

REVELATION

S O F T W A R E

A R E V E L A T I O N W H I T E P A P E R

**The world is not flat
and neither is your business.
So why should your data be?**

**Revelation Software's Post Relational,
Multi-Dimensional Database Environment.**



Contents

CONTENTS	3
INTRODUCTION	4
The Purpose of this White Paper	5
DISCUSSION	6
MultiValue – An evolution based on success	6
Revelation's ARev and OpenInsight	7
Maintaining a competitive edge over RDBMS	8
Supporting Embedded DBMS ISVs	9
Optimising systems for a modern world	11
Market Stats	12
CONCLUSION	13
APPENDIX A	15
APPENDIX B	18
APPENDIX C	23
APPENDIX D	27

COPYRIGHT NOTICE

© 2006 Revelation Software Limited. All rights reserved.

No part of this publication may be reproduced by any means, be it transmitted, transcribed, photocopied, stored in a retrieval system, or translated into any language in any form, without the written permission of Revelation Software Limited.

TRADEMARK NOTICE

OpenInsight is a registered trademark of Revelation Technologies, Inc. Advanced Revelation is a registered trademark of Revelation Technologies, Inc. OpenInsight for Workgroups is a registered trademark of Revelation Technologies, Inc. Report Builder+ is a trademark of Revelation Technologies, Inc.

Microsoft, MS, MS-DOS and Windows are registered trademarks of Microsoft Corporation in the USA and other countries. IBM, U2, UniData and the IBM e-business logo are registered trademarks of International Business Machines Corporation. Intel is a registered trademark of Intel Corporation. Lotus, Lotus Notes and Notes are registered trademarks of Lotus Development Corporation. All other product names are trademarks or registered trademarks of their respective owners.

Introduction

The world is not flat and neither is your business.

So why should your data be?

Throughout the world, in all businesses, in all industries, the challenge of best management of business data is increasingly critical to competitive advantage. As the emphasis moved towards the creation of large data warehouses for Business Intelligence (BI), the Relational approach to managing business data became the accepted norm for data storage. This can mostly be attributed to Dr Codd (more on Codd later) and a handful of corporations having based their business models on the supply of Relational Database Management Systems (RDBMS) software.

There are circumstances when the normalised, cross-referenced, relational data structure using flat lists is the optimal choice, such as warehousing data for BI. However, there are as many circumstances where a multi-dimensional data structure will be the optimal choice. Examples of such circumstances include increased system performance, enhanced user experience or to provide the developer with a more flexible application development environment that is best suited to longer term business benefits and goals.

For instance, users often have data that involves the storage of recurring details and other nested data structures. In the RDBMS model this involves the creation and management of many data tables, their associated indexes and relationships between tables, often swelling the size of their database. This model can result in a collection of data tables that require complex joins, multiple indexes, countless table lookups to function, all of which all carry a resource overhead. In addition, they can require high levels of database management, control and administration which further limit responsiveness to meeting user needs, performance and functionality. Perhaps unsurprisingly, many of these applications typically require rewriting after only a few years and incur significant maintenance efforts and costs.

Appendix A contains a simple illustration of how a typical MultiValue Database Management System (MVDBMS) database structure might compare to an identical RDBMS.

So why don't businesses embrace database technologies that employ both relational and multi-dimensional ideals?

Well actually, through the use of MultiValue toolsets (also now known as Multi-Dimensional or Post-Relational technologies) like Revelation's OpenInsight, countless numbers of developers and many millions of users worldwide do both favour and rely heavily on such technologies.

The problems associated with the pure relational model were reviewed many years ago and consequently various extensions to that model were created and adopted by businesses needing to manage and process complex data stores. This includes those businesses offering commercially developed solutions for use in vertical markets. Such markets include: Healthcare, Finance, Government Departments, Construction, Manufacturing, Payroll and Human Resources,

to name but a few. While looking at this list, it is not hard to realise that most people's lives in the UK for instance, are regularly touched by a MultiValue based system.

But still this technology is not well supported by the authors and suppliers of RDBMS software. Over the last two decades, Revelation Software has consistently provided professional application developers with toolsets that enable them to easily embrace this Post-Relational or Multi-Dimensional business need.

The Purpose of this White Paper

By releasing this white paper, Revelation Software is responding to renewed interest in Multi-Dimensional (MultiValue) Database Management Systems (MVDBMS) and is further promoting its flagship development environment, OpenInsight, as one of the IT industries best kept secrets. Through its suite of application development and deployment tools, the company offers business a wider choice which draws on a successful, proven, and rich application development heritage.

This paper will provide the reader with an understanding of the benefits of a MultiValue database driven solution. It will also take a glance at the benefits of using MultiValue technologies over pure relational systems. It will hopefully fuel the growing debate surrounding the benefits of MultiValue systems over Relational systems.

Throughout this paper, we will use one of the world's leading MultiValue Windows and Linux GUI development environments. Revelation's OpenInsight will be used to demonstrate how MultiValue based solutions help businesses globally to embrace the best of both worlds, in terms of application development, database management and overall system performance. This is especially true when embedded in vertical applications, where special data management considerations need to be addressed and minimal database administration is mandatory.

Discussion

MultiValue – An evolution based on success

Back in the early days of computing, getting information into and more importantly, out of large complex systems built using fixed serial streams of pre-formatted inputs and outputs was both time consuming and costly. To make information more readily available and useful, new data management formats were required. One such breakthrough came during the 1960's. Furthermore, that technology is still very relevant and relied on in today's modern business environment each and every day.

During the Vietnam War, the US Government funded a project called the General Information Retrieval Language System (GIRLS). The intention of the system was to enable U.S. Army GIs on the ground to create their own reports. This enabled them to quickly track which helicopter parts were on order, due in, etc. and keep the fleet in the air. The chief point to this new system was that the query language should be as natural, or 'English' like as possible. This project was a success and, it resulted in the implementation of the new technology to the GIs by Dick Pick and Don Nelson in 1965.

Identifying the business potential of the new technology, Dick Pick obtained a license to further develop it following the end of the war. Under the control of Pick, it was developed into a combined multi-dimensional DBMS and Operating System that was subsequently commercially released in 1973 by Microdata and their British distributor CMC.

Pick evolved into many versions until later in the 1970's, the PRIME Corporation developed a DBMS based on the Pick model and integrated it into their Prime Operating System¹. Known as PRIME Information, it quickly became one of the leading versions of Pick available at the time.

