



Microsoft IT deploys a federated Business Intelligence Portfolio

Published February 2014

The following content may no longer reflect Microsoft's current position or infrastructure. This content should be viewed as reference documentation only, to inform IT business decisions within your own company or organization.

Read about the system design principles and development methodology Microsoft IT's Service Deployment & Operations (SDO) organization followed to create a federated business intelligence (BI) portfolio that interconnects a number of existing BI data sources and provides internal customers with a single source of cleaned, simplified data. This paper also describes SDO's ongoing efforts to enhance the federated BI portfolio with a set of component layers that support an array of data analysis and reporting capabilities to help the business make data-driven decisions.

Situation

Microsoft IT's Service Deployment & Operations (SDO) organization wanted to support the company's *One Microsoft* initiative to provide end-to-end views of key internal processes by streamlining how it provides BI data to internal customers.

Solution

SDO federated some of its BI data sources into a BI portfolio that presents cleaned BI data through a single portal. This "one stop shop" enables SDO to maintain its current BI environment while simultaneously evolving the BI portfolio into a more strategic end-state that addresses longer-term business priorities.

Benefits

- Increased efficiency
- More scalable platform
- Improved adherence to standardized processes
- Unified data retention

Products and Technology

- Microsoft System Center 2012 Configuration Manager
- Microsoft System Center 2012 Operations Manager
- Microsoft Excel Power Pivot
- Microsoft SharePoint Power View and Power BI

Situation

To help guide their strategic decision-making processes, enterprises feed massive amounts of data through business intelligence (BI) systems to gain insight into the operational state of the business, identify trends, and plan for the future. But it has to be the right kind of data: too much data can be a liability if there's no way to use it efficiently, and inaccurate data increases operational risk or risk of undermining credibility of actionable data. How can an enterprise IT organization ensure its BI systems generate a clean and simplified data stream that the business can convert into valuable, actionable information? What components should be incorporated to help analyze data and present results?

Like many other IT organizations in large enterprises, the Service Deployment & Operations organization (SDO) within Microsoft Information Technology (Microsoft IT) manages over 1,000 line-of-business applications and more than 30,000 servers. The organization uses a large number of different monitoring mechanisms to identify outages, to flag potential service failures, and to provide BI data to internal customers for analysis.

Within the organization, many semi-independent teams build and maintain these apps using standards that are best suited to their own business needs. Because of this, BI has an additional challenge to bring all the data together into a standardized model. The resulting assemblage of systems and tools presented a number of challenges for those who wanted to produce accurate and actionable BI reports:

- **Disconnected and scattered data:** The lack of a true end-to-end integration of the various systems limited the organization's ability to provide reusable data to multiple groups. Data came from many different operational systems and scanning engines, and each data store used different processes because they corresponded to different transactional systems for different business units. People sometimes didn't know who to contact to learn about the processes and to obtain permission to access the data. It could also take time to understand the differences between two different sources that stored similar data.
- **Lack of adherence to standards:** Each BI development effort was driven by a team's need to support a specific business requirement as quickly as possible. Although this helped the business evolve quickly, the accelerated development process meant that teams occasionally would not adhere to standards concerning how to treat the data or how to use it. This lack of adherence to standards also made it difficult to communicate effectively about the data. Different analysts would devise different technical explanations for the data; other teams sometimes had to invent different methods to reproduce the results when they didn't have access to the same data stores.
- **Unscalable:** Because the different data sources had different objectives when they were created, they were implemented differently. Some didn't provide robust reporting capabilities; others wouldn't expose the same types of data—or might not have any capacity for growth. Transformations between data sets were also difficult to perform when each team's data was structured differently and based on different schemas.

Solution

As part of Microsoft IT's efforts to support the company's *One Microsoft* initiative to provide end-to-end views of key internal processes, SDO decided to federate some of its BI data sources into a BI portfolio. In this case, *federation* refers to interconnecting the existing data sources into an overall solution in order to present a single, clean data stream to internal customers without requiring any large-scale migration of data. This BI portfolio would also add value by including intuitive dashboards and scorecards to help users analyze their data quickly and accurately.

