Tableau for the Enterprise:

An Overview for IT

Authors: Marc Rueter, Senior Director, Strategic Solutions

Ellie Fields, Senior Director, Product Marketing

May 2012

Introduction

A new generation of business intelligence software puts data into the hands of the people who need it. Slow, rigid systems are no longer good enough for business users or the IT teams that support them. Competitive pressures and new sources of data are creating new requirements. Users are demanding the ability to answer their questions quickly and easily. And that's a good thing.

Tableau Software was founded on the idea that data analysis and subsequent reports should not be isolated activities but should be integrated into a single visual analysis process—one that lets users quickly see patterns in their data and shift views on the fly to follow their train of thought. Tableau combines data exploration and data visualization in an easy-to-use application that anyone can learn quickly. Anyone comfortable with Excel can create rich, interactive analyses and powerful dashboards and them share them securely across the enterprise. IT teams can manage data and metadata centrally, control permissions and scale up to enterprise-wide deployments.

This overview is designed to answer questions common to IT managers and administrators and help them support Tableau deployments of any size. In this document we cover:

- Tableau Architecture
- Deployment Models
- Security
- Scalability
- System Administration
- Data Strategy
- Metadata Management
- Mobile Deployment

Architecture

Tableau has a highly scalable, n-tier client-server architecture that serves mobile clients, web clients and desktop-installed software. Tableau Desktop is the authoring and publishing tool that is used to create shared views on Tableau Server.

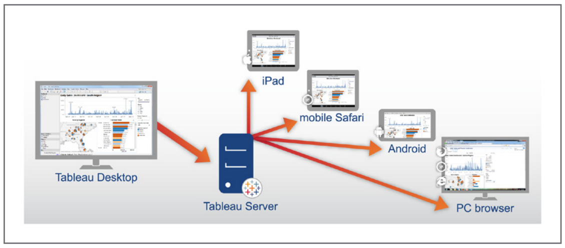

Figure 1: *Tableau provides a scalable solution for creation and delivery of web, mobile and desktop analytics.*

Tableau Server is an enterprise-class business analytics platform that can scale up to hundreds of thousands of users. It offers powerful mobile and browser-based analytics and works with a company's existing data strategy and security protocols. Tableau Server:

- Scales up: Is multi-threaded
- Scales out: Is multi-process enabled
- Provides integrated clustering
- Supports High Availability
- Is secure
- Runs on both physical and Virtual Machines

The following diagram shows Tableau Server's architecture:

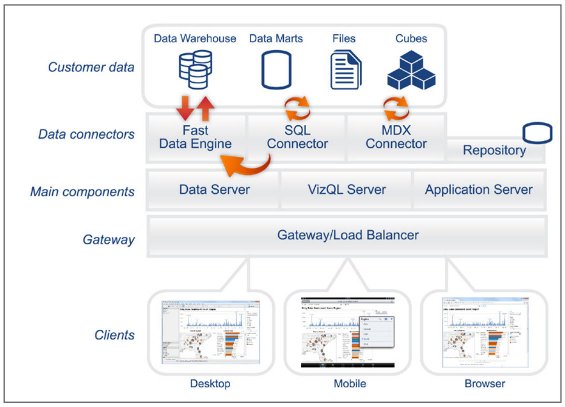

Figure 2: *Tableau Server architecture supports fast and flexible deployments.*

Below, we explain each of the layers, starting with customer data.

Data Layer

One of the fundamental characteristics of Tableau is that it supports your choice of data architecture. Tableau does not require your data to be stored in any single system, proprietary or otherwise. Most organizations have a heterogeneous data environment: data warehouses live alongside databases and Cubes, and flat files like Excel are still very much in use. Tableau can work with all of these simultaneously. You do not need to bring all your data in-memory unless you choose to. If your existing data platforms are fast and scalable, Tableau allows you to directly leverage your investment by utilizing the power of the database to answer questions. If this is not the case, Tableau provides easy options to upgrade your data to be fast and responsive with our fast in-memory Data Engine.

Data Connectors

Tableau includes a number of optimized data connectors for databases such as Microsoft Excel, SQL Server, Oracle, Teradata, Vertica, Cloudera Hadoop, and many more. There is also a generic ODBC connector for any systems without a native connector. Tableau provides two modes for interacting with data: Live connection or In-memory. Users can switch between a live and in-memory connection as they choose.

