*see*POWER Technical White Paper

Abstract

With *see*POWER™, Compudigm International™ have created a world-class Business Intelligence solution that delivers a new generation of Data Visualization tools for in-depth business and market analysis. By presenting business information as intuitive visual patterns, *see*POWER leverages the remarkable ability of the human visual system to detect and analyze patterns in pictorial representations of huge volumes of data. *see*POWER represents a radical innovation in the way that executive decision makers can see and understand their business.

*see*POWER can be configured to provide business solutions to companies in a wide variety of industries. Please refer to our brochures and web site for more details on how *see*POWER can be applied to these industries.

This Technical White Paper discusses the high level architecture of *see*POWER, and looks at how it complements and enhances traditional data mining tools to bring a powerful new paradigm to business analysis.

Introduction

Today, Business Intelligence is a world of charts, graphs and reports - the same as it has been for the last twenty years. The existing tools are adequate for analyzing a small number of Key Performance Indicators (KPIs), or producing a fixed set of standard reports.

However, the standard range of Business Intelligence tools is not well suited for analysis of large volumes of business data, which is typically required by modern organizations.

Compudigm's *see*POWER uses Data Visualization to present a complete and coherent picture of how an organization operates and interacts with its customers.

In today's technocentric and data oriented world, businesses have managed to amass more corporate data than they can analyze, often using disparate datasets.
Data Visualization provides the capability that businesses need to extract meaningful information from all this data, and thus is emerging as a leading business data analytical solution.

Data Visualization uses the power of the human eye and brain to discern relationships by presenting complex data as multi-dimensional color images and animations, with the end result providing the viewer with a qualitative and quantitative understanding of the information's content.

A central component to Data Visualization is that it provides management with an intuitive understanding of their business. *see*POWER enables management to quickly interpret KPIs such as net revenue, profit, turnover, margin, customer spend and so on.

*see*POWER takes advantage of Sybase's relational database for data warehousing, Adaptive Server IQ with Multiplex (IQ-M), to obtain superior performance and schema flexibility. By combining the power of IQ with Compudigm's unique data visualization algorithms, seePOWER produces visualizations that are both information-rich and intuitive.

Traditional Database Mining versus Data Visualization

The current standard data analysis technique for large volumes of data is Data Mining or On-Line-Analytical-Processing (OLAP). OLAP tools typically use a datacube model – an intrinsically abstract, mathematical construction. Thinking of data as multi-dimensional hypercubes works well for database theorists, but is inadequate as an intuitive metaphor for executive decision-makers.

Traditional data mining tools concentrate on data aggregation and slicing and dicing of the data cube, and presents the resulting data in some combination of numerically based charts, graphs and reports. While this is adequate for businesses that are content with a fixed set of standard analyses, these tools are of limited use when searching for unknown features in an uncharted expanse of data. Take one slice too many out of a hyper-dimensional datacube and chances are you'll end up lost in a sea of numbers.