Dick Pick eventually founded Pick & Associates, later renamed as Pick Systems and licensed what was known simply as 'Pick' to a variety of software manufacturers and vendors who have since produced different 'flavours' of Pick. Pick Systems later became Raining Data, supplying a flavour of Pick called D3. Another leading flavour includes IBM's DB2 product suite (incorporating UniData and Universe).

During the early 1980's the IBM Personal Computer (PC) was released and quickly adopted by business in general, bringing new power to employee's desktops. At this time, some programmers working at Boeing decided to do an Information style implementation for the PC. They created a version of Pick that was to become known as Revelation B. The company formed as Cosmos and was later bought by one of its largest customers. It then became known as Revelation Technologies, Inc. The company went on to grow rapidly owing to the flexibility and rapid application development characteristics of the toolsets, coupled with the MVDBMS. Revelation's toolsets include Revelation and Advanced Revelation, both for the MS-DOSTM platform and more recently OpenInsight for the WindowsTM and Linux platforms.

¹ During the 1960's and 1970's it was common for both operating systems and DBMSs to be bundled together and offered with the hardware as a 'total' high value proposition solution.

Meanwhile, during the mid 1980's RDBMSs began to emerge and quickly became the dominant model for DBMSs on midrange and mainframe computer systems. However, unlike 'Pick' based systems, these new systems were independent of the operating system and ran on VMS and UNIX. It was this OS independence and the simplistic 'tabular' nature of the underlying database in the form of rows and columns, which lead to the rapid adoption of RDBMS for general-purpose databases across the DBMS market.

These new RDBMSs were primarily modelled on the theories of one Dr. Edgar Codd (http://en.wikipedia.org/wiki/Edgar_F._Codd), which were published as 'Codd's twelve rules'.

These rules had two main objectives. Firstly, they were designed to define what is required from a DBMS for it to be considered purely 'Relational'. Secondly, it was a way of preventing Codd's vision of the Relational Database model from being diluted by database vendors scrambling to repackage their solutions under the Relational heading. In the end, the rules were so strict that even today, popular 'relational' DBMSs fail on many of Codd's rules.

Even though RDBMSs were becoming more and more fashionable, many Independent Software Vendors (ISVs) discovered that the Relational model was missing desirable features and was an inefficient method for storing complex data of a hierarchical, or nested format. For these business data needs, the Pick model remained the superior option.

This meant that systems running on VMS, UNIX and IBM's new PC, were fast becoming more prevalent and a company called VMark produced a Pick based DBMS that could be run on UNIX, as did another company called Unidata, with a product of the same name and optimised to run on the PRIME Information DBMS. Revelation in the meantime, continued to enhance their Revelation product line to take advantage of the emerging PC market.

All of these products flourished owing to the simplicity of their use, their performance and for their flexibility. Another key attribute that was important at the time, was the ability to finely tune the memory and data storage to best suit the application, optimising the use of costly disk space and system resources. As is still that case today, these MultiValue solutions carried an extremely small footprint compared to RDBMSs. Even today, OpenInsight retains its small footprint of around 60Mb (or lower) for the development suite itself.

Revelation's ARev and OpenInsight

Since Revelation's earliest days back in 1982, the company has provided fully integrated application development suites combining functional DBMS with a proven, robust and flexible Pick based database. Initially branded 'Revelation', the products were optimised to run on PCs and which were being readily adopted by all types of business in the early 1980's.

By 1987, and now branded as 'Advanced Revelation', the company found themselves with an award winning application development environment for DOS based single and multi-user systems, with a global user base.

Throughout this time, the underlying database remained truly multi-dimensional (Pick based) while taking advantage of the idea of mapping data in rows and columns through the use of powerful dictionaries and a nested data structure. At the same time, Revelation continued to provide its developers with the ability to create underlying data structures that adhere to the Relational model. In doing so, they give their community the best of both worlds (see Appendix A).

This new technology proved extremely popular with both corporate businesses and ISVs alike. Businesses found that they were able to create departmental systems quickly and easily, very often by power users who knew the business better than anyone else. On the other hand ISVs who needed a fully integrated package with a 'set and forget' embedded database and flexible DBMS, found the technology delivered these requirements, whereas they were not readily available from typical pure RDBMS available at the time.

The MultiValue structure remained the basis of Revelation's product suite throughout the 1990's. This continued during the companies move to adopt the new Windows operating system, which was being adopted to ease the use of computers right across the business.

Now with a fully featured Windows™ application Development Suite branded under the OpenInsight label, the company once again found itself with a toolset ideally suited to modern business needs, without the development and administration overhead often associated with the pure RDBMS model.

Unlike many DBMS providers over the last two decades or so, Revelation have continued to specialise in the MultiValue market place and have built a reputation of providing its clients with product enhancements that take advantage of emerging technologies, just at the right time.

This has been demonstrated time and again through the release of network products to maximise performance and stability as network systems grew, without any impact on the ISVs applications. This is also illustrated through the introduction of web functionality (CGI, COM and Java based) just as businesses began to embrace the web. More recently support for XML and web services have been introduced and new operating systems such as Linux supported.

All the while, the company have enabled their ISVs to do flourish in a very changeable business environment. Their database driven solutions could evolve without the need for regular remodelling and rewrites. The benefit to Revelation's developer community has and continues to be, a progressive business model designed to reduce costs of development and deployment of modern systems, all based around the flexible Multi-Dimensional (MultiValue) Pick model.

Maintaining a competitive edge over RDBMS

Codd's idea, based on mathematical set theory, was to store data in cross-referenced tables and provide an efficient way to catalogue data. This data storage model governed the relationship between the database and the application, rather than providing an efficient data storage model that reflected the needs of the business and its working practices.

There is much talk about the pros and cons of both RDBMS and MVBMS, but it remains interesting that some of the pure relational theorists are now seeing major RDBMS suppliers begin to embrace the MultiValue (Multi-Dimensional) data model within their offerings.

This change in the thinking of RDBMS suppliers can be attributed to the restrictions forced upon RDBMS system developers and users. Whilst not initially apparent these restrictions soon gain visibility when the business model itself calls for the storage of complex repeating groups nested over several levels. Obvious applications include Bill of Materials, complex Purchase Ordering systems and even clinical data tracking. These structures can be mimicked in a relation structure but all too often the cost of multiple external joins is heavy in terms of processing required and the time needed to perform them. It frequently occurs that visitors to web sites wait patiently for results only to have the site timeout on them under the sheer burden of a complex SQL query.

On the other hand, OpenInsight is built on a model that natively supports the nested table model and contains MultiValue Variable Length data fields. While pure RDBMS do not support the de-normalised data structure that OpenInsight and other MultiValue system embrace, OpenInsight is specifically optimised to exploit this flexible data storage and retrieval model.