Live connection: Tableau's data connectors leverage your existing data infrastructure by sending dynamic SQL or MDX statements directly to the source database rather than importing all the data. This means that if you've invested in a fast, analytics-optimized database like Vertica, you can gain the benefits of that investment by connecting live to your data. This leaves the detail data in the source system and send the aggregate results of queries to Tableau. Additionally, this means that Tableau can effectively utilize unlimited amounts of data – in fact Tableau is the front-end analytics client to many of the largest databases in the world. Tableau has optimized each connector to take advantage of the unique characteristics of each data source.

In-memory: Tableau offers a fast, in-memory Data Engine that is optimized for analytics. You can connect to your data and then, with one click, extract your data to bring it in-memory in Tableau. Tableau's Data Engine fully utilizes your entire system to achieve fast query response on hundreds of millions of rows of data on commodity hardware. Because the Data Engine can access disk storage as well as RAM and cache memory, it is not limited by the amount of memory on a system. There is no requirement that an entire data set be loaded into memory to achieve its performance goals.

Tableau Server Components

The work of Tableau Server is handled with the following four server processes:

Application Server: Application Server processes (wgserver.exe) handle browsing and permissions for the Tableau Server web and mobile interfaces. When a user opens a view in a client device, that user starts a session on Tableau Server. This means that an Application Server thread starts and checks the permissions for that user and that view.

VizQL Server: Once a view is opened, the client sends a request to the VizQL process (vizqlserver.exe). The VizQL process then sends queries directly to the data source, returning a result set that is rendered as images and presented to the user. Each VizQL Server has its own cache that can be shared across multiple users.

Data Server: The Tableau Data Server lets you centrally manage and store Tableau data sources. It also maintains metadata from Tableau Desktop, such as calculations, definitions, and groups. The published data source can be based on:

- A Tableau Data Engine extract
- A live connection to a relational database (cubes are not supported)

Read more about the Data Server in the section Data Strategy below.

Backgrounder: The backgrounder refreshes scheduled extracts and manages other background tasks.

Gateway/ Load Balancer

The Gateway is the primary Tableau Server that routes requests to other components. Requests that come in from the client first hit the gateway server and are routed to the appropriate process. If multiple processes are configured for any component, the Gateway will act as a load balancer and distribute the requests to the processes. In a single-server configuration, all processes sit on the Gateway, or primary server. When running in a distributed environment, one physical machine is designated the primary server and the others are designated as worker servers which can run any number of other processes. Tableau Server always uses only one machine as the primary server.

Clients: Web Browsers and Mobile Apps

Tableau Server provides interactive dashboards to users via zero-footprint HTML and JavaScript (AJAX) in a web browser, or natively via a mobile app. No plug-ins or helper applications are required. Tableau Server supports:

- Web browsers: Internet Explorer, Firefox, Chrome and Safari
- Mobile Safari: Touch-optimized views are automatically served on mobile Safari
- iPad app: Native iPad application that provides touch-optimized views and content browsing
- Android app: Native Android application that provides touch-optimized views and content browsing
- Android browser*:* Touch-optimized views are automatically offered in the Android browser

Clients: Tableau Desktop

Tableau Desktop is the rapid-fire authoring environment used to create and publish views, reports and dashboards to Tableau Server. Using Tableau Desktop, a report author can connect to multiple data sources, explore relationships, create dashboards, modify metadata, and finally publish a completed workbook or data source to Tableau Server. Tableau Desktop can also open any workbooks published on Tableau Server or connect to any published data sources, whether published as an extract or a live connection.

Deployment Models

Tableau can be configured in a variety of ways depending on your data infrastructure, user load and usage profile, device strategy, and goals. Tableau Server can be clustered with any number of machines. Below are six examples of common configurations.

Simple Configuration

For many customers, a single server with the recommended minimum hardware configuration, 8 CPU cores and 32GB of main memory, will provide good performance. This type of configuration is useful for a proof of concept for a larger deployment, or for a departmental server. Tableau recommends running two instances of each major process: Data Server, Application Server, VizQL Server and Backgrounder on a single server 8-core deployment of Tableau Server.

3-Server (24-Core) Cluster

Environments with heavier user load will require clustering additional servers. In a 3-Server configuration, the Gateway or primary machine will host the Backgrounder, Repository and Extract Host, and will send Application Server requests to the worker machines. An administrator can configure the number and type of processing running in the system to support heavy or light extract usage and other characteristics.