System Design Principles

SDO followed these key system design principles when developing its BI portfolio:

- **Leverage original data sources and existing investments by federating current BI systems.** SDO wanted to ensure that data was obtained from original data sources only; the system must not have chained dependencies (where a data source that was used actually pulled its information from another source). The business logic for interpreting data needed to be implemented just once in order to minimize the need for teams to learn and re-implement the business logic, to reduce the potential for the logic to start morphing, and to help drive accountability to the business process owners. The system design also needed to leverage existing BI investments, so federating the disparate data sources into acting as a cohesive whole was a must.
- **Develop a framework for the future that enhances the current federated portfolio with reusable, extensible, and automated platform that is based on industry best practices.** For SDO, this meant designing templates based on industry best practices that:
 - Model data around business processes (not just operational systems data sources)
 - Make it easier for people to access and consume data
 - Make the data secure
 - Assure data quality
 - Are optimized for analytics and reporting
- **Create “best of” data sets by leveraging multiple sources for key data points.** By leveraging multiple sources for key data points, data can be merged and analyzed for quality and completeness. The system must be able to pull all the data sets together, and then perform data forensics to ensure data quality when required (such as when two sources of data show one value, but a third source shows a different value).
- **Use smart data strategies to minimize refresh latencies and to aggregate and converge data sets.** Providing a standardized data delivery model enables SDO to establish a common metrics platform for usability, including:
 - Common business logic used to decide how to match and integrate the sources
 - Enabling data in the model so that the reports can utilize common filters
 - Extended metadata, including information on deviations, first, min, max, percentiles, outliers, and more.
 - Common client access technologies, such as repositories with common SQL APIs, standard reports, dashboards, and scorecards.
 - Synchronize with sources' update cadence to enable near real-time data
 - Smart aggregation strategies that deliver optimized, high user query performance
 - Ultimately create a central location for management of metrics and KPIs
- **Deliver with partner and SME collaboration.** Subject matter experts (SMEs) must be involved in the data model design decisions in order to ensure that all business logic is correctly captured and implemented. As the owners of the business provide definitions for how to interpret data and how to calculate KPIs and metrics, SMEs must provide those requirements and be involved in the validation process as part of solution testing.

Development Methodology

SDO followed an Agile process methodology for development, which is an industry-standard set of values, principles, and practices that emphasizes an iterative approach to development in order to realize business value quickly. When compared to the traditional Waterfall method of starting a

project with a single prioritization and working linearly to deliver the end product, Agile's initial prioritization produces a roadmap with the work divided into actionable chunks called *sprints* that deliver incremental solutions. The start of each sprint is an opportunity to re-evaluate priorities, ensuring that the next set of deliverables maps to any shift in business direction. Agile's ability to continually adjust the development pipeline helps ensure that the end product addresses the changing needs of its customers and delivers a positive return on investment.