For example, reading or writing a record within an OpenInsight based system requires just 1.2 disk reads for a record using a primary key, or around 4 disk reads for a foreign key record. A similar record in an RDBMS could typically result in an order of magnitude more disk reads as multiple tables are queried, joins performed and locks managed.

Furthermore, OpenInsight's variable length field architecture² provides users with enhanced data storage efficiencies and also enables the application footprint to be diminutive. Disk space is automatically managed without requiring additional work by the system developer and unnecessary system administration at the business level. More importantly, this flexibility ensures that users are not forced to truncate or lose data, something that is frequently seen as an irritation to users and can also be the sign of a poorly designed and implemented system.

While countless RDBMS developers continue to labour with system modelling and normalisation of their data structures, OpenInsight developers are able to benefit from a significant reduction in development times, through a simple and easy to understand data model. In addition, the ability to add dictionary items and complete other similar system enhancements without the need to rebuild the database and then commit those to the live system without the need of costly downtime, or scheduling out of hours work, is seen as a major competitive advantage for many of Revelations ISVs over their suppliers offering RDBMS based systems.

Supporting Embedded DBMS ISVs

The RDBMS and MultiValue worlds differ greatly, while delivering similar results. It is common knowledge that RDBMSs have flourished within the corporate world; however the same can be said for MultiValue technologies in the embedded database world.

There are many reasons why RDBMSs are favoured by businesses the world over, but one of the more readily accepted reasons can be attributed to readily available resources. When a corporate business identifies a need to create an in-house database driven application, they normally take a look at the people-resource available and all too often go with the 'preferred' solution of the individual heading up the team, or undertaking the work.

² Appendix B takes a detailed look at the flexible nature of OpenInsight's database architecture.

Other considerations can include the future direction of the DBMS to be used, the pedigree of the DBMS supplier, the product's ease of use and the availability of ready trained personnel. Essentially, the decision to opt for a particular DBMS in business is based mostly on the evaluation of available resources, rather than evaluating the right product for the project and then hiring the developer and administration resources later.

All too often, business leaders adopt a technology that their staff knows, rather than the DBMS that is the right tool for the job.

This approach often works where there are no concerns about obtaining and maintaining a competitive edge, also where there is a lower emphasis on costs. After all 'I have committed to the developer's salary, so who cares about how long the job will take. IT projects take months and cost a small fortune anyway – right?

Not necessarily. Put yourself in the position of an ISV looking for a flexible DBMS, that will enable the ISV to grow their clients systems in line with their client's business.

The ISVs priority in choosing a DBMS is its 'suitability to the project'. The underlying database will be invisible to the end-user. It should not require a dedicated Database Administrator (DBA) to keep the system optimised and running. The end-users will not have the ability to fix issues with the system, the database must therefore be of a 'set-and-forget' nature and it must just run and run without the need for daily reboots, data storage management and other labour intensive jobs. In addition, the DBMS must be fully integrated, inexpensive to maintain and straightforward to deploy, either by the user through simple install scripts, remotely through periodic updates on CD, or through the use of self installing web downloads.

Another important element of the ISV's decision is competitive edge. The need to remain competitive within the ISV market place is often high on the ISV's agenda throughout the life of their business. Where systems within a vertical market are similar in functionality, it is often price and the commitment of the ISV to their customer base that is the deciding factor when it comes to ordering.

Whereas businesses usually opt for the most easily resourced DBMS, ISVs who build their systems using a MultiValue technology, like Revelation's OpenInsight, usually base their decisions on these evaluation factors. In doing so, they invest in a market place with some of the best possible support resources. For instance, Revelation provides their ISVs with direct access to authors of the DBMS, both in the USA and in Europe on a daily basis.

ISVs choosing a MultiValue DBMS know that they have tools that enable them to enhance their solution to meet their markets requirements quickly and easily. This often helps to regularly demonstrate an ability to meet tight development deadlines. This is especially important in typical MultiValue markets such as finance and healthcare, where legislative changes can force incredibly tight deadlines. This ability to respond flexibly to changing business requirements was graphically illustrated in the UK Pensions market where recently a corporate user of OpenInsight was one of the few suppliers to meet UK reform legislation changes. They credited the speed of their response to their use of OpenInsight.

ISVs in the MultiValue market place opt for a MVDBMS based on the environment best suited to their business solution and their market place. This decision usually provides them with a competitive edge, lower costs of ownership and optimised database driven solution.

Optimising systems for a modern world

Over the last couple of decades, businesses have seen Information Technology evolve at an alarming rate. No longer do businesses have a hand full of PCs, or dumb terminals running against huge expensive mainframes and islands of data used by a handful of managers.

Today, businesses benefit from huge amounts of data processing power right across the business, at all levels. Islands of data are now integrated data systems that provide invaluable central information stores and that business intelligence is shared across the globe. On demand services are now common place, the web continues to deliver new and exciting resources to businesses, and business owned systems are rapidly being replaced by service oriented hosted systems.

RDBMS suppliers have continually enhanced their solutions to meet these needs but they have remained confined by the limitations of the RDBMS model. On the other hand, the MultiValue approach is still proving to be the right solution for many ISVs who need to manage complex data structures and maximise performance across ever wider and busier networks.

Extensions to the RDBMS world have included the introduction of Extensible Mark-up Language (XML) and more recently Asynchronous JavaScript and XML (AJAX). These new technologies are designed to enhance the user's experience and provide more power to the developer of rich applications run over the web.

MultiValue DBMS are a perfect fit for XML. For example, Revelation's OpenInsight makes extensive use of a flexible string based filing system known as Linear Hash, which is modelled on the traditional Pick methodology, as discussed earlier. This method of storing data uses delimiters to tell the DBMS where records begin and end and where MultiValue data is located within the string. This methodology has been the case and benefited Revelations ISVs for over two decades.

When you take a look at the 'modern' XML methodology, you notice that it simply contains tags to describe the content in terms of what data is being described, much like MultiValue DBMSs such as OpenInsight. The result is a 'perfect fit' without the need for a major rethink and rewrite of the DBMS.

As businesses become ever more dependant on the World Wide Web as a delivery vehicle of corporate data and communications with its customer base, new web techniques continue to be developed and reinvented. One such technique, known as 'AJAX' brings together JavaScript and XML into one new technique aimed at further enhancing data driven solutions over the web. Once again, MultiValue toolsets like OpenInsight, have quickly been able to take advantage of a new modern emerging technology to the benefit of MultiValue developers and users.