*see*POWER complements traditional OLAP tools by allowing users to perform a data survey as a precursor to datamining (Figure 1). *see*POWER's ability to present huge volumes of data as a complete and coherent picture ensures that interesting features of the data stand out as prominent visual patterns that are readily identified by business decision-makers. Once the features of interest have been identified with the visual data survey, data mining can begin in earnest.
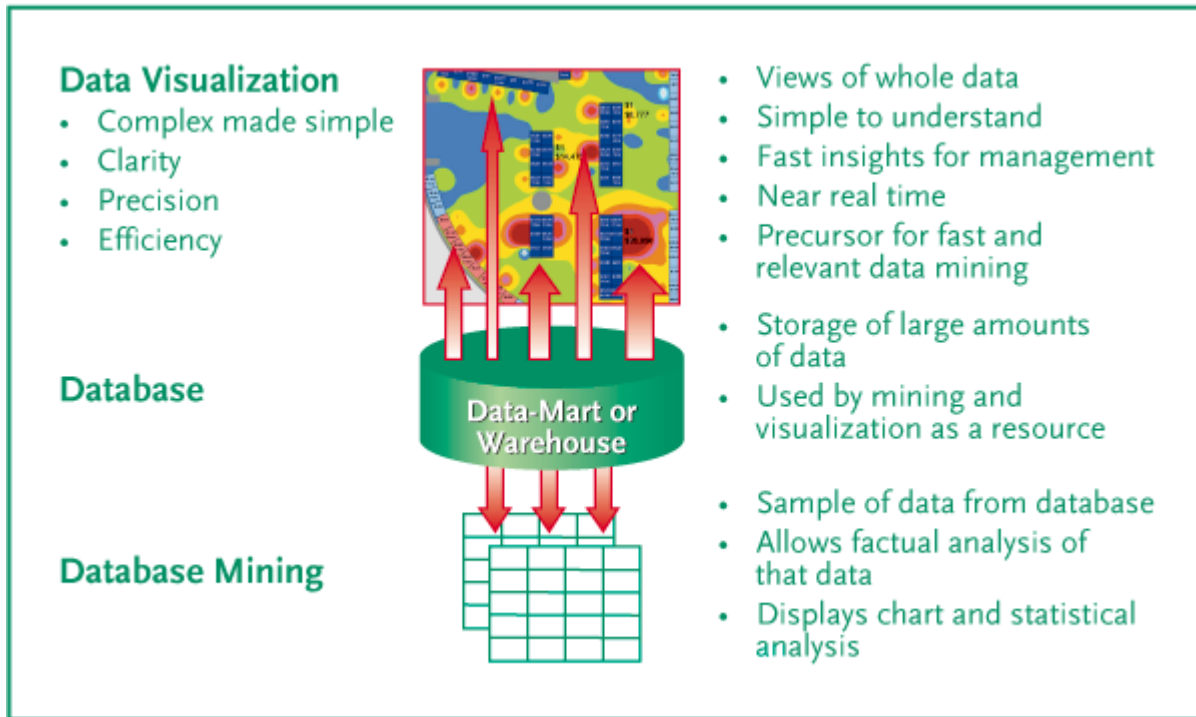
Figure 1: Database Mining versus Data Visualization

*see*POWER Architecture

Key Elements of *see*POWER

Compudigm's *see*POWER is a multi-tier, distributed application designed to be highly scalable, extendible and configurable (Figure 2).

*see*POWER uses two distinct types of middle-tier servers – the *see*STORE server and the *see*OBJECT servers. The functions of these servers are described below.
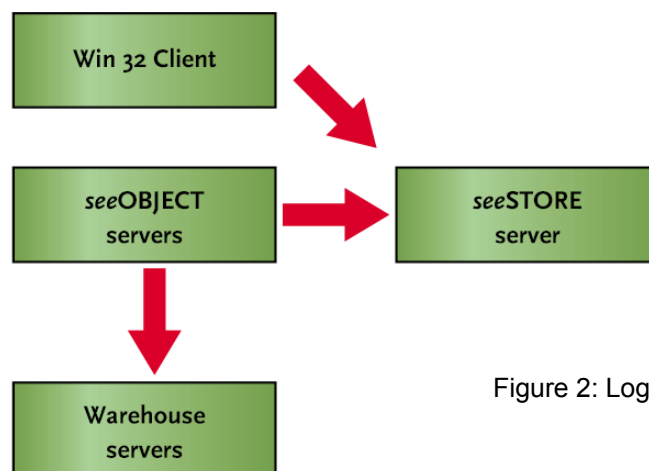

Figure 2: Logical view: n-tier architecture

SEEP●WER™

A central concept in *see*POWER architecture is the *see*OBJECT – the interactive business objects created by *see*POWER, which end users create at the client using simple business-oriented wizards, and through which users view their business data.

*see*OBJECT Servers

The *see*OBJECT servers are responsible for creating the *see*OBJECTs. They respond to queued job requests by querying the data warehouse and constructing the requested *see*OBJECTs. The resulting *see*OBJECTs are placed in the *see*STORE, where clients may view them.

*see*STORE Server

The *see*STORE server acts as a central object repository for *see*OBJECTs. Clients are essentially *see*STORE browsers
– the user navigates through the *see*STORE, selects the *see*OBJECT they wish to view, and the client invokes the appropriate *see*OBJECT viewer.

