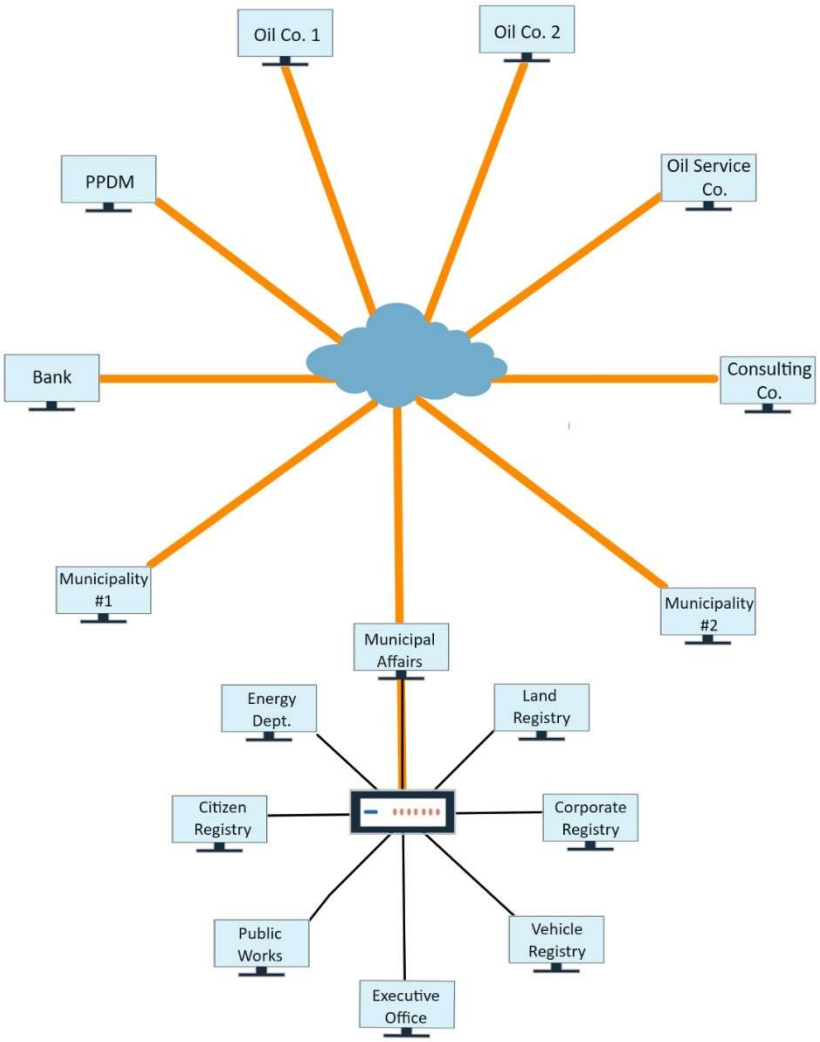

5.1 Sample Enterprise Systems

Below is an overview of each system, highlighting the data they manage, import, and export.

5.1.1 Government Subnet Systems

1. **Citizen Registry System:** The Citizen Registry manages all citizen data within the government, including temporal relationships such as marriages and parent/child connections. Additionally, it handles its own financial subledger, projects, activities, and

equipment. Updates to citizen records trigger automatic exports to subscriber systems, while financial and budget transactions are continuously shared with the central government system. Asset information is also sent to the central government registry.