Agile vs. Waterfall

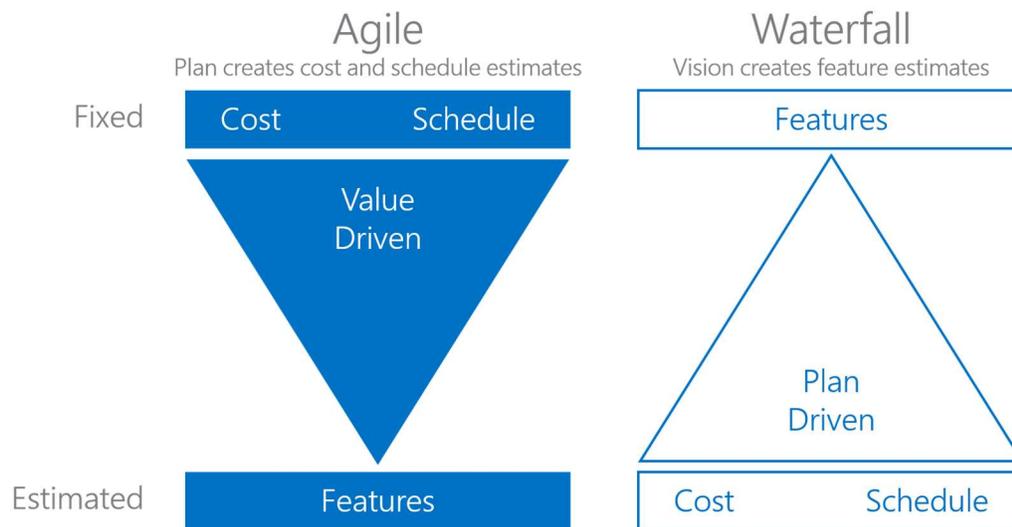


Figure 1. Comparing Agile and Waterfall development methodologies.

Solution Architecture

Improving BI processes is an ongoing effort for SDO. In this section, we compare the federated solution that is currently in use with the long-term strategic architecture that is under development.

Current Federated Solution Architecture

The current BI environment is a portfolio of a number of small, stand-alone BI solutions that have been implemented by different groups over time. SDO has federated these separate BI systems together through a single portal. This "one stop shop" view simplifies user access and reporting, even though the solutions are still technically discrete systems.

As shown in Figure 2, the current system is designed to comply with the *Leverage original data sources and existing investments* design principles by pulling data from a portfolio of federated systems—including multiple sources for the same data. As the data is being pulled in, duplicate data items are compared for accuracy. Where there are gaps, data items from secondary sources are filled in.

After the data is fully assembled, the data is moved to an operational data store where reporting views are created. From the operational data store, data is also used to populate a dimensional model for daily historical and trend reporting.

The current federated solution contains a number of different analysis engines. Each engine analyzes the real-time and historical data stream for different business process—making available facts,

metrics, and report view details to service owners, program managers, subject matter experts (SMEs), and end users.

