PreferredStrategies
BI & REPORTING SOLUTIONS FOR JD EDWARDS

WHITE PAPER

Governance and Security in a JDE Self-Service Data and Analytics Environment

As more and more business users require access to data in order to perform their jobs, the burden on IT to handle requests for access and specific data views has become increasingly challenging. Supporting the business with inbound data requests and dashboard views of data takes critical IT staff away from other daily support tasks.

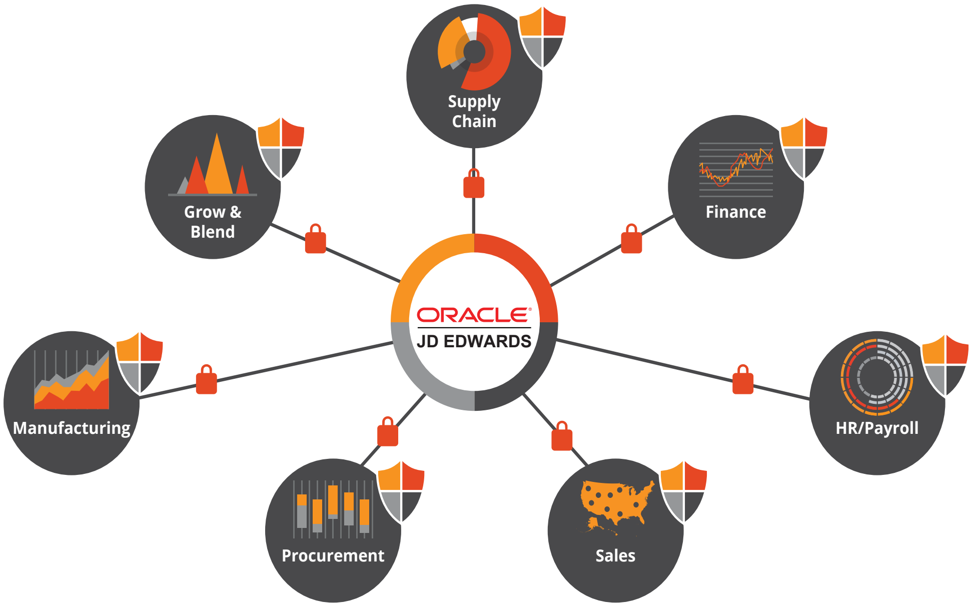

Modern BI tools such as Microsoft® Power BI have alleviated some of the reporting and ad hoc analytics visualization workload for IT. However, this has created a new challenge for the organization summed up in three simple data governance questions:

1. Who should have access to data?
2. What data should those people have access to?
3. Which attributes should they be able to see in that data?

Who should have access to data?
This question is naturally very business dependent, but an increasing reliance on data to run core systems within an organization affects everyone from HR to manufacturing. For most business users reporting—answering the question "what happened in my business?"—is table stakes. Understanding trends in reported data is far easier to do visually than in Excel table format. Data or underlying business anomalies are easier to identify using modern BI tools.

Because organizations rely so heavily on data to understand trends, it makes sense to equip all business users with modern BI tools rather than continuing to rely on Excel. Visualization software allows business users to more easily identify and act on trends and anomalies before they become an issue. And not all trends or anomalies are bad; some might identify positive sales patterns, or a reduction in in-bound support calls, etc.

As business users become more proficient with modern BI tools and those tools become ever more capable, trends in data can be used to feed predictive models. Power BI offers Quick Insights, essentially machine learning or AI algorithms that run data science models, such as a linear progression, K-Means function and the like, against historical data. These powerful tools can turn regular business users in to citizen data scientists and provide some remarkable insights to the business.

What data should those people have access to?
The next question is easier to answer. Access can be granted by a wide range of factors, seniority, geography, product division, operational function, and so on. We find companies that foster corporate-wide transparency often make data widely available across the organization. And this makes a lot of sense. If you have two business divisions, what downside is there in letting the other business division have visibility into performance metrics? Indeed, there may be real upside to this kind of access. Economies of scale, vendor or supply chain commonality, shared distribution and shipping costs, and more may be revealed by these kinds of cross-company insights.

Corporate performance is a company-wide activity. It's not uncommon for Key Performance Indicator dashboards to be displayed prominently around the company in breakrooms and open-plan office spaces. This helps the organization understand where they are against metrics and fosters a team effort environment.

The other key consideration here is what data can be changed, a question distinct from which users have write privileges to source data, although that is a very important decision, too. Instead, the question focuses on what changes an individual can make to a visualization that has been shared with them. In modern BI tools this is most commonly delivered by filter controls that let a user simply select the data that is important to them. In a JD Edwards sales visualization, for example, the sales director might be able to see her whole region and then drill down into a specific territory or product category. The VP of North American sales will be able to see data for their whole region, whereas the Chief Revenue Officer might be able to see global sales data with the ability to filter and run reports to meet specific sales insights.