2. **Corporate Registry System:** This system tracks all corporations and their key resources (e.g., Presidents). It automatically exports corporate data updates to other government systems and external entities. The Corporate Registry integrates citizen data from the Citizen Registry and exports financial transactions, budget details, and equipment to the central government system.
3. **Vehicle Registry System:** The Vehicle Registry maintains a temporal record of all vehicles within a state or province. It records vehicle details and revenue transactions related to updates. When a vehicle's status changes, updates are sent to the organization that owns the vehicle. Financial, budget, and equipment information is exported to the central government system.
4. **Public Works System:** The Public Works system coordinates projects and activities for other government departments. It maintains an extensive equipment registry, including heavy equipment. It accepts activity data from external consulting firms and other departmental systems involved in its projects. Public Works exports project information, activity updates, and financial transactions to relevant departments and the central government.
5. **Land Registry System:** This fully temporal system tracks land parcels, titles, and registered interests (e.g., mortgages). The Land Registry provides a legal framework and monitors subdivisions and land aggregation plans. It handles financial transactions associated with land changes and exports land registry data to interested entities, such as municipalities, banks, and oil and gas companies. It also exports financial data, budget transactions, and equipment details to the central government.
6. **Municipal System:** Supporting the municipal department, this system aggregates taxation data from all municipalities and integrates updates from the Citizen and Corporate Registries. Financial transactions, budget data, and equipment details are exported to the central government system.
7. **Energy Department System:** The Energy Department accepts temporal well data from oil and gas companies, based on the PPDM industry model. With over 2,000 tables, this system demonstrates how large datasets are handled. It exports financial transactions, budget information, and equipment data to the central government.
8. **Central Government Office System:** The central system manages reference table updates and houses a consolidated financial ledger encompassing all government finances. It maintains a central equipment registry to facilitate equipment sharing between departments. The system uses a data warehouse generation capability that consolidates data across all departmental systems into a unified data warehouse. It ensures that data governance principles are maintained by importing records from the systems that govern them. The Master Data Query Engine allows rapid access to all relevant data, such as citizen or land parcel details.

5.1.2 External Systems

1. **Municipal Taxation Systems (2):** Each system manages temporal data for all properties within its jurisdiction, including rate payers, mill rates, assessments, and revenue. These systems integrate land title updates from the Land Registry and export taxation data to the Municipal System for aggregation. They also accept updates from the Citizen, Corporate, and Land Registries.
2. **Banking System:** This system demonstrates how banks receive updates for land parcels where they hold mortgage interests. It is automatically updated with citizen or corporate changes related to properties under mortgage.
3. **Consulting System:** Used by consulting firms engaged in government projects, this system exports project activity and financial data to the Public Works department. It includes a full financial system, equipment registry, and contact management functionality. Activities from this system are aggregated with departmental activities to provide a comprehensive view of project costs and progress.
4. **PPDM System:** The PPDM system showcases the management of the oil and gas industry model and reference values shared across companies. It handles its own financial, equipment, and project activity records and exchanges data with other oil and gas systems.
5. **Oil and Gas Systems (2):** These systems feature 2,000 tables based on the PPDM model, populated with scrambled data from live environments. The systems exchange data for partnerships and send relevant well data to the government's Energy Department system. Additionally, they share reference table updates with other oil and gas systems. Their internal financial systems are integrated with projects and equipment management.
6. **Oil and Gas Service Company System**
7. This system supports oil and gas well maintenance. It exports activity updates to the relevant oil and gas companies and manages its own financial, equipment, and activity data.

5.1.3 Conclusion

The POC environment is a sophisticated demonstration of our enterprise systems integration capabilities. By effectively orchestrating a network of interconnected systems, it showcases the potential for improved data governance, operational efficiency, and real-time data exchange within a complex governmental landscape. This POC serves as an exemplary model, illustrating how tailored enterprise architectures can foster collaboration, enhance accountability, and optimize resources in public administration and other sectors.

This detailed integration highlights the platform's adaptability, handling high data volumes and complex transactions while ensuring data integrity across all systems. We believe this POC can inspire future developments in enterprise system architecture, potentially transforming how organizations manage and share data.

5.2 TEST CASES

The following test cases verify that the 3D Platform and the POC environment are operating as expected. These scenarios illustrate the platform's capabilities in handling data integration, synchronization, and management across multiple systems in complex environments.