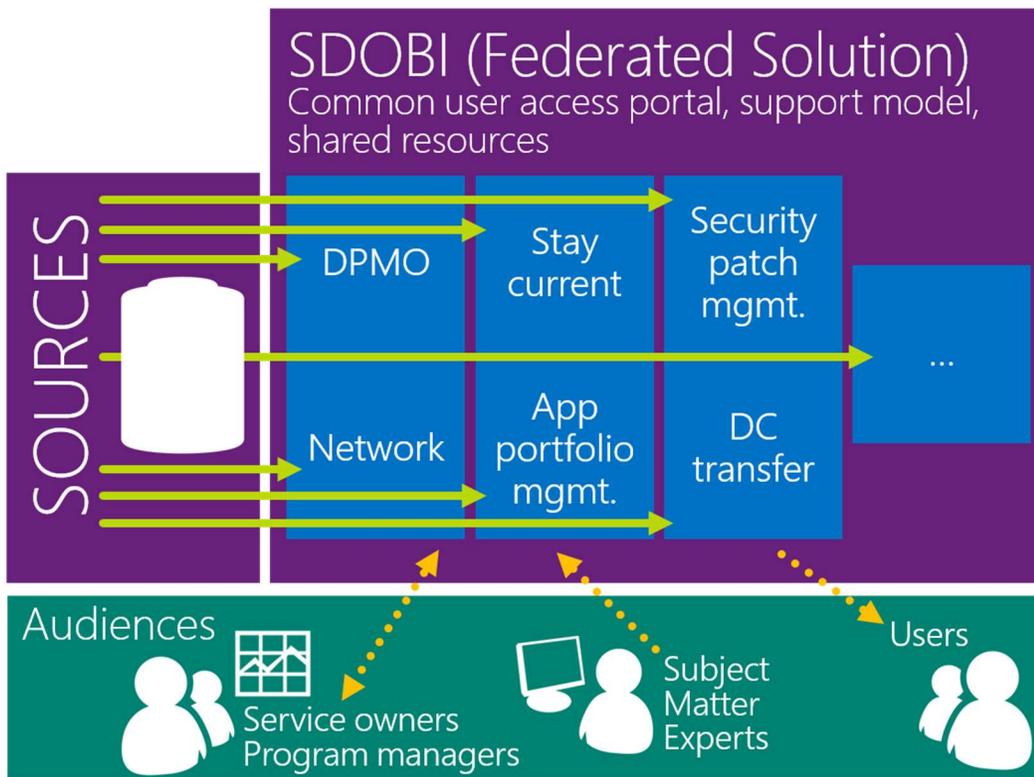


Figure 2. The current federated BI solution.

Today's federated solution enables SDO to re-purpose its existing BI systems and to provide an easy-to-access portal for users. Using it, SDO can maintain its current BI environment while simultaneously evolving the BI portfolio into a more strategic end-state that addresses longer-term business priorities (see below).

Strategic Solution Architecture

In contrast to the previous section that discussed the current BI environment,

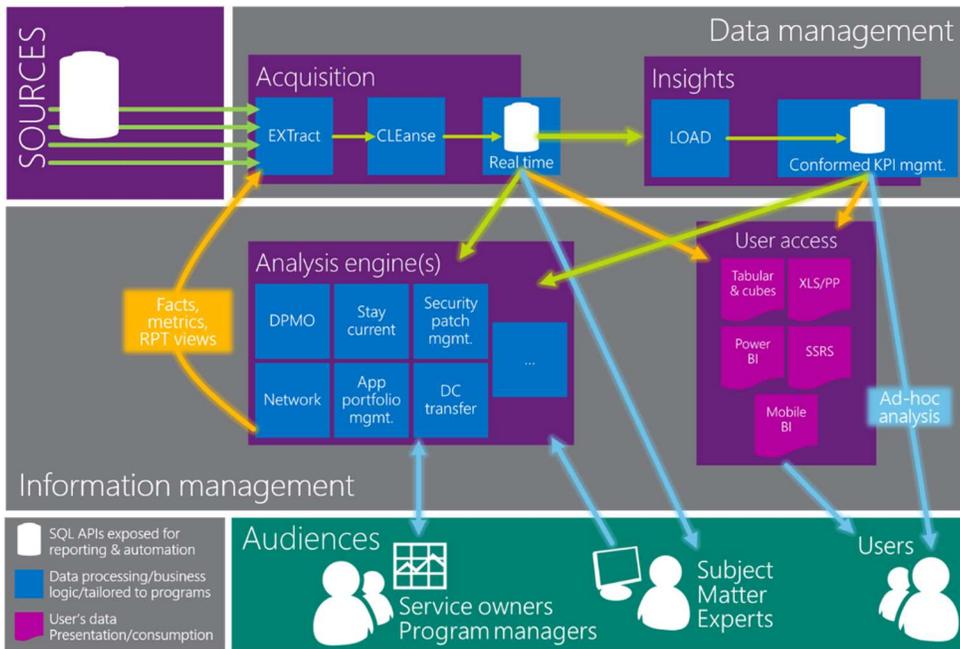


Figure 3 illustrates the strategic end-state that SDO envisions for its BI solution. This architecture that is currently under development contains two layers, each with a set of interoperating components. The solution’s layers and components are described in the more detail in the following sections.