Which attributes should they be able to see in that data? This question is perhaps the easiest to answer because it is dependent on job function. The HR team should be able to access employee records, payroll, benefits administration, etc. The finance team naturally needs access to all financial data. The interesting part is when cross organization insights are required.

For example; for a divisional or operation leader to administer bonuses and raises for his organization, it's valuable to give him access to some data so he can maintain overall corporate guidance against raises or bonuses. This can be handled using data security, showing individual pay grades or titles with proposed raises across the organization, while obscuring individual employee names.

This type of data security is established and maintained by a corporate data governance policy that can administered by IT or, as is common more recently, by a dedicated data steward.

The Evolving Role of the Data Steward

The Wikipedia definition of a data steward is quite accurate:

A ***data steward*** is a role within an organization responsible for utilizing an organization's data governance processes to ensure fitness of data elements - both the content and metadata. Data stewards have a specialist role that incorporates processes, policies, guidelines and responsibilities for administering organizations' entire data in compliance with policy and/or regulatory obligations.

According to a report published by Boston Consulting Group (BCG) in 2017 based on first party research, over 71% of companies have established protocols that govern data access which are enforced by a data steward.

The data steward can act as an intermediary between the business and the IT organization, alleviating some of the data access requests to IT by granting access on a reinforced access policy. Working together, IT and the data steward must establish clear lines and protocols for escalation when access may be denied or modified. Similarly, relinquishing access to former employees or third party user of the corporate data is another typical role.

In many respects the data steward acts as a librarian or docent to corporate data. The data steward should provide descriptions of data sources or data sets; essentially business glossary terms that reveal the true meaning behind obscure data attribute names. Preferred Strategies QuickLaunch™ is essentially a virtual data steward in this capacity. QuickLaunch provides a comprehensive business glossary for JD Edwards to allow business users across the organization to find and use the JDE data that is important to them. Here are some examples of the "translated" business glossary from JDE data as seen by QuickLaunch:

Data Type	JDE Metadata Description	Actual Business Term
Table Name	F0101	Address Book
Field Name	MCMCU	Business Unit
Data Field: Date	107181	6/30/07
Data Numeric	15025	150.25

Building and Governing a JDE Data and Analytics Environment

How you expose JDE data to business users plays an important role in governing and securing the environment that enables self-service data exploration and use. Connecting modern BI seats directly to your JDE environment is not advisable—especially if you have lots of users—for a number of reasons, including compliance and the impact on your core system by concurrent users. There are various intermediary data storage solutions, such as a data warehouse or cloud-based data repository, to which you would then attach your BI seats.

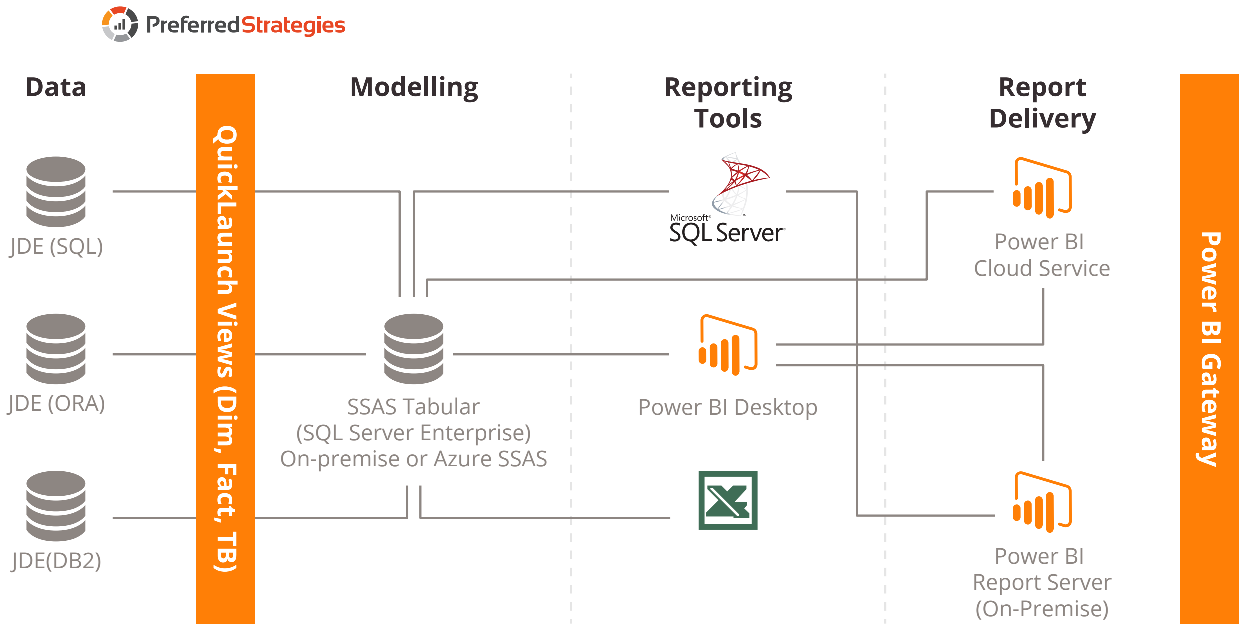

A typical JDE environment showing Preferred Strategies QuickLaunch integration providing self-service reporting and analytics across the enterprise under a governed and secure data model.

