

MAP Your Future

A Modern Analytics Platform
for Growth and Change

The logo for DSI, consisting of the letters 'DSI' in white, bold, sans-serif font, centered within a solid green square.

The Modern Analytics Platform must be agile.

Gone are the days where implementing business intelligence involved a rigid chain of requests to a small team of specialists who extract data from an on-premise, centralized environment. This traditional approach to analytics risks alienating the business and stifling innovation.

The growth and capabilities of analytics tools have accelerated rapidly to meet the needs of information workers. The first wave of analytics tools was very much focused on enabling business users to create their own ad hoc analysis from governed data sources. Information workers can now load, manipulate and analyze millions of rows of disparate data without the support of the information technology (IT) organization. These capabilities are

empowering workers to venture out on their own and tackle more and more complex business problems, often circumventing security and data governance policies enforced by IT.

Data-driven organizations need a Modern Analytics Platform that empowers their information workers and satisfies the needs of a diverse set of end users. Rigid processes to deliver enterprise-wide reports to business users are being replaced by an agile architecture, driven by new data sets and new data types. The volume of data is growing as businesses capture it at a finer level of detail and therefore it is becoming increasingly difficult to predict the projected data volume even a few years out. The variety of data sources is also multiplying – data from social media

and the Internet of Things (IoT) are certain to be joined by other diverse sources of structured and unstructured data. Velocity will no longer be a measure of how quickly the data is being created but how fast the requirements and sources are changing.

Data-driven organizations need a roadmap for this ever-changing landscape if they are to build the foundation for a Modern Analytics Platform. This platform must enable them to adapt dynamically to the increasing volumes and varied data sources while leveraging their existing investment in the traditional data warehouse. It must also enforce security and data governance policies.

Traditional analytics can only take us so far.

The traditional analytics platform offers many benefits and still plays a role in the Modern Analytics Platform. However, Big Data and the demands by business users for real-time intelligence and reporting have exposed its limitations.

The traditional analytics platform is still the place to centralize data integration and manage data quality. It can also support any reporting application the organization chooses and deliver consistent data across those applications without any data modelling. However, it structures data sources rigidly, and integrating new sources into the data repository takes development time and resources. There's no support for ad hoc data sources, which is essential for today's business.

By limiting themselves to this rigid structure, organizations ultimately find they cannot keep up to the data demands of the business units, who then start building their own ad hoc reporting solutions outside the governance of the IT department.

The business unit approach has its own benefits and limitations. While it can handle dynamic data sources, support purpose-specific reporting and dynamic business logic, there's no centralized integration, and each report requires complex data modelling. It also requires deep knowledge of source data, and the business logic is replicated on a report-by-report basis, thereby raising the risk of inconsistent business logic and requiring a great deal of distributed maintenance.

The most troublesome aspect of the business unit approach is it ultimately works in parallel with the traditional data warehouse, so there's no interaction between the two.



The best of both worlds.

In an ideal world, both models are combined, leveraging the strong data governance of the traditional data warehouse and combining it with the dynamic characteristics of the ad hoc model. This hybrid approach is well-suited for the velocity challenge and supported by recent enhancements to SQL Server Analysis Services (SSAS) and the introduction of Power BI within the Microsoft Analytics suite.

This Modern Analytical Platform (MAP) involves a **three-stage** approach to the maintenance and adoption of new data sources and business logic into the analytical platform.

- 1 Identification of ad hoc data source / business logic**

New data sources would initially adopt the ad hoc architecture and get mashed together with enterprise data in a presentation tool, such as Power BI. The sources and business logic are defined in the tool and assessed regularly to determine whether ad hoc analysis should be retired or formalized within the IT organization.
- 2 Enterprise adoption of new data source and business logic**