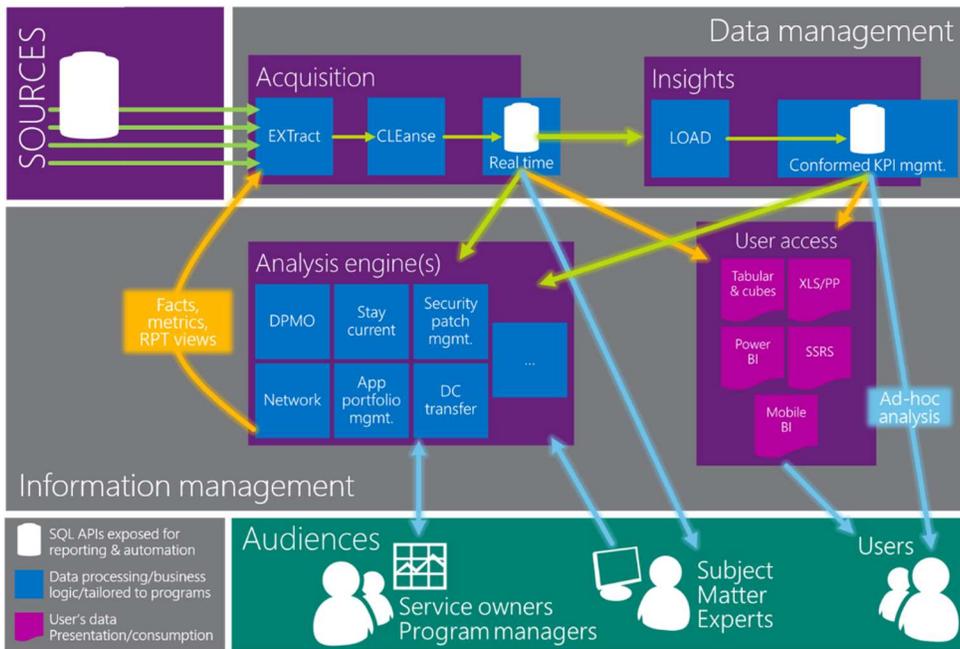


Figure 3. Architectural diagram of the BI Portfolio at its planned end-state.

Data Management Layer

The data management layer addresses the *Create “best of” data sets by leveraging multiple sources for key data points* design principle by using data forensics to clean the data stream. This layer then makes both near-real-time and historical information available for people’s use in a variety of outputs, including reports, scorecards, and dashboards.

- **Acquisition component:** This component acts as the system's federated data library, processing the data that streams in from many different federated sources to ensure data quality. Acquisition also includes automated business rules to help ensure that SDO implements the correct logic only once. It links, consolidates, and curates the data that is extracted, processed, and exposed to users. This extracted and cleansed data is stored in a database for near-real-time access by authorized users. Historical data is passed to the Insights component on a daily basis (see below).
- **Insights component:** This component acts as the system's key performance indicator (KPI) repository, holding historical data and trend information. The component then takes this information and combines it with the cleaned real-time data stream to be processed with analysis engines or to be presented in user reports.

Information Management Layer

This layer addresses the *Develop a framework that is reusable, extensible, automated, and based on industry best practices* design principle. It shares data that, for example, enables Problem Management to do sandbox development of data models that analyze and assess risks detected in the IT environment. SDO's SMEs use Excel Power Pivot to perform initial *ad hoc* model development. When a useful model or view is prototyped, the Excel file is handed to the BI group to create a permanent, automated report using SQL Server Reporting Services (SSRS), high volume logic in Tabular/Cubes (SSAS), and/or Power BI solutions.

- **Analysis Engine(s) component:** This component embeds the current federated system while enabling new development as pluggable components. In this new strategic architecture, the federated model's components are repurposed, becoming abstracted as independent Analysis Engines within the new framework. Additional components developed in the future will follow the new design, interoperating with the legacy ones through the new architecture. This hybrid model enables SDO to utilize both the old and the new, and optimizing resources by re-building older components only when the business requires a significant enhancement or change in functionality. This modular and flexible component supports a number of different analysis engines. Each engine corresponds to a different business process or program, analyzing the real-time and/or historical data stream, sending facts, metrics, and report view details back into the Data Management layer's Acquisition component to reintegrate the additional data produced into the common stream. This component's extensible design is evidenced by the variety of engines SDO has developed to analyze the data. Some examples of engines (related to key business processes/initiatives) that are currently in use include:
 - **Stay Current:** An analysis engine that compares scanning data from servers with standards, and thus assessing KPIs measuring deviations from policies and strategies (applies to OS, SQL, drivers, agents, etc).
 - **Security Patching Management (SPM):** This engine assess the security risk of not having security patches deployed on servers.
 - **Defects Per Million Opportunities (DPMO):** Derived from Six Sigma industry standards, this analysis engine automates calculations of KPIs measuring potential operational risks on servers, in order to ensure continuous improvements towards minimizing IT risks over time.