Scalability

Multiple Windows NT *see*OBJECT servers may be run in parallel to provide unlimited scalability on business object generation, constrained only by the ability of the warehouse to supply data (Figure 3). When run against Sybase's Adaptive Server IQ-M, the system scales indefinitely.
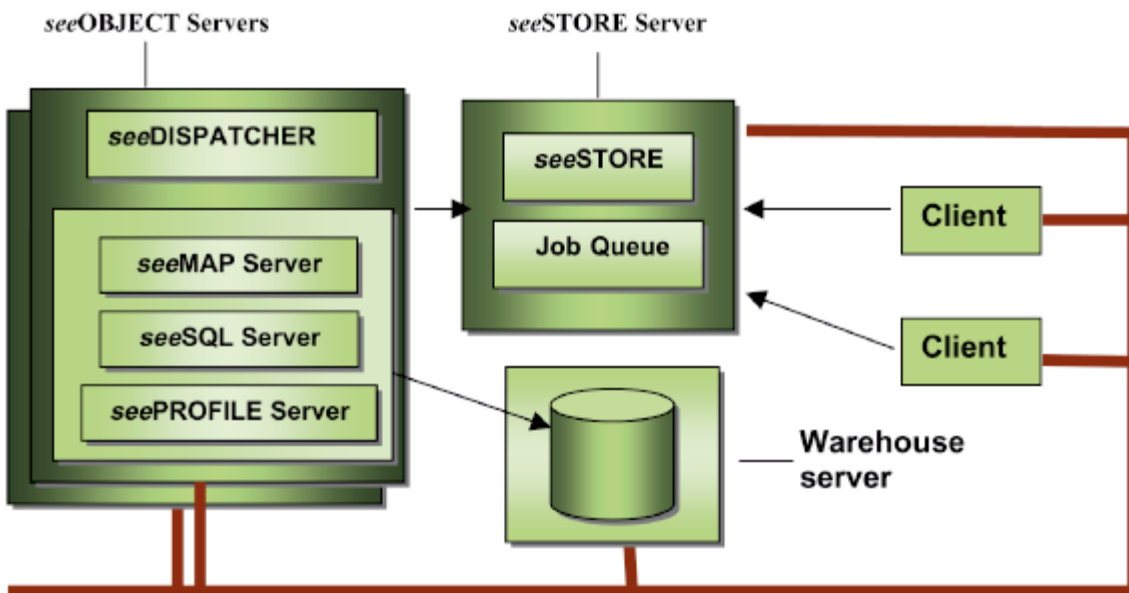

Figure 3: Physical view: Scalability through multiple *see*OBJECT servers

SEEPOWER™

Job Queue

*see*POWER uses a persistent Job Queue on the *see*STORE to decouple the clients from the servers. This ensures a robust system, because clients can continue to request jobs even when the *see*OBJECT servers are unavailable, and the jobs will be actioned when the servers come back online.

It also allows the clients to behave asynchronously – they do not have to wait until the job is complete – the user simply submits a job request, and may then browse to another *see*OBJECT. Jobs may contain scheduling information to allow them to trigger repeatedly.