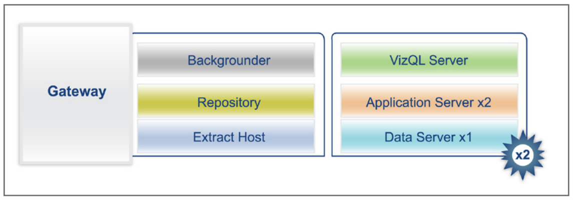

Figure 3: *Configuring of a simple 3-server cluster.*

5-Server 40-Core Cluster

More worker machines can be added to a cluster to support heavier data usage or higher user load. In a larger cluster using data extracts, you might choose to isolate the repository and extract host on one machine, the backgrounders on another, and let the VizQL and Application servers reside on the other worker machines. Different configurations are available to support different workload profiles.

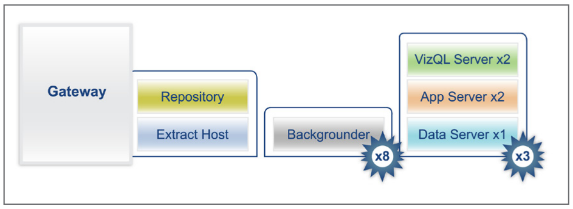

Figure 4: *Configuration of a 5-server cluster optimized for many data extracts.*

High Availability Cluster

Tableau's High Availability feature helps IT organizations meet SLAs and minimize downtime. Tableau's High Availability solution provides automatic failover capabilities for the repository and data engine components. A minimum of 3 nodes are required in a High Availability environment. A primary node serving as the Gateway and load balancer and 2 additional nodes hosting the active processes. Gateway failover is a manual step. On failover, Tableau Server sends email alerts to specified administrators.

Virtual Machine or Cloud-Based Deployment

There are no special considerations when running Tableau Server on virtual machines or in a cloud deployment. Virtual machines can be used to virtually limit the number of cores available to Tableau Server or to provide disaster recovery via the virtual machine itself. When running Tableau in the cloud, bear in mind that Tableau Server requires static IP addresses.

Security

As organizations make more data accessible to more people, information security becomes a critical concern. Tableau Server provides comprehensive security solutions that balance the variety of sophisticated requirements with easy implementation and use.

Tableau's enterprise-level security features manage authentication, permission, data and network security. Together, these capabilities provide a complete security solution that will serve the needs of a broad and diverse user base, whether internal to the organization or external on the Internet. In fact, Tableau Server has passed the stringent security requirements of customers in the financial services, government, and healthcare sectors.

Authentication – Access Security

The first level of security is to establish the user's identity. This is done to prevent unauthorized access and to personalize each user's experience. This process is typically referred to as 'authentication'. It should not be confused with 'authorization' which is covered in the section titled 'Permissions – Object Security'.

Tableau Server supports three types of authentication plus the option to allow anonymous (un-authenticated) access to the system:

1. Microsoft Active Directory,
2. Local authentication managed by Tableau Server,
3. Trusted authentication that creates a trusted relationship between Tableau Server and one or more web servers.

Tableau provides automatic login timeouts that can be configured by administrators.

Roles and Permissions - Object Security

In Tableau, a role is a set of permissions that is applied to content to manage how users and groups can interact with objects such as projects and published content. Published content such as data sources, workbooks, and views, can be managed with permissions for the typical actions of view, create, modify, and delete. Projects control the default permissions for all workbooks and views published to the project. Administrators can create groups such as "Finance Users" to make permission management easier. The use of projects can be used on a single server where support for multiple external parties (multi-tenants) is needed.

Roles provide a default permission structure to differentiate users. For example, a user may be assigned the role of Interactor for a particular view, but not for all content. And, a user with a Viewer role can see a particular view but does not have the ability to change the view. There are over 20 parameterized customizations available to help manage object security. These role-based permissions do not control what data will appear inside of a view.

Data – Data Security

Data security is becoming increasingly important, especially for organizations needing to meet regulatory requirements or who deliver content externally. Tableau offers flexibility in helping organizations meet their data security needs in three different ways: implement the security solely in the database, implement security solely in Tableau, or create a hybrid method where user information in Tableau Server has corresponding data elements in the database.

When a user logs into the Tableau Server, they are not logging into the database. This means that Tableau Server users will need to have credentials to log in to the database in order for the database level security to be applied for them. These login credentials can be passed using Windows Integrated Security (NT Authentication), embedding the credentials into the view when published, or prompting for specific user credentials.

Tableau also provides a User Filter capability that can enable row-level data security using the username, group, or other attributes of the current user. The filter appends all queries with a 'where' clause to restrict the data and can be used with all data sources.

Network – Transmission Security

For many internal deployments, network security is provided by preventing access to the network as a whole. However, even in these cases it is important to securely transmit credentials across the network. For external deployments, transmission security is often critical to protect sensitive data, credentials and to prevent malicious use of Tableau Server.