Appendix C includes Social Software's case study in which they made use of OpenInsight and XML to create a dynamic XML hub, enabling various organisations within criminal justice to exchange data between disparate data sources. Using OpenInsight's MultiValue DBMS, Social Software was able to achieve the desired functionality, at a fraction of the cost that some RDBMS suppliers were indicating.

MultiValue DBMSs are also extremely well suited to processing data over the World Wide Web. For example, OpenInsight's underlying database is and always has been, fully optimised for use over the web. Using Pick's model of locating data in buckets and accessing that data through the use of a complex, but efficient algorithm, OpenInsight is able to retrieve data for the web, as it does for any application, by simply going straight to the record required and returning the data with only one or two disk reads. Unlike with many RDBMS, there is no querying of various tables and working with groups of data spread over the disk. Each request is refined, concise and efficient – a prerequisite for modern web applications.

OpenInsight also fully supports distributed processing. This is achieved when central data engines run on low use machines on the network. This helps businesses to make optimal use of their investment in computing resources and to equally distribute the processing workload across their systems.

One such optimised web system running in Australia is processing over 2,000,000 database transactions each day with only a dozen OpenInsight engines. In addition as soon as products are entered into the database from good inwards they are available on the web for sale. In such an environment, speed of data retrieval and reliability are key. Further details about the TVSN case study can be found in Appendix D.

Market Stats

Studies analysing the various DBMSs being Object Oriented, XML, Relational, or MultiValue (Post Relational, Multi-Dimensional, etc.) are few and far between. However, a competitive analysis report by IDC entitled 'Worldwide Embedded Database Management Systems 2003 Vendor Shares' show the Relational market to be the large sector with around 78% market share and the MultiValue market to be in second position with around 19% market share.

However, while RDBMSs continue to maintain the lion's share of the market, the MultiValue share greatly outstrips the two where growth is concerned. As businesses look to ever flexible DBMS with proven track records and the ability to blend with the business, we see RDBMSs with a growth in 2003 of around only 10% and MultiValue with a whopping 27%. Much of this growth can be attributed to the take-up of these technologies by ISVs looking for DBMSs that fulfil their requirements for a flexible DBMS that is optimised for embedded databases that run equally as well in LAN, WAN and Web environments and have demonstrated an ability to scale accordingly.

With a projected market potential of around \$2.2billion in 2010 and a projected growth rate of around 11% (3% higher than that of RDBMSs) the future for MultiValue systems is set to look very promising. With leading RDBMS suppliers also looking to embrace MultiValue concepts in one form or another, these forecasts could easily be exceeded over the next few years.

Conclusion

Those readers who have not experienced the MultiValue sector of the IT industry, may well be asking 'Why have I not heard, or come across this technology previously'.

Well you probably have already come across it, but not known that you have. Why? Quite simply, because there is no reason for end-users to be told about, or need to know about, the database driving their purchased 'boxed' solution. In addition, there is no benefit, or need for the ISV to promote their chosen DBMS supplier and in many cases give away the secrets of their competitive edge.

This ISV layer between the MultiValue DBMS suppliers and their ultimate end users is the main cause for businesses in general being ignorant to this flexible and proven technology.

In Revelation's case, this is further compounded by the fact that single user runtimes (SURs) are provided to users free of charge. For this reason, Revelation's ISVs (known as Value Added Resellers (VARs) in the Revelation developer community) do not need to register SUR systems with the company. For instance, there is one West Coast USA Revelation client with over 12,000 SUR systems and no reason to register them with Revelation.

The decision to opt for a RDBMS, a MVDBMS, or any other DBMS is a complex one. However, the decision should depend primarily on the proposed usage of the system. As one Revelation VAR explained a few months ago – It's like opting for an 18 wheel truck, or a high powered sports car.

If you need to process corporate-wide data that is structured and not subject to change (such as a corporate data-warehouse or Business Intelligence (BI) system), then opting for one of the leading RDBMS, or the high maintenance 18 wheeler which simply trundles from A to B along wide highways is usually correct. However, if your system needs to be flexible, supportive of your changing business processes with nested data structures and be easy to maintain with little, or no database administration (such as a departmental solution) then a MVDBMS like OpenInsight, will no doubt be the most appropriate, much like the nimble sport car, zipping effortlessly along winding country roads.

There are many choices available when it comes to choosing an embedded DBMS, including several MultiValue DBMS offerings. However, you will require a fully featured and integrated toolset such as OpenInsight, which has a solid pedigree and proven track record. Furthermore, along with the benefits that comes standard with a Pick based MultiValue DBMS, you will need tools to define and manage the integrated database, coding tools to build the business logic layer and solid debugging tools to aid the development process. Then you will require front end design tools for both desktop and web based applications, XML tools and data warehousing tools (to take advantage of your corporate 18-wheeler BI solution). Deployment tools are required to enable the distribution of your application to users and maintain it with upgrades and patches.

OpenInsight contains all of these key features and more. In addition, through its repository driven architecture, this leading Windows and Linux development environment also fully supports team development and the sharing of resources throughout the team and between applications.

Revelation application development environments fully support the Pick (MultiValue) paradigm and through OpenInsight, fully support Windows application development techniques and industry standard communication protocols. They have a proven track record of delivering business critical systems across all industries for over two decades. OpenInsight provides the ability to scale from just a handful of users, to thousands of users largely through evolution, rather than redesign. With systems currently running on every continent (including Antarctica) and with over 2,000,000 users worldwide, Revelation are now spreading the word once again and letting the MultiValue secret out.

Though the continued development of their DBMS environments, Revelation have demonstrated the ability to support and exploit emerging technologies on a 'just in time' basis. In doing so, they ensure that their ISVs invest in solid technologies, rather than vapourware technologies.

Revelation will continue to support latest operating systems and routes to market that enable their ISVs to remain successful, while remaining a leading supplier of GUI application development tools based on the founding MultiValue, Post-Relational, Multi-Dimensional technology. Furthermore, the company will continue to use it's own products to build it's own products, for example OpenInsight is now built using OpenInsight, proving the power, reliability and flexibility built into the system and ensuring that every single Revelation VAR has a toolset that remains fully integrated and able to deliver market leading solutions.

For more information or to experience the toolset please contact your local office (details on the last page of this paper) or visit www.revsoft.co.uk.

Appendix A

OpenInsight's database structure compared to a typical Third Normal Form (3NF) RDBMS.