Configurability

A comprehensive meta data layer allows *see*POWER to be configured to run in any business environment capable of providing an SQL interface to its business data. Setting *see*POWER up to run in a new business environment is a simple matter of mapping the generic *see*POWER concepts to the business domain through a simple, wizard driven setup application.
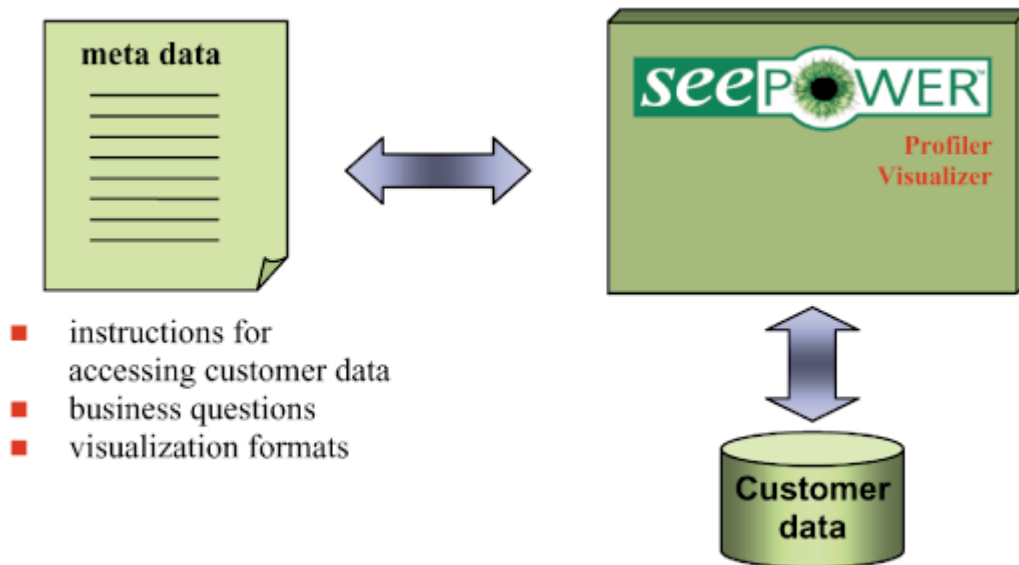

Figure 4: Configurability of *see*POWER via meta data

Extendibility

New kinds of *see*OBJECTS may be added to the system by adding new plug-in *see*OBJECT viewers to the client, and new *see*OBJECT servers to the middle tier.

The object-orientated nature of system applications means that infrastructure elements depend only on the generic *see*OBJECT abstraction and are not perturbed by the introduction of new *see*OBJECT subclasses to deal with new kinds of business applications.
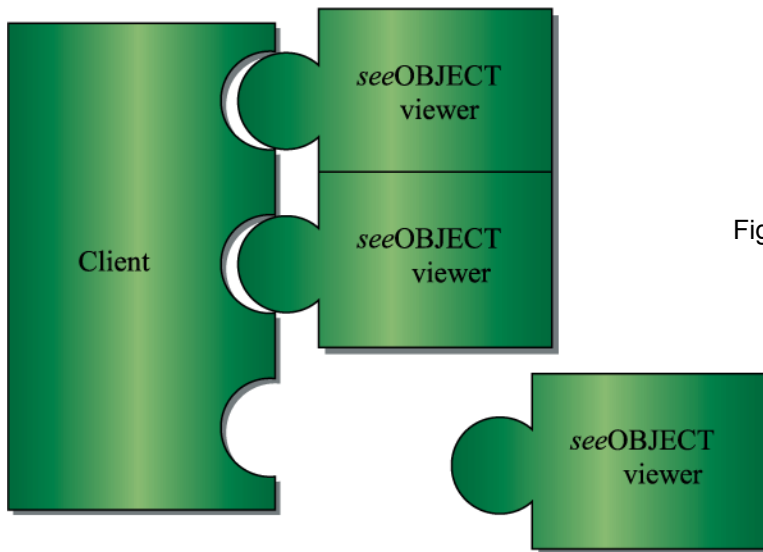
Figure 5: Extendibility of *see*POWER