In order to perform these functions, each engine contains the following optional subcomponents:

- **MDM:** Master Data Management (MDM) subcomponent that stores information needed by the engine but which can't be discovered programmatically, such as classification of data needed to group elements in reports, which don't exist on any transactional data source available to SDO (complements the business process' transactional systems at a higher business analytical level). The service owners related to the business process that the engine automates are responsible for maintaining these MDM reference data.
- **Logic:** Logic about how to link sources, consolidate, and aggregate data.

- **KPIs:** Logic about how to calculate measures, metrics, and comparison to business targets.
- **Reports:** Specific views of the data needed by this business process in each particular data engine.
- **User Access component:** This component consumes data from the previously described layers and presents information to users. The User Access interfaces are embeddable in portals, which simplifies how to review automated reports built from the two data stores. This component provides analysis service cubes, tabular models, and data mining models, allowing users to work with Excel, Power BI, SSRS, tabular models, and other reporting tools.

Audiences

The strategic architecture is designed to support the needs of different users' roles (people may perform several roles at different points of the business process), including:

- **Service owners and program managers:** These roles drive accountability for the business processes, and thus define all logic, rules, conditions, and reference values that the solution must manage, implement, and deliver. They are the stakeholders and drivers of the BI solution from the business.
- **SMEs:** Subject Matter Experts (SMEs) are people whose roles implement and/or execute the business process. They perform tasks, operate systems, monitor execution, drive improvements, analyze results, and drive decision making. They access the real-time data store directly in order to perform operational task execution, troubleshooting, or *ad hoc* analytics.
- **End users:** End users can access the data and run different reports through a portal interface, or they can perform *ad hoc* analyses by accessing the historical data store directly. Key roles in the end-user community include management (such as team managers, business managers, and VPs), function managers (including incident managers and problem managers), and business analysts. These various roles consume information from cubes, canned reports, dashboards, scorecards, portals, visualizations, SQL APIs, and mobile consumption solutions. They can also produce their tailored *ad hoc* Power BI analytics from these data.

Benefits

By implementing a new BI framework with a strong foundation on industry best practices which supports evolving from a federated BI portfolio, SDO was able to gain the following benefits:

- **Increased efficiency:** By building a BI portfolio that federates existing BI Solutions, SDO is able to present users with a single, clean, easier-to-use data access and stream. Moreover, federating the existing solutions means that SDO can manage the BI capabilities without needing to undergo any large-scale data or process migration.
- **More scalable platform:** Basing the BI portfolio on a federated model also simplifies the process of onboarding additional data sources, which significantly enhances the system's capacity for accelerated growth.
- **Adherence to standardized processes:** The new BI portfolio organizes and presents data in a known, consistent fashion. Now, different groups have access to the same data models and can base their analyses off of a known, clean data stream. Furthermore, the new system enables SDO to cross link related data to produce more complex information, such as measuring organizational compliance to standards by making correlations between compliance of a server to the amount of support tickets it generates.
- **Unified data retention:** The solution provides a credible historical account of business events that users feel comfortable to consume to drive business decisions.

TIP: For detailed information about the business processes that SDO uses to leverage the BI reporting, see the [companion business case study in Microsoft IT Showcase](#).