The following graphically represented scenario is an over simplified one, but clearly demonstrates just how quickly a typical standard RDBMS meeting third normal form rules can become over complex and inefficient. The same simple scenario will be used to demonstrate how a MultiValue DBMS, such as OpenInsight, remains concise and efficient.

The scenario shows a specific example with: 4 employees, 2 departments and 2 projects. The Revelation paper that this is taken from, details how a typical RDBMS would attempt to model the various types of relationships that are available, however we will only look at the table structure and summarise the indexes required. The full paper is available from Revelation Software on request.

We will use OpenInsight for our example, as it provides us with a simple data model based on the way that people in business think.

The business rules are:

- 👁 There is a business
- 👁 The business has EMPLOYEES
- 👁 The employees have PHONES
- 👁 Each employee has a DEPARTMENT
- 👁 Employees may be assigned to PROJECTS
- 👁 Each employee has a manager in a hierarchy up to the president.

The scenario results in:

- 👁 **OpenInsight** - 3 tables with 8 rows and 1 index with 10 values
- 👁 **3NF RDBMS** - 6 tables with 27 rows and 12 indexes with 48 values

Sample OpenInsight Scenario

3 tables with 8 rows and 1 index with 10 values

Employee	Name	MV:	MV:	Department	MV:	Manager	
ID		Phone	Phone Types		Projects	Is	Of
1	Van Wagner, John (aka Jack)	555 - 1111 555 - 2222 555 - 3333	Work Home Fax	Sales	Jones Smith	2	
2	Browne, Arlene	555 - 4444 555 - 5555	Home Work	Sales	Smith	3	1
3	Carruthers, Sandy	555 - 6666 555 - 7777 555 - 8888	Work Mobile Fax	Sales	Jones	4	2
4	Jameson, Kim	555 - 9999	Work	President	Jones Smith	4	3 4

Employees table: 4 rows, indexed on name (10 values)

Department ID	MV: Employees	Supervising Employee
Sales	1	3
	2	
	3	
President	4	4

Departments table: 2 rows, no indexes. Use this employee list as an index to Employee by Department.

Department ID	MV: Employees	Managing Employee
Smith	1	3
	2	
	4	
Jones	1	3
	3	
	4	

Projects table: 2 rows, no indexes. Use this employee list as an index to Employee by Department.

Sample Third Normal Form RDBMS Scenario, based on SQL.

6 tables with 27 rows and 12 indexes with 48 values.

Employee ID	Name	Department
1	Van Wagner, John (aka Jack)	Sales
2	Browne, Arlene	Sales
3	Carruthers, Sandy	Sales
4	Jameson, Kim	President

Employees table: 4 rows, indexed on Employee ID and also on Name.

Employee ID	Phone	Phone Type
1	555 – 1111	Work
1	555 – 2222	Home
1	555 – 3333	Fax
2	555 – 4444	Home
2	555 – 5555	Work
3	555 – 6666	Work
3	555 – 7777	Mobile
3	555 – 8888	Fax
4	555 – 9999	Work

Phones Table: 9 rows, indexed on the key [Employee ID and Phone] (9 values) and on Employee ID (4 values).

Department ID	Supervising Employee
Sales	3
President	4

Departments table: 2 rows, indexed on the key [Department ID] (2 values).

Project ID	Managing Employee
Smith	2
Jones	3

Projects table: 2 rows, indexed on the key [Project ID] (2 values).

Project ID	Employee ID
Smith	1
Smith	2
Smith	4
Jones	1
Jones	3
Jones	4

Project Employee table: 6 rows, indexed on the key [Project ID and Employee ID] (6 values), also on Project ID (2 values) and also on Employee ID (4 values).

Employee ID	Manager ID
1	2
2	3
3	4
4	4

Employee Manager table: 4 rows, indexed on the key [Employee ID and Manager ID] (4 values), also on Manager ID (3 values) and also on Employee ID (4 values).

Many databases locate data definitions in the same table as the data itself. They set aside some space at the beginning of the file, in which the data is described. However, when the developer needs to add new columns to the table, they have to recreate the file, copy all the data and perform other generally high maintenance chores. Other issues arise if end-users want to be able to do their own queries and see dates in their local format, assuming some are American and others English.

While this can be perceived as a 'developer only' benefit, the end-users benefit when it comes to modifying an already deployed application. Very often an end-user needs to have a new column (field) to capture a new piece of information, or a calculated column that automatically brings some data together, removes data entry duplication, or just makes data entry easier. With a Revelation-based system the developer can simply define the new column in the dictionary, drop the prompt onto the appropriate form and provide an update to the users. This enables developers to make extremely rapid changes to a bespoke system (customised for the business or purpose-built) and in doing so bring immediate business benefits to the desktop, rather than having to wait days while the system is changed, tested and deployed – saving on downtime and the cost of the change.

OpenInsight Table Builder

File Edit Indexes Details Equates Window Help

Table Builder - CUSTOMERS

Table Builder

Use this screen for creating or modifying dictionary column definitions.

Data Columns

Position	Column Name	Data Type	Key	Not Null	Default	S/M	Len
KEY 0		INTEGER	Yes	No	SEQKEY	S	14
1	LNAME	VARCHAR(255)	No	No		S	20
2	FNAME	VARCHAR(255)	No	No		S	20
3	ADDRESS1	VARCHAR(255)	No	No		S	20
4	ADDRESS2	VARCHAR(255)	No	No		S	20
5	CITY	VARCHAR(255)	No	No		S	20

Calculated Columns

Column Name	Data Type	Formula
COMPANY_XREF	VARCHAR(255)	CALL XREF(COMPANY),202C2D2F5C,"",1")
CUSTOMER_NAME	VARCHAR(255)	@ANS = {FNAME}""(LNAME)
CUSTOMER_NAME_XREF	VARCHAR(255)	CALL XREF(CUSTOMER_NAME),202C2D2F5C,"",1")
INVOICE_TOTAL	DOLLARS(15)	INV_NOS = (INV_NOS)total = slate(INVOICES), INV_NOS,

Page 18 of 31

Hashed Retrieval System

"Hashed Retrieval System" describes the mechanism by which Revelation-based systems locate rows that have been stored in the tables. It does not matter how well databases store information, if users can not retrieve their data, it is not worth storing!

There are several traditional ways of retrieving information from a database and most of them have some flaws. The simplest is "Sequential Access Method" and involves the system reading through the file until it finds what the user is looking for. There are application variants of this such as Indexed Sequential Access Method (ISAM) and BTREES but none are as efficient as Revelation's "Hashing" system.

As an example:

Think for a moment about when you want to get a client file out of the filing cabinet.