Warehouse Data Flows

The data warehouse is updated daily or more frequently from the range of transaction databases (Figure 6). While the databases and warehouse can be any platform, Compudigm has found Sybase IQ to be the most reliable and efficient.
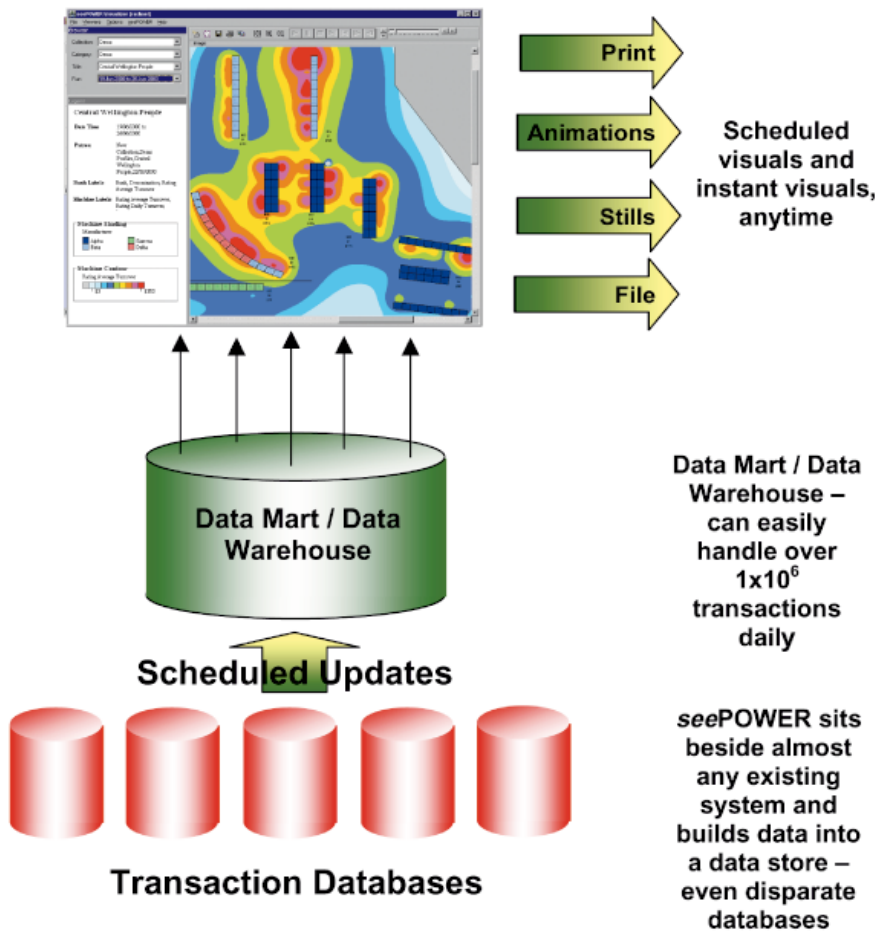

Figure 6: Warehouse Data Flows

seePOWER™

System requirements to run *see*POWER

The minimum specification required to run *see*POWER effectively will vary depending on a number of factors, including:

• the size of the underlying data mart / data warehouse,

• the number of concurrent users,

• the number and type of jobs being sent to the *see*OBJECT servers. In particular, animations require greater performance.

For good performance, it is recommended that the *see*STORE server be run on a dedicated server that has high availability – i.e. there should be no other CPU intensive applications running on that server. At least one separate *see*OBJECT server is required. *see*OBJECT servers are CPU intensive, so in general, the faster the server, the faster your *see*OBJECT requests will be available.

It should be noted that while the *see*POWER clients and servers require Windows NT or Windows 2000 platforms, the data mart / data warehouse may sit upon a number of different platforms.

Recommended minimum specifications:

1 to 5 users

Machine	CPU minimum	RAM	Free Disk Space	Operating System
Client	PIII 500	128MB	100 MB	- Windows NT (Service Pack 5 and above)
				- Windows 2000
*see*STORE Server	PIII 700	512MB	16 GB	- Windows NT (Service Pack 5 and above)
				- Windows 2000
*see*OBJECT Server	PIII 700	512MB	16 GB	- Windows NT (Service Pack 5 and above)
				- Windows 2000

5 to 10 users

Machine	CPU minimum	RAM	Free Disk Space	Operating System
Client	PIII 500	128MB	100 MB	- Windows NT (Service Pack 5 and above) - Windows 2000
*see*STORE Server	PIII 700	1GB	16 GB	- Windows NT (Service Pack 5 and above) - Windows 2000
*see*OBJECT Server1	PIII 500	256MB	16 GB	- Windows NT (Service Pack 5 and above) - Windows 2000
*see*OBJECT Server2	PIII 500	256MB	16 GB	- Windows NT (Service Pack 5 and above) - Windows 2000

Customers will be responsible for obtaining the appropriate operating system licenses.

All hardware should lie on a 100 Mbit network.

In addition, a data warehouse or data mart server is required. Consult your warehouse vendor.

This document was produced in September 2001. The specifications and architecture may have changed after this date. Please contact Compudigm for the latest specifications.