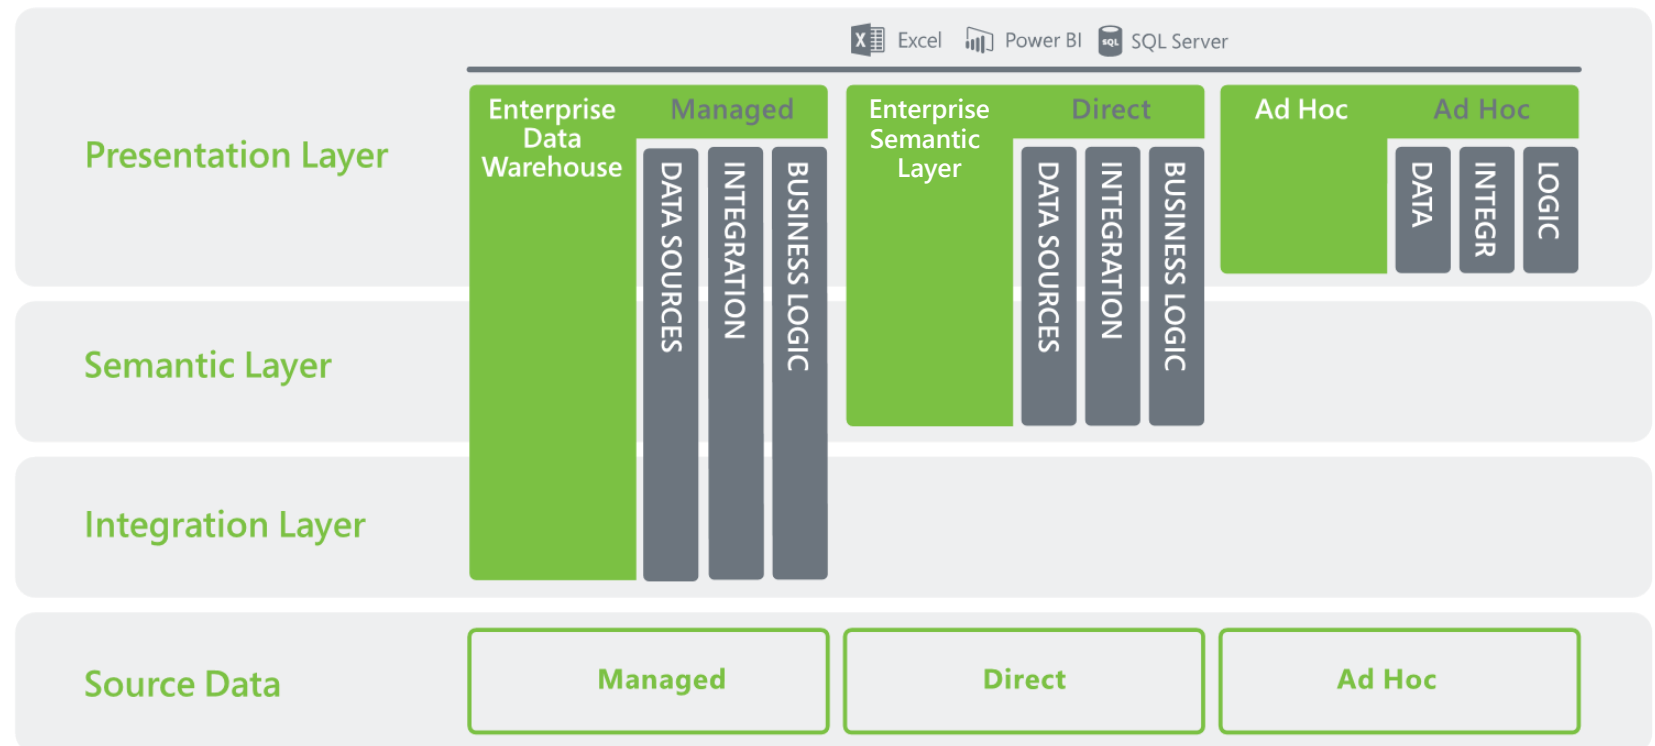
At this point, IT takes ownership. Normally, this takes months of design and development, so as an interim stage, a good approach is to update the semantic layer to address the needs of the business and to avoid a significant volume of analysis in the ad hoc stage. The full adoption process incorporates the new data source into the semantic layer model, replicating the final structure of the enterprise data warehouse as much as possible. This integration should lead to modification of all existing reports to use the semantic layer SSAS instead of the ad hoc source.
- 3 Enterprise integration of new data source and business logic**

Not all sources and logic must be incorporated, but those that should be included are ones that require integration with other data sources, complex transformations and delta processing logic. An integration process incorporates the new data source into the enterprise data warehouse, including the standard ETL / data quality processes as defined, as well as the business logic. Done correctly, this stage should not affect existing reporting or analysis, but could take months to complete depending on the priorities and schedule of the enterprise data warehouse development.

This MAP will bring us to an analytical platform consisting of the following layers.

This analytical platform enables Microsoft Power BI and SSAS to leverage the same underlying code base and data model and can be hosted on-premise or on Azure.

- 1 Enterprise data repository (Managed): This is the traditional enterprise architecture.
- 2 Enterprise semantic layer (Direct): This includes the traditional enterprise architecture, as well as interim data sources, interim integration logic and interim business logic.
- 3 Ad hoc (Ad Hoc): This includes the traditional enterprise architecture, as well as the interim semantic layer, ad hoc data sources, ad hoc integration logic and ad hoc business logic.



Data governance and security.

Overall data governance and security converge at the semantic layer and will represent the single point of access for all reporting.

The semantic layer exposes an integrated data model to the users, and no modelling is necessary when accessing this data source. It also acts as a central repository for all KPI definitions, so there's no need to replicate them in each report. The semantic layer also centrally defines the security model.