You may have 4 draws labelled A-G, H-M, N-T, T-Z.

You're looking for Otigbah.

Do you go to the first draw and start looking through sequentially? That's what a 'sequential access' method would do.

Do you go to the third draw and start looking sequentially? That's what an 'indexed sequential access' method would do.

Do you go to the third draw and glance at the middle file, seeing that it is for Robertson you know that Otigbah must be before the middle. So you glance at the middle file of the first half set of files. Then seeing this is for O'Doul you know that Otigbah must be after the middle of the first half. So now you glance to the middle of the top half of the bottom half... eventually you find the file. That's what a BTREE lookup would do.

Or do you, like our hero the hashed retrieval system, open up the third draw and go straight to the right file?

How does it do this? By an algorithmic design called "Hashing". Essentially the steps are simple:

When a file is created it is split internally into "data buckets" called "frames". Imagine you were given a file and a pile of names to store in the file. You need to be able to get them back quickly, so you could split the file into 26 data buckets and put the As in the first one, the Bs in the second, the Cs in the third and so on. Then if someone asked for Zachariah you would know to go straight to the 26th data bucket.

What you would have created there would have been what we programmers call a "mod 26 file", or to be more formal – a modulo 26 file. The modulo is simply the count of data buckets in a file.

Now hashing is a bit more complicated than just making 26 buckets but it effectively does the same thing.

Imagine the hashing algorithm as a Black Box. You tell it how many data buckets you have and you tell it the name of the person you want retrieved – Otigbah say. The Black Box then does some clever arithmetic and says – oh that will be in data bucket 123,456 so off you can go and get it, going straight to the place on disk where the information can be found.

This works because all of the frames are of a known size. Where the data becomes too large for the frame, an overflow frame is created for the user automatically.

There are other systems similar to Revelation's, but most need to have their files manually recreated larger from time to time to minimize overflow. However with a Revelation-based system, this is not a problem for users to worry about as Revelation use something called "Linear Hashing with Partial Expansion". That means that the files automatically resize as they are used, thus eliminating the maintenance overhead of periodic resizing and as a result reducing the cost of database administration.

Variable Length Data Structure

Most databases require developers to predefine the structure of a row before data is added to a table. This is normally referred to as a "fixed length data structure". To provide a real world example :- tacking Customer details. The developer would need to create a "Table" called say "CUSTOMERS" then tell the computer the structure that they wanted this table to have, for example

- ▶ Customer number – primary key – 10 digits
- ▶ Entry Date – DD/MM/YY
- ▶ Name – 30 characters
- ▶ Address Line 1 – 30 characters
- ▶ Address Line 2 – 30 characters
- ▶ Address Line 3 – 20 characters
- ▶ Town – 20 characters
- ▶ State – 2 characters
- ▶ Zip – 5 characters
- ▶ Phone – 12 characters

In the above example, each order recorded (normally called a "Row" but sometimes called a "Record") would be 10+8+30+30+30+20+20+2+5+12 or 167 characters long. Each separate piece of information recorded (e.g. Town) would be called a "Column", or sometimes a "Field".

This works very well when the row is of a known structure, as if a column is shorter than a definition, it can just be padded with spaces to ensure that all columns start and stop where expected.

So for example, customer 1234 recorded on the 12th May 2002 as Revelation Software at 99 Kinderkamack Road, Suite 109, Westwood, NJ 07675 on 201 594 1422 would be stored as:

```
1234▼▼▼▼▼12/05/02Revelation Software▼▼▼▼▼99 Kinderkamack
Road▼▼▼▼▼Suite 109▼▼▼▼▼
▼Westwood▼▼▼▼▼NJ07675201 594 1422
```

In the above example the system uses a total of 87 bytes out of 167 – so nearly half of the space on disk is wasted.

However developers run into problems if the data doesn't match with their preconceived expectations. So for example, say that the developer needs to use the same structure to store some information about a German client, say 9000 Ingenieurbüro für Projektentwicklung GmbH. at Gorkistraße 12 99084 Erfurt with a telephone number of 00 49 3615 62 18 51.

This would be stored as:

```
9000▼▼▼▼▼12/05/02Ingenieurbüro für ProjektentwiGorkistraße
12▼▼▼▼▼
▼▼▼▼▼Erfurt▼▼▼▼▼
▼▼▼▼9908400 49 3615 6
```

In the above example, the system uses a total of 72 bytes out of a total of 167 and more importantly the user loses data!

Worse still, if the system users need to enter multiple address lines, this information could be lost as well. This is very common when working with American written applications, or when dealing with American companies for mail order who are used to Street Address, Town, State and Zip. A classic example of this was when Revelation, in their old London office, ordered an item for delivery from a supplier using such a system with limited address functionality. Asking suppliers to record "Revelation Software Limited, Northumberland House, 11 The Pavement, Popes Lane, Ealing, London, W5 4NG found them getting quite flustered. And when you see what actually turns up on the printed label, you can see their systems have problems!

Naturally the developer could make all of the columns much bigger than they think their users will need, but it can often be hard to get this right. Imagine you were a developer working in America. For years, zip codes were 5 digits in length. You would just assign 5 characters to the column and that would be fine until a decade or so ago, when the Postmaster General decided that 5 was no longer enough and added a hyphen and another 4! Now ALL of the record structures are wrong and they need to be laboriously changed and the data read and rewritten, creating a headache for both the developer and the end-user.

This is where Revelation's "variable length data structure" excels. With this sort of structure, whilst the table is still defined, the length of the columns is not. Instead the developer just defines the position of the column in the record. When data is added to the row it is just recorded as it is typed, then when the end of the column is reached, the system puts in a special "End of column" character before starting the next column. Traditionally these are referred to as "Field Marks".

MultiValue Data

A major benefit to opting for a Revelation-based application is the use of 'MultiValue Data'. Taking the previous example, a Revelation developer will drop the multiple address lines (Address Line 1, Address Line 2, Address line 3) and simply use one MultiValue column called 'Address'.

So in the above example the Revelation developer would define the structure of the data table as:

- ▶ Customer number – primary key – field 0
- ▶ Entry Date – DD/MM/YY – field 1
- ▶ Name – field 2
- ▶ Address - MultiValued - field 3
- ▶ Town – Field 4
- ▶ State – Field 5
- ▶ Zip – Field 6
- ▶ Phone – Field 7

Normally in database applications, developers can only store one value in one column/row combination, but if developers can utilise "MultiValues", it means that a column can accept multiple values, with no limit to how many values there are. So in the example above, the developer would be able to get rid of three address columns and just have one that can contain the contents of the previous three. Each line of the address then becomes a "MultiValue", within a "MultiValued" column. These values can themselves be application variable in length, because the system will use another delimiter to show where each value ends! These characters are traditionally referred to as "Value Marks".