There are three main network interfaces to the Tableau Server, though Tableau pays special attention to the storage and transmissions of passwords at all layers and interfaces.

1. The client-to-Tableau Server interface defaults to standard HTTP requests and responses but can be configures for HTTPS (SSL) with customer supplied security certificates.
2. Tableau Server-to-database uses native drivers whenever possible and uses generic ODBC adapters when native drivers are not available.
3. Secure communication between Tableau Server components is only applicable in distributed deployments and is done using a stringent trust model to ensure each server receives valid requests from other servers in the cluster.

In addition to these network interface security capabilities, Tableau provides additional safeguards. There are a variety of encryption techniques to ensure security from browser to server tier to repository and back, even when SSL is not enabled. Tableau also has many built-in security mechanisms to help prevent spoofing, hi-jacking, and SQL injection attacks, and actively tests and responds to new threats with monthly updates.

Scalability

Tableau Server is highly scalable, serving the largest enterprises and up to tens of thousands of users. General Motors, Wells Fargo, eBay and Bank of America are some organizations that are using Tableau. Ray White, a large real estate company, uses Tableau to serve reports to 10,000 real estate agents.

Since 2009 Tableau Server has been running at a high scale in Tableau's own data centers to power Tableau Public, a free service for online visualization of public data. Tableau Public supports over 20 million distinct users, and as of April 2012 was serving up over 800,000 views per week. In fact, Tableau Public hit the record of over 94,000 views in one hour in late 2010.

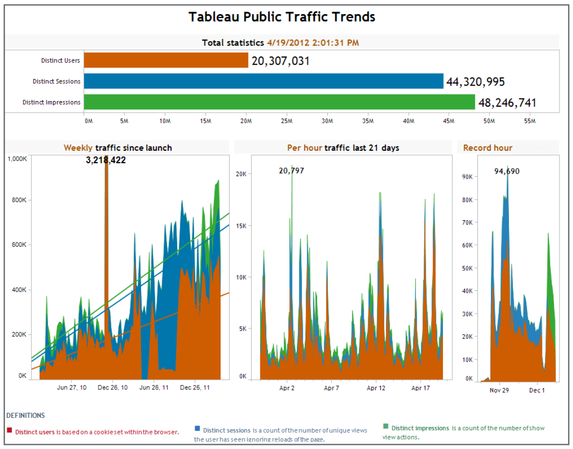

🟥 Distinct users is based on a cookie set within the browser.

🟦 Distinct sessions is a count of the number of unique views the user has seen ignoring reloads of the page.

🟩 Distinct impressions is a count of the number of show view actions.

Figure 5: *Tableau Server has scaled up to extremely high loads as the infrastructure for Tableau Public.*

Every environment is unique and there are many variables that impact performance. Factors that affect the scalability of a Tableau deployment include:

- Hardware considerations: Server type, disk speed, amount of memory, processor speed, and number of processors. More is better.
- Architecture: Number of servers, architecture design, network speed/traffic, data source type, and location.
- Usage: Workbook complexity, concurrent user activity, and data caching.
- Software configuration: Configuration settings of Tableau Server.
- Data: Data volumes, database type, and database configuration.

We periodically perform scalability tests on Tableau Server. Ask your Tableau Account Manager for the most recent scalability test results.

Environmental Optimizations

A single machine will be the fastest configuration for small user counts as adding additional machines introduces latency between the machines. The exact load at which distributed configuration becomes more important will differ based on the load type, network speed, hardware, database performance, and other variables.

For optimal performance, it is best to set the number of application and VizQL™ processes to be equal.

Best Practice Optimizations

In addition to an environment that is optimally designed, there are best practices that can be used to greatly improve performance and reduce the average response time.

Use extracts – Extracts store data locally so that users can access the data without making direct requests to the database. They can be easily filtered when users don't need the detail, significantly improving response time. If extract databases become too large, they can be offloaded to a local extract engine on a separate machine.

Schedule updates during off-peak times – Often, data sources are being updated in real-time but users only need data daily or weekly. Scheduling extracts for off-peak hours can reduce peak-time load on both the database and Tableau Server.

Avoid 'expensive' operations during peak times – Logging in and publishing, especially large files, are two very resource-consuming tasks. While it may be difficult to influence login behavior, it's often easy to influence publishing behavior. Ask users to publish during off-peak hours, avoiding busy times like Monday mornings.