Because the semantic layer represents the single access point for all data access / reporting and analysis, the deployment of objects into this layer is critical to the enforcement of data governance and security. This means all deployments into or changes to the semantic layer will require approvals from both the business and the security office.



Building up

We must also address volume of data. Traditionally, this involves hosting on-premise and requiring advance knowledge of the nature and size of the data you will be storing and analyzing. Essentially you are guiding your hardware and software licensing purchases based on an uncertain future state that leads to expensive overprovisioning or falling short of capacity and having to re-platform. The cloud solves these issues by providing a fully extensible environment, enabling the business to purchase/rent and scale up the capacity as needed. The hybrid approach can easily be moved to the cloud with little changes to the existing components:

Component	On-Premise	Azure	Comments
Data Sources	Any Existing	Any Existing	Remain On-Premise (not part of Analytical Platform)
Integration / Prepare	SSIS On-Premise	SSIS run by Azure Data Factory	Migrate Code As-Is to Azure Data Factory running existing SSIS packages (PaaS Service)
Enterprise Data Warehouse	SQL On-Premise	Azure SQL	Migrate Database As-Is to Azure SQL (PaaS Service)
Semantic	SSAS On-Premise	Azure SSAS	Migrate Model As-Is to Azure SSAS (PaaS Service)

The result is the same hybrid model in Azure but with the benefits of Platform-as-a-Service (PaaS):

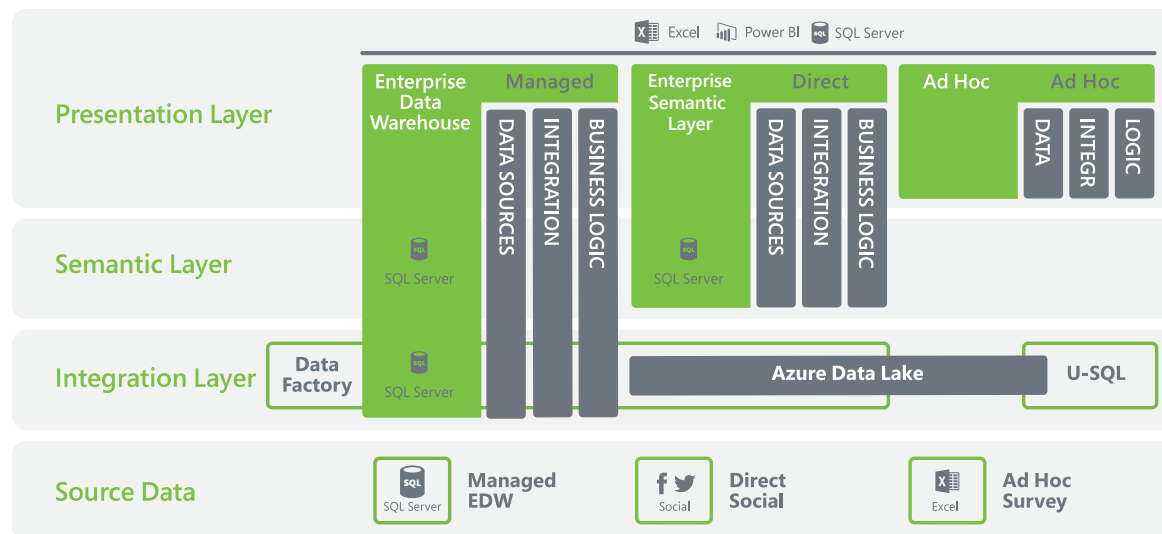
- Backups and geo-replication are fully automated
- Services are automatically updated and patched, and can be turned off or paused when not in active use
- There is no fixed hardware or software costs – you pay for what you use – and you can scale up or down services as necessary

Building out

The final leg of the MAP deals with the variety of data sources. Now additional Azure services can be integrated:

- Big Data from IoT devices, social media data, research data and other unstructured data can be loaded to the data lake
- Social media data, Azure streaming and Event Hubs can be streamed into the data lake
- Machine learning such as sentiment analysis for social media and predictive analysis for financial data can be applied

Below is an example of how a variety of customer feedback could be integrated into the Modern Analytics Platform.



This solution leverages many Azure services but most critically is Azure Data Factory which is used to orchestrate the ETL/ELT processes across multiple data sources and multiple ETL/ELT applications.

It's time to hit the road.