So let's look at how the two examples above would look in an application variable length, MultiValued filing system.

```
1234■12/05/02■Revelation Software■99 Kinderkamack Road²Suite
109■Westwood■NJ■07675■201 594 1422
and
9000■12/05/02■Ingenieurbüro für Projektentwicklung GmbH■Gorkistraße
12■Erfurt■■99084■00 49 3615 62 18 51.
```

In the above examples ■ is a field mark and ² is a value mark. Note that a lot less disk space is used and more importantly, no valuable data is lost.

Appendix C

Social Software Limited



The Challenge to build an XML Data Exchange

Youth offenders have been hitting the news headlines of late, not because of what they are doing, but because of concerns over how we are dealing with them in the UK. There has been a sharp rise in youth custodies since Christmas 2005, the trigger is the recently ratified 'ASBO' (Anti-Social Behaviour Order). A key element of the Crime and Disorder Act 1998, the ASBO was designed to put a stop to anti-social behaviour (for example, graffiti and littering, excessive noise late at night, intimidation, drunken behaviour and dealing in drugs) and was first put into use in 1999. Now more and more 'breach of ASBO' indictments are being directed at youths, while simultaneously the Youth Justice Board's (YJB) budget for custodies is being reduced.

The average ASBO costs more than £5,000 to process and three months to obtain. Over half of those sentenced in court for breach of an ASBO receive a custodial sentence. It is for these reasons that the courts and the Home Office agree that being 'locked up' should be a last resort in the case of an ASBO breach. Needless to say, the tackling of anti-social behaviour and indeed all aspects of youth offence are high on the agenda of both national government and local authorities, both of whom have taken a strong line to improve the individual assessment of each case to ensure that the punishment best fits the crime.

In fact, the Home Office has long been committed to improving communications between the various parties involved in criminal justice, running a series of pilot projects over the last two years aimed at improving the process, from the initial police notice, all the way through to the courts and after care. Technology has been identified as a key enabler for a smooth running, efficient system.

In a recent breakthrough, a pilot just completed in Warwickshire has given tremendous clarity to the measures that need to be put in place to improve the way youth offenders are dealt with in the UK.

The Challenge: To Achieve the Exchange of Data

The aim of the Warwickshire pilot was to actually show that it would be feasible for different agencies involved in the processing of offenders to have maximum visibility of all of the case facts on demand. This applies to lawyers, the police, the youth offending teams (YOTs) and the courts, among others. At the moment, each agency maintains their own information in separate case management systems, in a proprietary format. This information is exchanged via post, fax and telephone, resulting in a slow and inefficient system that is prone to inaccuracies and double, or triple entry of data.

YOIS Warwickshire Pilot

The Warwickshire pilot represents the culmination of around ten years of work concerning the most practical way forward with data exchange in the Criminal Justice system. Commissioned by the government's dedicated information technology department for criminal justice, CJIT, its objective was to link three of the organisations typically involved in the criminal justice process. By doing so, this would prove that any number of organisations in the chain could share information freely and openly. The three selected participants were the police, the courts and the Youth Offending Team.




Diagram 1: Spider infrastructure

To give some further background, existing arrangements consisted of a bi-lateral electronic link between the police and the courts. However, to link together every organisation in this way using point to point connections would be slow to develop, be very expensive and hard to maintain (diagram 1).

What was required was a completely new information architecture, centred around a hub that would allow any given organisation a way to link in and exchange relevant data from any other organisation on the intranet (diagram 2).

Youth Offending Team systems specialist Social Software Limited (SSL) was called in to take on the joint work with CJIT of creating the first three-way link. SSLs vast experience of working

with Criminal Justice and Youth Offending Team data was key in finalising the design of the message links. Of the country's 155 YOT organisations, SSL was already providing advanced case management solutions to 110.

Faced with the challenge of continuously developing and enhancing systems to accommodate new thinking and technologies, as well as supporting for example the complex needs of Youth Offending Teams, SSL chose to steer away from popular development database packages on the market, on account that they were not sufficiently rapid, or flexible, involved inappropriate compromises and unnecessary costs of enhancements in practice. Instead they opted for a solution based upon a MultiValue database environment called OpenInsight, which would essentially allow them to create a flexible, standards based facility ideally suited to meet the need to exchange rich XML data.

"The choice regards design environment was a cornerstone of the project. Looking forward to many tens of years into the future, factoring in scalability and increasingly complex data message requirements, the materials we started with were of paramount importance," stated Derek Girling, Senior Consultant at Social Software Limited.




Diagram 2: hub-enabled infrastructure

"When you are dealing with information about people, where there may be differing and valid views about the relevant facts, there are few set responses. The data tends to be text heavy and has few fixed rules. This presents special challenges regarding data structures and the enormous flexibility of the Revelation Software design product; OpenInsight, made it much easier to create a fluid application that would mould to the changing needs of data exchange into the future."

The YOIS pilot focussed mainly on the request and receipt of Pre-Sentence Reports (PSRs) for youth offenders. When a youth is found guilty, the magistrate requests a PSR from the YOT. A PSR request message is sent to the YOT and recorded in the database against the young person's case file and the corresponding YOT worker is alerted by the administrator to provide the PSR. Once the PSR is completed (following interviews and research), it is stored on the YOT server. A message is then sent to the hub to record the fact that the document is now available. When required the hub will request that YOIS generates an outgoing PSR message which gets passed through the hub to the Magistrates Court.

All of these communications take advantage of fully-encrypted electronic messaging. For example, an encrypted message is sent from the Magistrates Court via the Criminal Justice System Exchange hub. The YOIS system detects this new message and initiates a process to decrypt and consume the message before transferring data into the YOIS database.

"Communications between organisations that lie within the Criminal Justice System Exchange network need to be secure. Also the system design is such that a single encrypted electronic message can inform several people in different roles, bringing additional efficiencies," explained Girling. "This has obvious advantages over fax and mail-centric procedures, which can not guarantee that information is picked up and read only by the intended recipients, or indeed enable data to be reused e.g. for reporting purposes"

To ensure that all messages on the system could be read by existing, as well as future, computer applications, the YOIS system utilised XML (extensible Markup Language). XML provides a common language that enables all parties to agree and work on the contents of data message irrespective of their differing application environments.