Another approach is to implement the Preferred Strategies Direct Connect method. Amongst other benefits to JDE self-service and data and analytics applications, this approach provides a layer of governance on your JDE data by only exposing data to your users based on your specific access policy.

When combined with Microsoft Analysis Services and the tabular or cube model, the governance of calculated values or metrics is also possible, as described later in this white paper.

JDE Security at a Company or Business Unit Level
In a multi-divisional or business unit level it is often desirable to secure data at a more granular level. Defined as JDE Row Level Security this is often handled externally from the JDE environment by employing an access security layer such as LDAP or Active Directory. With Preferred Strategies Quicklaunch this row level security can be extended to become an integral part of the analytics solution.   This user or group level security means that users of QuickLaunch tabular model will only see the data they are authorized to view, whereas other users may see the full model.

Historically data and analytics tasks have been handled by IT. The business user would request a report, and the IT department would source, cleanse, enrich and transform the data using Extract, Transform and Load (ETL) processes. Because the IT person is knowledgeable about data she can see when there are data errors and perform cleansing, normalization or other processes as needed prior to transformation. This tribal knowledge is usually not transferred from one user to another or captured as necessary steps when using a specific data source. When that person leaves the company, the tribal knowledge leaves with her, requiring the next member of the IT department to learn the process and nuances of data sources all over again.

The concept of master data management (MDM) was introduced to alleviate some of this tribal knowledge lock-in. The concept is straightforward; if there is a data source that is accessed regularly, such as customer data or sales data, then perform these cleansing and normalization tasks and store a new, clean, managed data set that others can access. Operational processes must be put in place to implement Change Data Capture or Incremental Data Capture functions on the source dataset, and to ensure that the same processes are applied and then update the master dataset—the trusted data—for that source.

Beyond the data is the need for clarity around data attributes or column headers in a database. These are often obscure, such as those in JD Edwards ERP systems, whereby a business user would be unlikely to know the meaning of the attribute name. Preferred Strategies QuickLaunch leverages the JDE dictionary by exposing the translation from obscure attribute descriptions to BI tools.

Trusted Data Sources [image of data with shield here center page with text wrapped around it]

In the new paradigm of self-service data and analytics, the importance of trusted data becomes even more paramount so that business users across the organization can feel confident about data integrity. This can be achieved by hosting the curated or trusted data in an environment such as a data warehouse or cloud hosting solution such as Microsoft Azure or AWS and connecting BI tools like Microsoft® Power BI to that data. The Preferred Strategies QuickLaunch Direct Connect and Data Warehouse option provides this level of data governance. The Analysis Services Tabular/Cube option provides for an even greater level of governance and security, namely, metrics or calculated values governance.

The business user will know that the insights being generated are derived from data that the whole organization believes to be the most accurate—the real power behind trusted data. In this scenario, any debate is about the insights being derived or presented, and not about the accuracy of the data itself.

Trusted Calculated Measures and KPIs
This same level of governance and trust extends beyond source or derived data into calculated measures or KPIs, such as cost of goods sold (COGS). COGS, or sales margin, often requires multiple attributes to be combined together, such as raw material costs, handling costs, manufacturing costs, packaging costs, documentation costs, shipping costs, etc. In some cases there may be as many as 30 or more attributes, leaving plenty of margin for error. And yet COGS is a critical number to get right if you are trying to understand profitability. Manufacturers with multiple divisions—especially where a division has been acquired—often find that calculated values are derived differently between business units, again undermining the trust of the derived business insights and reports.

The good news is that there are tools on the market to help JD Edwards customers identify and manage trusted data. Preferred Strategies QuickLaunch, for example, leverages over 10 years' of knowledge built up around JD Edwards and provides governance for calculated measures and a host of other business-specific calculated values. When using QuickLaunch, the entire organization can be assured that the when they drag Revenue or Quantity on Hand onto a visualization canvas the underlying data and calculations will be correct. These calculated values are "locked down" by a data governance policy that prevents users from changing the underlying data attributes or the calculation.

How many presentations have you been in when reviewing analytics or reporting on data-driven insights where the first question from the audience is, "Where did you get this data?" Trusted data answers this question—unequivocally!

You Can Handle The Truth!

Contrary to the opinion of Colonel Nathan Jessup, portrayed by Jack Nicolson in the 1992 movie ***A Few Good Men***, not only can you handle the truth, but you must as an essential aspect of a self-service data environment.

For business leaders who rely on reports to gain analytical insights, the accuracy of the source data is of utmost importance. If the source data is suspect, then the resulting outputs must be given an equal amount of suspicion.

Governed, trusted data is an essential element to ensure that business users across the organization handle the truth about their data.