Cache views – As multiple users begin to access Tableau Server, the response time will initially increase due to the contention for shared resources. However, as each request comes into the system, the view can be cached and renders much more quickly when the next user requests the same view.

System Administration

The process of system governance and the role of the Tableau Server administrator are much like that of any other application. However, in Tableau, administrators can be assigned to either System or Content Administrator roles. System Administrators have complete access to all software and functions within Tableau server. They can then assign select users to the role of Content Administrator who will manage users, projects, workbooks, and data connections within the group they are assigned to. This allows each group to better manage their own needs.

Key areas the Administrators will be responsible for are:

- Software installation
- Software upgrades
- Monitoring performance and server utilization, and system tuning
- Processes that support security, backup and restore, and change management
- Managing users, groups, projects, workbooks, and data connections

Although Tableau is extremely flexible and can handle tens of thousands of users and more, its server administration tasks are very much part time. In fact, after the initial setup, most organizations find they spend very little time on Tableau Server administration. The time required will likely depend on the number of users, the frequency of user changes, and whether the admin provides any level of user support.

Data Strategy

Every organization has different requirements and solutions for its data infrastructure. Tableau respects an organization's choice and fits existing data strategy in two key ways: first, Tableau can connect directly to data stores or work in-memory; and second, Tableau works with an ever-increasing number of different data sources.

Access a Variety of Data Sources

At its very simplest, Tableau connects to a single data source with a single view whether the data is in large data warehouses, data marts, or flat files. The view can be a join of multiple tables within that data source which could be a:

- **Relational database –** Multiple tables within a single schema in relational databases such as SQL Server, Oracle, Teradata, DB2, and Vertica, can be joined together.
- **Multi-dimensional (OLAP or "cube") database –** Technologies such as SQL Server Analysis Services and Essbase
- **Access MDB file –** Multiple tables within the Access database can be joined together.
- **Excel spreadsheet –** Each tab in the spreadsheet is treated as a single table and different tabs can be joined in the same way relational database tables can.
- **Flat files –** Files using the same delimiter (comma, tab, pipe, etc.) residing in the same Windows folder can be treated as individual tables within a database. For Access, Excel, and flat files, the default database engine is the Microsoft Jet engine.
- **Hadoop –** Data in a HDFS using HIVE and Apache Hadoop from Cloudera."

Users have the ability to define joins between tables as long as those joins are supported by the database. If all of the data needed is in a single database management system (DBMS), such as Oracle, SQL Server, or Teradata), the database administrator (DBA) can create a database view pulling data together from various schemas or the user can create a logical view of the data with custom SQL.

Data can be stored in any structure including transactional (3rd, 4th, or 5th normal form), denormalized "flat" forms, and star and snowflake schemas. The Tableau view performance is directly related to the speed of the underlying structure of the database. While multi-dimensional databases generally perform best, a relational database with a clean star schema or an analytics-optimized database will perform higher than most other highly-normalized transaction-oriented databases.

Using Extracts for Efficiency and Offline Access

Tableau can connect directly to data or bring data in-memory. If you have invested in fast, analytics-optimized databases, Tableau will connect directly with an optimized connector to let you leverage the value of that investment. If you have a data architecture built on transactional databases or want to keep analytical workloads off your core data infrastructure, Tableau's Data Engine provides an analytics-optimized in-memory data store. Switching between the two is simple.

By default, Tableau provides a "real-time" experience by issuing a new query to the database every time the user changes their analysis. While this can have its own benefits, it can also be an issue if datasets are large or data sources are underperforming or even offline. When data is not constantly changing, real-time queries create an unnecessary workload.

In these cases, Tableau offers an extract capability that brings back data from the initial query and stores it on the user's local machine. The extract is stored in Tableau's columnar database that is highly compressed and structured for rapid retrieval. Extracts can be created from all database types except multi-dimensional databases (cubes).

Using Tableau extracts can greatly improve the user experience by reducing the time it takes to requery the database. In turn, extracts free up the database server from redundant query traffic. Extracts are a great solution for highly-active transactional systems that cannot afford the resources for date-time queries. The extract can be refreshed nightly and available offline to users during the day. The ability to access data offline can be helpful for users who are traveling or otherwise off of the network.

Extracts can also be subsets of data based on a fixed number of records, a percentage of total records, or a filtering of the data. Tableau can even do incremental extracts that update existing extracts with new data. Extract subsets can speed up development time. Developers can use a small subset of data to build a visual application and they won't have to wait for a delayed query response every time they make a change.