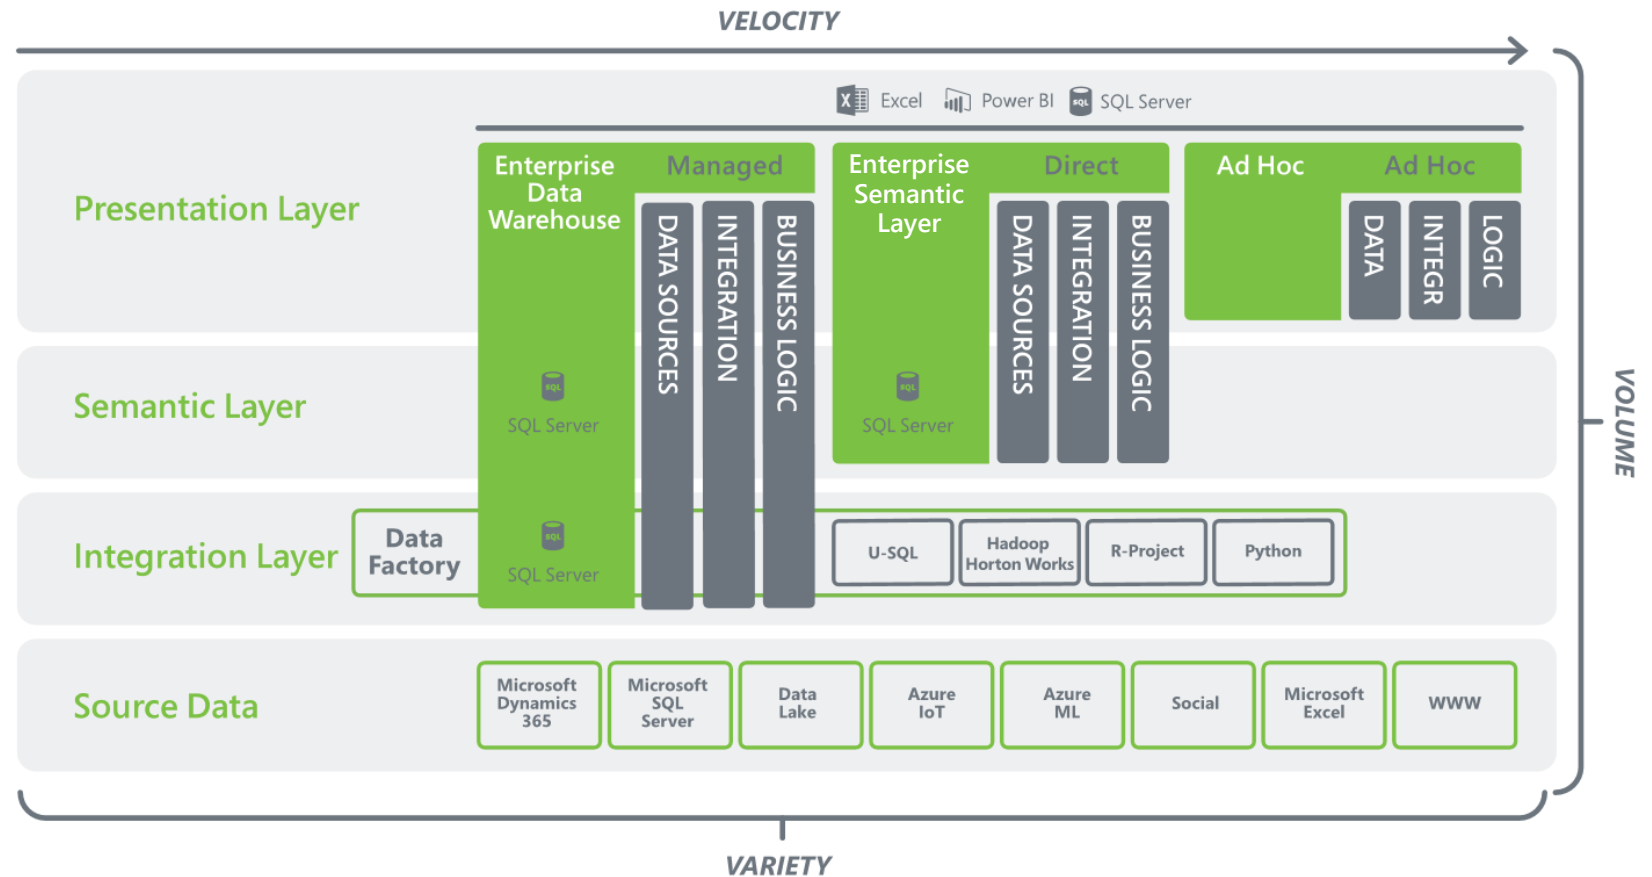
DSI can help migrate your business to this modern platform while leveraging the investment you have already put into your existing data warehouse. There's no need to start from scratch.



Evan Ross, Principal & CTO

To explore these opportunities for your organization, let DSI provide a free half day consulting workshop. Get in touch with us at sales@dstrat.com.

This MAP provides a clear path to move your analytical platform from a traditional, on-premise enterprise data warehouse to a Modern Analytics Platform that deftly handles the challenges presented by the variety, volume and velocity of data requirements, as depicted below





Microsoft
Partner



- Gold Application Development
- Gold Cloud Platform
- Gold Data Analytics
- Gold Datacenter
- Gold Data Platform