1. **Subdivision and Property Sale**
   A property is subdivided and sold to a newly married couple. The couple registers a mortgage, triggering updates across the following systems:
   - The Land Registry (parcel and ownership details).
   - The Bank (interest registered against the property).
   - The Oil and Gas Company (due to a well on the property).
   - The Municipality (parcel updates).
   - The Central Government (financial transactions).
2. **Municipal Tax Roll Creation**
   The Municipality receives updated parcel information and creates new tax rolls and assessments. These updates are exported to the municipal department within the government.
3. **Marriage, and Name Change**
   A couple gets remarried and the new partner who changes their name. Updates are automatically broadcast to relevant systems:
   - The Land Registry (for property title updates).
   - The Corporate Registry (since one party is a CEO).
   - The Bank (holding a mortgage on a property related to the parties).
   - The Municipality (for tax records).
   - The Public Works Department (as one party is an employee). Financial transactions for registration fees are automatically exported to the central financial system.
4. **Temporal Update to a Land Parcel**
   A temporal update is made to a land parcel. The test verifies that all systems subscribed to this record receive the update correctly.
5. **Adding a System to a Subscriber Group**
   A system is added to a subscriber group, automatically receiving any record groups the group is subscribed to. The test checks that updates occur both within the subnet and to external systems.
6. **Subscription to a Three-Level Record Group**
   A subscription is created for a record group with a three-level structure (e.g., employee details, contact information, and journal entries). The test ensures that changes are correctly exported to the subscriber system.
7. **Table Subscription with Filtering**
   A subscription is created for an entire table with filtering criteria applied. The subscription spans both local and external systems. This test confirms that table subscriptions and filtering function correctly.

8. **Temporal Data Subscription Across Multiple Systems**
   A subscription is made for temporal data across multiple systems. The test validates that temporal updates are consistently applied across all subscribing systems.
9. **Project Subscription and Redaction Capability**
   A subscription is created for a project, including associated activities and resources. The test also verifies redaction capabilities, ensuring employee details are generalized to their company. Data is exported to both local and external systems.
10. **Adding a Subscriber to a Subscription Group**
    A subscriber is added to a subscription group, automatically receiving all subscribed data, including contact, address, GL rate, and rate type information.
11. **Special Character Handling**
    Data is updated with special characters and exported to all subscribing systems. This test ensures that special characters do not cause issues during data export.
12. **Undo Functionality Validation**
    The undo functionality is tested, verifying that transactions can be reverted and exported correctly to local and external systems. The test involves undoing changes from the special character test.
13. **Lower-Level Record Group Updates**
    A lower-level member of a record group is updated. The test ensures the update is exported while filtered data (e.g., home addresses) is excluded. Changes are verified across subscribing systems.
14. **Flow-Through of Temporal Updates**
    A temporal update is made and verified across four systems—three local and one external. The test ensures consistent update propagation.
15. **Deleting a Temporal Record Series**
    A temporal record series is deleted, and the deletion is verified across all subscriber systems.
16. **Undoing a Delete Operation**
    The test verifies that undoing a delete operation restores the record in all subscribing systems.
17. **Inserting and Referencing New Values**
    A new address type is added to a local system and then referenced on an address. The test ensures that the main record and address type are properly picked up and exported to the destination system.
18. **Three-Level Subscription (Parent, Child, Grandchild)**
    This test involves a parent system, a child system subscribing to the parent's data, and a grandchild system subscribing to the child's data. Changes made at the parent level flow through to the child and grandchild systems.
19. **Handling Table Modifications**
    The test ensures that table modifications in both the parent and subscriber systems are detected and notified. Approved changes are automatically applied, such as adding a column with default values or non-null constraints.
20. **Negative Test Scenarios**
    Various adverse scenarios are tested, including dropped columns, changed constraints, or

disconnected databases. The system's ability to handle these disruptions without data loss is verified.

21. **Data Warehouse Generation and Inquiry Engine**
    The data warehouse generation function is tested to ensure it creates a merged data model and populates it correctly. The Master Data Inquiry Engine is validated to extract master data and related records across multiple levels.

5.3 CONCLUSION

Through these detailed test cases, the effectiveness of the 3D Platform's capabilities to integrate, manage, and communicate within various enterprise systems is thoroughly validated. Each test case is designed to scrutinize specific functionalities, ensuring that the POC environment operates under expected parameters, maintains data integrity, and provides a reliable framework for future deployments.