Best Practices

When designing and implementing a new federated BI portfolio, SDO developed and implemented these best practices:

- **Develop a framework that is reusable, extensible, automated, and based on industry best practices.** For SDO, this meant ensuring that every component used to build out the BI platform is reusable: database schemas, design patterns, ETL (Extract, Transform, Load) packages, logging, auditing, and more. Although such an approach requires some time at the beginning to set up, having standards and the practices in place enables SDO to reuse and adapt much more quickly than could be done without this type of framework. SDO made extensive use of the Kimball Group's methodology and best practices for Dimensional Modeling that offers a full life cycle for managing BI solutions and detailed guidance on designing an ETL architecture.

TIP: For more information about the Kimball Group, visit <http://www.kimballgroup.com/>.

- **Federate instead of consolidate.** Avoid re-inventing the wheel: if your current BI systems are adequate for your current needs, federate your existing BI assets and use a single set of standards to cleanse and model the resulting data stream. This information can then be expressed as a single source to the downstream business. SDO was able to leverage its diverse set of existing data sources and federate them into a single cleaned data stream for downstream consumption.
- **Use an Agile process for development and engage with business stakeholders.** In contrast to a Waterfall methodology, following Agile development helps you speed the delivery of business value, which goes hand-in-hand with the Kimball Group's principles. Following this approach enabled SDO to partner more often with the business in all phases of the solution's development. By regularly engaging with SMEs and end users, SDO is able to continue to add analysis engines and enhance the user interface as it evolves the platform. This in turn ensures close alignment to business priorities at all times and continuously delivers incremental business value.
- **Establish agreed-upon standards.** Enforcing standards clarifies how data is expressed and communicated—and ultimately enhances reusability. As an example, when SDO receives a new requirement from a different business group, the standards-based platform streamlines the development and shortens the time required for the business to receive the new deliverable. For example, when building the packages that extract and clean the data, SDO can use templates that have been developed and make a few adjustments to connect to the new data source—a much faster process than when everything had to be custom-designed.
- **Deliver trusted data sets, integrate and bridge data, metrics, share mastered data.** If your data resides in multiple sources, develop a process to clean and simplify it before you expose it to any downstream BI processes. For each data source, SDO works with the appropriate team to optimize and extract the data to ensure it presents the right level of information for the user.
- **Provide an extensible, user-friendly access layer that simplifies ad-hoc analytics and reporting.** To better support how users work with BI, SDO has developed a centralized portal that provides a single point of access to all BI information. In this portal, reports and data access are organized by business process, and each process includes contextual explanations of its function, including FAQs, definitions, and related information.

Conclusion

Because SDO has developed a federated BI portfolio that is based on industry best practices, SDO is able to present a *hierarchy of truth* where internal customers have a single authoritative source for all

BI data. The reusable, extensible, and automated nature of the solution extends its capabilities beyond simple reporting; in fact, this platform advances the organization's capability for data-driven decisions. This in turn helps Microsoft IT drive a zero-defect culture by reducing the risk of defects and improving service delivery throughout the company.

SDO continues to refine and release updates to the BI portfolio on a monthly basis. As the business continues evolving over time, all mature data and information systems will be fully supported, feeding their key business processes into the BI portfolio. SMEs will perform their experimentation, proof of concept, and piloting in a supported sandbox. As new BI implementations are built on top of the new framework, federated systems will eventually be phased out—resulting in a true consolidated BI platform.

Resources

[Microsoft IT improves business processes with a federated Business Intelligence Portfolio](#)

[Microsoft IT](#)

[The Kimball Group](#)

For More Information

For more information about Microsoft products or services, call the Microsoft Sales Information Center at (800) 426-9400. In Canada, call the Microsoft Canada Order Centre at (800) 933-4750. Outside the 50 United States and Canada, please contact your local Microsoft subsidiary. To access information via the World Wide Web, go to:

<http://www.microsoft.com>

<http://www.microsoft.com/microsoft-IT>

© 2014 Microsoft Corporation. All rights reserved. Microsoft and Windows are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. The names of actual companies and products mentioned herein may be the trademarks of their respective owners. This document is for informational purposes only. MICROSOFT MAKES NO WARRANTIES, EXPRESS OR IMPLIED, IN THIS SUMMARY.