Extracts are necessary to share packaged workbooks. Tableau's packaged workbooks (.twbx file type) contain all the data that was used making it both portable and shareable with other Tableau users. Packaged workbooks can also be processed by Tableau's Reader which gives users an interactive experience but with static data, and without the security measures of Tableau Server.

If a user publishes a workbook using an extract, that extract is also published. Future interaction with the workbook will use the extract instead of requesting live data. If enabled, the workbook can be set to request an automatic refresh of the extract.

Lastly, keep in mind that the amount of temporary disk space used to build an extract can be significant. An example of this would be a star schema with a long fact table and many dimensions, each with many, long descriptive fields.

Tableau Data Server

The Tableau Data Server is the component by which Tableau Server provides sharing and centralized management of Tableau Data Extracts and shared proxy database connections. This allows IT to make certain data sources available to users of Tableau Server without duplicating extracts or even data connections across workbooks. This means the organization has a way to centrally manage:

- Data connections and joins
- Calculated fields (for example, a common definition of "profit")
- Field definitions
- Sets and groups
- User filters

At the same time, users can extend the data model by blending in new data or creating new definitions. The centrally managed data will not change, but users retain flexibility.

Published data sources can by of two types:

1. Tableau Data Extracts: Users can connect directly to a published data extract. An organization may choose this approach to provide users with fast, self-service analytics while taking load off critical

systems. Centralized data extracts also prevent a proliferation of data silos around an organization. Data refreshes need only be scheduled once per published extract and users across the organization will stay up to date with the same shared data and definitions.

2. Shared Proxy Connections: Users can connect directly to live data with a proxy database connection. This means that each user does not have to set up a separate connection, making it easier to get started working with data. The user also does not need to install any database drivers, reducing the load on IT to distribute drivers and keep them up to date.

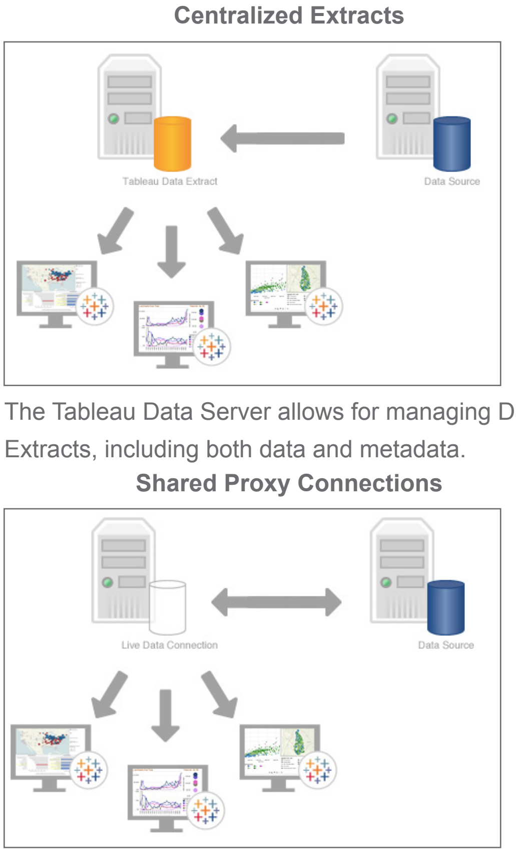

Centralized Extracts

The Tableau Data Server allows for managing Data Extracts, including both data and metadata.

Shared Proxy Connections

The Tableau Data Server can also support live proxy connections.

Figure 6: *Tableau supports centralized management of both live data connections and extracts.*

Report Governance

As data and information continues to grow, information governance becomes critical. It's important that users only have access to information they are authorized to see. The best way to do that with Tableau is to create projects. Projects isolate views and limit individual users to seeing only the views for which they have permission.

Many organizations choose to create a project for each business unit, like Sales, or for each logical business function, like finance. Once a project is created, users or user groups can be associated with the project. Users not associated with a project do not see any of its views. This process creates what is sometimes referred to as a "Chinese Wall" between views.

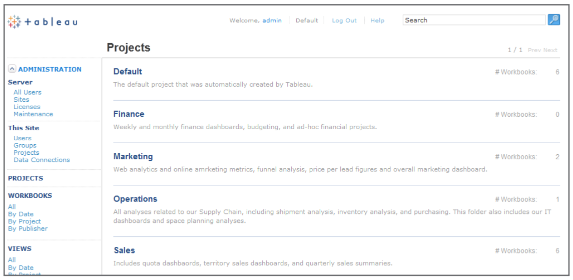

Projects

Default — # Workbooks: 6
The default project that was automatically created by Tableau.