6 TOOLS TO ASSIST WITH THE CREATION OF SYSTEMS

The 3D Platform is equipped with a comprehensive suite of tools designed to automate and streamline the creation of new systems within the model. Driven by a data dictionary, these tools significantly reduce the effort required to generate and populate new systems, allowing development teams to start with fully populated datasets. As systems evolve, new datasets are continuously generated in alignment with updated models.

- **Reverse Engineering Existing Models:** Tools are available for reverse engineering existing data models. This process updates the central dictionary with table definitions and foreign key references, using user-friendly terminology to standardize table and column names. An abbreviation engine assists in consistently applying abbreviations across attribute and table names.
- **Design New Models:** Designing new models involves mapping existing data into the core data structure and extending it to meet the unique requirements of each business unit. As systems are migrated, this information is captured, enabling AI to assist in automatically mapping future systems with minimal manual analysis.
- **Generate the Database:** Utilities facilitate the creation of new databases based on the latest data models. They ensure that the data dictionary definitions pass integrity checks before the database is generated, adhering to the new model specifications.
- **Convert the Data:** The data conversion tool is a powerful engine that simplifies the data conversion process. Key features include:

    - Merging data from two tables into one within the new model.
    - Allowing for manual transformations based on business rules, such as converting a set of values into foreign key references.
    - Splitting data from one table into multiple tables in the new model.
    - Merging duplicate master data from multiple systems.

- Transforming composite keys into the new 3D primary key format.
- Standardizing the handling of temporal data.
- Separating classification data into distinct reference tables. While the data conversion engine handles many transformations, manual data conversion steps are incorporated to manage exceptions and unique cases.

- **Generate Integrity Constraints:** The system programmatically applies data integrity constraints, such as foreign key, null, and unique constraints, ensuring data consistency and reliability within the database.
- **Generate Views:** The system generates views to enhance data browsing, utilizing user-friendly names and ensuring that master data is accurately represented through foreign key translations. This makes the data more accessible and understandable for users.
- **Refine and repeat:** This phase involves refining the data dictionary and system functionalities to correct any errors and incorporate new requirements and then repeating the entire data conversion process to generate a new data for testing. Testing with live data in the newly created forms ensures that data conversion and system functionalities align with user needs and business rules.

By segmenting the development process into manageable phases, the pathway to a comprehensive and robust Enterprise System is streamlined and the associated costs and risks are reduced.

7 EDUCATIONAL VIDEO SERIES OVERVIEW

Introduction

The educational video series covers a wide range of topics, from basic introductions to advanced features of the 3D Platform. Each video is crafted to address the specific needs of different audiences within an organization, from IT professionals to executive managers.

Video Topics Covered

- **Architecture Introduction:** A comprehensive guide to Enterprise Systems Architecture, highlighting crucial aspects like change history, master data queries, and financial modules, essential for auditing, debugging, and data exchange.
- **Data Dictionary:** Discusses the critical role and upkeep of data dictionary information, including detailed explanations of column purposes and naming conventions.
- **Data Conversion:** Describes the comprehensive process of integrating new systems into the model, emphasizing reverse engineering and data validation.
- **Primary Keys:** Details the design of primary keys within enterprise systems, their assignment, and their wide-ranging applications.
- **Temporal Data:** Explores the management and practical applications of temporal data within database design.

- **Change History:** Investigates the operational facets of change history, its initiation, and its contributions to security and data management.
- **Data Exchange:** Examines the shift from old to new paradigms in data exchange, detailing the processes of exporting, distributing, and importing data.
- **Record Governance:** Discusses the importance and implementation of record governance within system architectures.
- **Data Warehouse Generation:** Delves into the creation of data warehouses for advanced reporting and AI, focusing on schema aggregation and data integration.
- **Master Data Query:** Introduces innovative approaches to querying enterprise databases that facilitate rapid data aggregation for decision-making.
- **Financial Module:** Reviews the objectives and new developments in the financial module, addressing old challenges and showcasing new-generation routines.
- **Regen Environment:** Outlines the procedures for generating new enterprise system platforms based on revised data dictionaries.
- **Application Framework:** Describes the architecture of a global enterprise system network, focusing on an innovative framework that aids in cost reduction and enhances system longevity.
- **Database Design Introduction:** Sets the stage for in-depth discussions on database design.
- **Polymorphic Joins:** Unravels the complexities of polymorphic joins in database design, highlighting their utility in linking common functionalities to various database objects.
- **Recursive Structures:** Discusses the intricacies and applications of recursive structures in database design.
- **Sub-typing:** Focuses on the benefits of subtype structures in achieving cleaner database designs and enhancing record governance.
- **Core Models:** Examines the role of core models in managing universal organizational data.
- **Testing Environments:** Describes the testing environments developed for regression tests and proof of concept demonstrations, showcasing the advanced capabilities of the system.

Conclusion

This video series is designed for business owners, IT professionals, or anyone interested in the internal workings of the 3D Platform, offering a wealth of knowledge to help you develop efficient and integrated solutions for the contemporary business landscape.

8 NEXT STEPS

As we look to the future of the 3D Platform, our focus remains on enhancing integration capabilities, expanding the model to include new industry sectors, and ensuring our users are equipped to fully leverage the platform's potential. Here’s a structured plan to guide our next steps:

- **Technical Review with Assigned Resources:** Initiate comprehensive technical reviews with designated teams to identify and address any initial integration challenges, ensuring a seamless deployment and operational experience.
- **Incorporation of New Core Data Subjects:** Broaden the scope of our core data models to include critical business elements such as contracts and inventory maintenance. This expansion will enhance the platform's applicability across various operational needs.
- **Integration of New Industry Models:** Extend the platform's capabilities by incorporating industry-specific models. This step includes merging core data from these models with existing structures and testing models in the POC environment to ensure seamless data exchanges with other industry models.
- **Development of the User Interface:** Develop a highly parameterized user interface that integrates directly with the data dictionary, facilitating ease of use and reducing the learning curve for new users. The goal will be to automatically generate a user interface for 80% of the system directly from the data dictionary, with the remaining custom-developed to ensure flexibility and meet specific user requirements.
- **Establishment of a Governing Structure:** Establish a governance body to develop and enforce standards and policies that guide the creation of new systems and the integration of existing ones.
- **Educational Initiatives:** Increase the availability of educational materials and training sessions to ensure users at all levels can fully utilize the platform's capabilities.
- **Strengthening Community Engagement:** Foster a vibrant community of developers and users who can contribute to the platform's evolution, share best practices, and provide feedback on potential improvements.

9 FREQUENTLY ASKED QUESTIONS

1. **How does the platform handle scaling and performance with complex data exchanges?**

   The platform uses a custom eight-byte key strategy combining a system ID with a sequential record ID for efficient storage, indexing, and querying. The sequential assignment minimizes index fragmentation and ensures global uniqueness across systems, even in large-scale deployments. System IDs are coordinated similarly to TCP addresses, optimizing performance for real-time data exchanges while ensuring scalability.

2. **How does the platform ensure data security and privacy during data exchanges?**

   The platform supports secure data exchange within conglomerates via internal subnets and for global data exchange through a global server, ensuring data delivery to the intended recipients. Although advanced encryption standards are still in progress, the framework is designed to support secure transmission. By managing privacy concerns at the platform level rather than system-by-system, data security is more consistently enforced.

3. **How does the platform manage conflicts when multiple systems update the same data?**

   The platform uses record governance, where the system that creates a record automatically governs it. Governance can be transferred if needed, such as when an equipment is sold, or a member moves districts. This ensures that only the relevant system maintains authority over the data, preserving data integrity and resolving conflicts across systems.

4. **How easy is it to customize and extend the platform while maintaining core functionality?**

   Customization is straightforward. The data exchange system handles changes to columns and tables, alerting users when updates occur. All metadata is managed within the same data structures as regular data, allowing flexible imports into test environments for safe customization without breaking core functionality.