Finance — # Workbooks: 0
Weekly and monthly finance dashboards, budgeting, and ad-hoc financial projects.

Marketing — # Workbooks: 2
Web analytics and online amrketing metrics, funnel analysis, price per lead figures and overall marketing dashboard.

Operations — # Workbooks: 1
All analyses related to our Supply Chain, including shipment analysis, inventory analysis, and purchasing. This folder also includes our IT dashboards and space planning analyses.

Sales — # Workbooks: 6
Includes quota dashbaords, territory sales dashboards, and quarterly sales summaries.

Figure 7: *The Tableau content management interface allows easy report governance.*

Change Management

Tableau offers several options when it comes to performing change management with workbooks. Organizations that have existing change management tools should go ahead and use them to track changes in Tableau workbooks. Those without these types of tools can set up a manual change management process by either creating user folders on the network or using a backup server and performing a nightly backup that is saved. These can then be restored as needed.

As with any development process, moving work from development to production should follow stringent guidelines that include testing and approvals. One of the best and easiest ways to support the promotion to production is to set up a staging area project that parallels each production project. Staging projects can be set up either on a development server or directly on the production server. Users publish new work to the appropriate staging area and send a follow-up request to the team responsible for validation and promoting to production. When the data source behind the Tableau view used in development is different than production, the connection information will need to be changed.

Metadata Management

Most Business Intelligence platforms say they offer metadata functionality but require modeling the entire enterprise as a first step or they don't provide metadata capabilities at all. Tableau has taken a hybrid approach so that IT can add value by providing a rich metadata layer – yet business people are empowered to modify and extend it. This means the metadata layer does not require extensive modeling exercises before getting started. Tableau has been so successful in making metadata seamless, approachable and transparent, that customers often don't realize Tableau has a metadata layer.

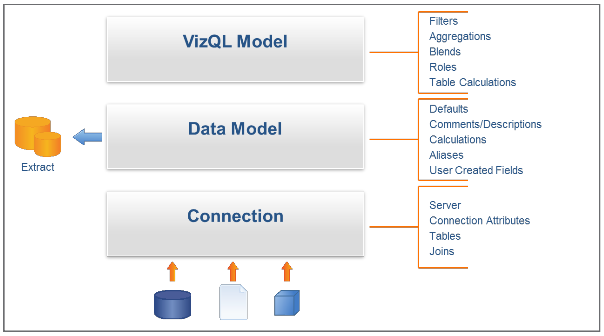

Figure 8: *Tableau's metadata management system allows for rich yet flexible metadata.*

Tableau's metadata system is a 3-tier system with 2 layers of abstraction and a run time model (VizQL Model). The first layer of abstraction is the Connection which stores information about how to access the data and what data to make available to Tableau. It includes attributes for the database, tables, views, columns, joins, or custom SQL used to access the data.

The second layer is the Data Model which automatically characterizes fields as dimensions and measures. When connecting to cubes, the fields are read directly from the metadata in the cube. In relational data, Tableau uses intelligent heuristics to determine whether a field is a dimension or a measure. The Data Model also keeps track of user-generated fields, such as data sets and calculations. Known together as a Data Source in Tableau, the Data Model is independent and insensitive to changes in the Connection.

The third layer is the VizQL Model which is unique to Tableau. The VizQL Model lets the user adjust the role and aggregation of the fields at run time. For example, a user can change a measure to a dimension, such as the calculation of 'average employee age' to using employee age as a dimension against the 'number of employees' measure. Many calculations and comparisons that are difficult to define in a typical data model are easy to define in the VizQL Model.

Tableau provides additional metadata flexibility. Users can combine data from different Data Sources into one hybrid model without any changes to the Connections or Data Models. Data Models can have dependency on other Data Models. And, a Connection can be used in multiple Data Models, a Data Model can be used in multiple sheets and multiple views can be used in a dashboard.

The real value in metadata is the ability to share and reuse components. Data Sources (Connections and Data Models) can be published to Tableau Server creating relationships to workbooks. This means changes in a master data source is automatically propagated to the workbooks that use that data source. Also, other users can use a Data Source as a starting point for their analysis. And, Data Sources can also be exported and shared as files.

Unlike connections, explicit changes to the Data Model must be made in Tableau Desktop due to the broader scope of changes, such as redefining a calculation. When a change is made Tableau automatically manages that change across all sheets in a Workbook. And, although the VizQL Model is insensitive to changes in the database, such as renaming a member of a column, Tableau Server is sensitive to renaming or removing columns used in a View. If an expected column is missing, the VizQL Model will temporarily remove it from the view.