5. **How does the platform ensure consistency and standardization across systems?**

   A governing body oversees the core structures, ensuring consistency while allowing organizations flexibility. Straying from core structures may limit interoperability and access to platform functionality. Organizations can balance customization with the need for consistent data governance.

6. **What is the onboarding process for adopting the platform, especially for migrating legacy systems?**

   Onboarding involves reverse engineering existing models using Excel sheets, automated database generation, conversion procedures, and automatic constraint/view generation. The process is scalable, starting with smaller systems and expanding to more complex environments as the platform is validated.

7. **What is your strategy for encouraging widespread adoption?**

   The initial strategy focuses on using the PPDM model to demonstrate benefits for the oil and gas industry. Beyond that, broader rollouts aim to challenge traditional IT practices by showcasing the platform's deep integration and seamless data exchange capabilities.

8. **How does the platform handle disaster recovery and fault tolerance?**

   The platform ensures data integrity through ACID-compliant databases, automated synchronization, and robust redundancy measures. Built-in data exchange tools support seamless recovery across systems, while additional custom interfaces can be developed for organizations with higher security and recovery needs.

9. **How is the global data exchange architecture structured, and how does it manage latency and redundancy?**

   The platform uses distributed nodes, edge caching, DNS-based load balancing, and regional synchronization for efficient global data exchanges. Redundancy is maintained through automated replication and failover strategies, ensuring high availability across different regions.

10. **How does the platform integrate with non-conforming systems?**

    The platform uses stub systems to bridge gaps between conforming and non-conforming systems. Once an integration is developed, it becomes reusable across multiple organizations, simplifying future integration projects and enhancing the platform's interoperability.

11. **Are there real-world case studies or pilots demonstrating the platform's effectiveness?**

    No live deployments exist yet, but the platform has been tested in a proof of concept (POC) environment with 16 fully functional enterprise systems. The POC validates the platform's capabilities across a wide range of scenarios, preparing it for real-world adoption.

12. **How does the platform ensure compliance with data sovereignty laws and regulations like GDPR?**

    The platform primarily exchanges low-risk data while providing controls for more sensitive information. Data redaction, classification, and auditing features allow organizations to maintain compliance while still benefiting from global data exchanges.

13. **What's your strategy for building a developer community around the open-source platform?**

    The platform will launch on an open-source repository with comprehensive documentation, contributing guidelines, and a community-driven roadmap. Developer tools, SDKs, and a plug-in marketplace will support contributions, while partnerships and early adopter programs foster collaboration.

10 SUMMARY

The IT industry has reached a breaking point in how we create and manage enterprise systems. The traditional approach—building systems by business function and addressing integration as an afterthought—has resulted in a landscape filled with siloed, non-integrated systems that are nearly impossible to connect. Organizations are burdened with costly, complex solutions that fail to scale and meet the demands of modern businesses, especially in large, multifaceted entities like governments and global enterprises.

The 3D Platform represents a paradigm shift in how systems should be designed, deployed, and integrated. By rethinking the foundation of enterprise architecture, this platform addresses the root issues that have plagued the industry for decades. Rather than implementing systems by function and struggling with integration later, the 3D approach organizes systems by business unit and allows them to seamlessly share data, ensuring that every piece of information is only entered once and made available to those who need it.

This platform is not proposed as a catch-all solution—it strategically focuses on central financial and administrative systems while enabling the continued integration of high-performance operational systems that drive large online businesses. By flipping the axis on how we think about systems, the 3D Platform offers a scalable, flexible, and future-proof alternative that can be tailored to the unique needs of any business unit while facilitating deep integration across entire conglomerates and supply chains.

The 3D Platform not only addresses today's integration challenges but also prepares organizations to fully leverage AI and advanced analytics. Organizations that embrace this paradigm shift will be better positioned to outcompete their rivals, as they will be able to unlock the full potential of their data without being weighed down by the complexity and inefficiency of outdated systems.

The time has come for the IT industry to adopt a new way of thinking—one that prioritizes integration and scalability from the start. The 3D Platform is the blueprint for the future of enterprise systems, enabling organizations to build smarter, more adaptable solutions that can meet the demands of a rapidly changing world.