Mobile Deployment

According to a 2011 report from InfoTrends , mobile knowledge workers account for 70% of the total workforce in the U.S. Enterprise adoption of mobile devices, especially ultra-mobile devices such as tablets, netbooks and smartbooks is soaring. Bringing business intelligence into the places where decisions are made and discussions are happening is the real promise of business intelligence. Users have come to expect the same experience on their mobile devices that they have had on their laptops and computers—business intelligence capabilities included.

Tableau delivers mobile business intelligence with the same power and simplicity as the rest of its solution. Tableau's author-once approach to mobile BI means that a dashboard on Tableau Server works automatically on both mobile devices and in a computer's web browser. There is no need for custom dashboard development as Tableau automatically detects the device and optimizes the visual output and capabilities accordingly. Users will see touch-optimized views on a native iPad app, mobile Safari (for iPad users who have not downloaded the Tableau app), and Android devices both as a native app and in the Android browser.

Tableau's mobile apps connect to any Tableau Server and provide native touch controls such as dynamic, finger scrolling, and pinch and zoom. Views are touch-enabled allowing finger control of filters, parameters, pages, and highlighting. Users of the Tableau iPad and Android apps also get touch-optimized content browse, with the ability to search for workbooks, save favorites and see recently used content.

Mobile device security is critical and is a concern for many organizations. Tableau Server enforces the same security for all views, including data-level and user-level security, whether they are served on the desktop or a mobile device. And since no data is stored on the device, there is little risk that lost or stolen devices will compromise data security.

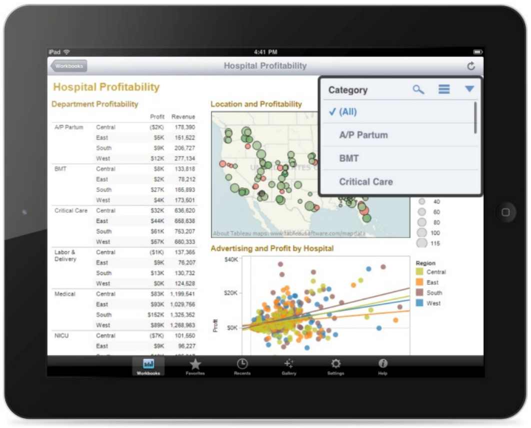

Hospital Profitability

Department Profitability

		Profit	Revenue
A/P Partum	Central	($2K)	178,390
	East	$5K	151,622
	South	$9K	206,727
	West	$12K	277,134
BMT	Central	$8K	133,818
	East	$2K	78,212
	South	$27K	165,893
	West	$4K	173,501
Critical Care	Central	$32K	636,620
	East	$44K	658,638
	South	$61K	753,207
	West	$57K	660,333
Labor & Delivery	Central	($1K)	137,365
	East	$9K	76,207
	South	$13K	130,732
	West	$0K	124,528
Medical	Central	$83K	1,199,541
	East	$93K	1,029,756
	South	$152K	1,325,352
	West	$89K	1,268,963
NICU	Central	($7K)	101,550
	East	$9K	96,227

Figure 9: *Tableau dashboards are automatically touch-enabled for mobile devices.*

Conclusion

Tableau Server provides a robust infrastructure that meets the security, scalability and architectural requirements of IT managers and administrators. It provides flexible deployment options that scale up to the largest businesses. It supports your data architecture decisions by allowing live connection to a variety of databases or in-memory analytics. And most of all, it lets IT managers get back to strategic IT by getting them out of the dashboard-creation and update cycle.

Tableau gives organizations what today's business demands: a self-service business analytics solution that is truly enterprise-ready.

Further Reading

For more depth on some of the subjects in this paper, please refer to:

- *Tableau Server Administrators' Guide*
- Whitepaper: *Metadata Management in Tableau*
- Whitepaper: *Rapid-Fire Business Intelligence*
- Knowledge Base: *Optimizing Tableau Server Performance*
- Whitepaper: *In-Memory or Live Data: Which is Better?*

End notes

[i] http://www.infotrends.com/public/Content/Press/2011/01.20.2011a.html

About Tableau

Tableau Software helps people see and understand data. Ranked by Gartner and IDC in 2011 as the world's fastest growing business intelligence company, Tableau helps anyone quickly and easily analyze, visualize and share information. More than 7,000 companies get rapid results with Tableau in the office and on-the-go. And tens of thousands of people use Tableau Public to share data in their blogs and websites. See how Tableau can help you by downloading the free trial at *www.tableausoftware.com/trial*.