The Benefits

The pilot has highlighted the numerous benefits of open systems communications for the criminal justice process. Primarily richer, more accurate, standards based sources of data, which in turn leads to better judgements and improved organisation of services.

Marc Radley commented; "with present procedures it is difficult to build an overall picture in order to fully evaluate any one case. For example, the police might cite a 'breach of ASBO' charge, recording all of the evidence, but not the thinking behind why that charge was made. A lawyer, however, is more interested in these circumstances. Single agency systems tend to present information according to the lowest common denominator of data available from a particular agency. When systems such as YOIS are enhanced and connected to a hub alongside several other systems, richer data becomes available and this promotes wider perspectives about using repositories of information in the future."

The CJIT hub is conceived as a 'hollow' hub, acting in a way similar to a telephone exchange routing requests from any point on a network to the correct locations. The advantages of this set-up are two-fold. Firstly, there is a security aspect, in that the hub is impervious to potential attack and acts of terrorism, as it doesn't physically hold any information. Secondly, this model enables the individual authorities themselves to retain ownership of their data, while at the same time-sharing it. The process can also move forward without incurring the time consuming and impractical task of mapping information and bespoke links for each disparate system.

From a logistical viewpoint, the process of connecting the YOIS system to the hub has demonstrated the approach is agnostic. The pilot successfully demonstrated the possibility to connect an organisation into the system at any stage of system or operational readiness. This affords a substantial degree of flexibility in rolling out the data exchange system nationwide, allowing this to be carried out incrementally. Essentially, it means getting data exchange up and running would not depend on all participants being ready to join simultaneously. In other words, the success of the system is not determined by its weakest link.

"Connecting YOIS to the hub in Warwickshire shows that we can deliver benefits with as few as three system participants, allowing for more and more participants to join as and when they are ready. This revolutionary concept for data exchange has potential for national, or even international, deployment," stated Radley.

The potential impact of enhancing and connecting YOIS is far reaching. Youth Offending Service data available via 'hubs' promises to deliver more complete information at a much faster pace to professionals involved at key points in the various services. This is bound to lead to better judgements concerning crime prevention, the response to risk of offending, the process of justice, correction and rehabilitation. This would mean more effective use of government funds and resources and could also result fewer incidences of repeat offending. Moreover, the integration of data and enhanced quality of information would assist in policy reform within the entire social welfare and justice system; these changes being backed up by reliable information and facts as opposed to poorly informed opinion or media pressure.

Summary

The Government's programme of reform and change in the criminal justice system includes a vision of a quicker, more integrated and effective process in a bid to provide more effective services and reduce re-offending. Integrating YOIS into CJSE exchange has been a clear signpost of the way forward, promoting further integration and the relevant agencies working together to help youth offenders, their families and their communities. The approach also demonstrates a strong lead about putting in place future proofed and cost effective data exchange also applicable to the associated domains of education, social services and health.

The Youth Offending Team is a part of a wider process of managing offenders, offending behaviour and victims involving local communities. CJIT's pilot project in conjunction with SSL in Warwickshire has proven that exchanging data with systems such as YOIS is part of the process of extending data available to the Criminal Justice System and the continuing success of government reforms aimed at promoting swifter justice and safer communities as well as saving the tax-payer money in the long term.

Above all else, CJIT and SSL have proved a point in fact – that it is indeed possible to open up the secure exchange of data for any number of organisations.

Appendix D



TV Shopping Network: Huge Growth With Small Costs

Australia's top TV retailer counts on Revelation Software

Normally, people don't put the phrases "control costs" and "expand IT function" together in the same sentence. When it comes time to talk about "expanding IT function," it's more likely you'll hear it paired with phrases such as "wow – that seems like an awful lot," or "you mean on top of what we've already spent?"

It's a tricky syndrome. As a business grows, its IT requirements grow and systems need to be upgraded and expanded. Poof! There goes your profit. Over the past few years, plenty of companies have sunk under the weight of their massive investments in technology infrastructure and IT services.

There's a company from Australia that knows how to come out on top of the IT investment conundrum: Australia's TV Shopping Network (TVSN). They've controlled costs even as their technology requirements have taken a quantum leap – and they credit Revelation Software with making it possible.

Television Shopping – It's a Hit!

In just eight short years, TVSN has grown to become Australia's leading television retailer with approximately 700,000 customers in its database. With gross revenues doubling every two years, the company is five times larger than its nearest competitor. Most importantly, TVSN is poised to start delivering significant profit to the bottom line.



"The trick is to grow the business without growing the cost base," says Rob Hunt, CEO at TVSN, "and for us, the key to keeping the cost base fixed is IT. Our IT platform generates efficiencies in the warehouse, the call centre, the way we order product, the way we analyze past sales and project future sales. At a most fundamental level, our IT platform manages our customer base and enables us to give them top quality service – and thereby maximize sales. And we count on Revelation Software to help us do all that."

Revelation – There from the beginning

TVSN grew out of a company called Demtel, an early player in television retailing. Demtel built its own customer database and retail operations applications using Advanced Revelation (ARev), the award-winning development environment for DOS. That was nearly 20 years ago. When TVSN was launched in 1995 and acquired Demtel, they engaged PowerForce Software – one of Australia's leading Revelation software developers – to upgrade and extend the system. PowerForce deploys a variety of technologies to suit their customers' needs, and Revelation products are among the favoured tools in their kit.

The launch of TVSN might have prompted some developers to recommend switching to a new platform, but not PowerForce. It was not so much an issue of cost, but simply that no technology was up to the task. "TVSN had a huge wish list of things to be done," said Kyle Amadio, Chief Systems Architect at PowerForce for Retail Systems. "We had to build the entire call centre and television studio management system and we had to do it in six weeks – just me and one other programmer. The only application environment that could handle it in the timeframe available was ARev."

Seamless Integration

Fast-forward a few years: TVSN has been growing rapidly and decides to upgrade their call centre with a new Interactive Voice Response (IVR) system integrated with a UNIX based Oracle database. The IVR system's primary role is to take orders and reduce the load on the call centre. A new vendor is brought in to implement the IVR system, but the project bogs down as costs rise and the vendor tries to convince TVSN to rebuild their entire technology platform. After several months, Powerforce is brought back to straighten things out.

"TVSN's entire operation lived in ARev," Amadio explained, "all their dispatching systems, finance systems, all the pricing information, the primary customer databases, the whole thing. Why would you want to junk all that just so you could add a new order processing system? You wouldn't. Because Revelation integrates so seamlessly and is such an easy environment to work in, we brought in the Oracle based voice system without junking any of the existing database, or applications. That not only saved TVSN a huge amount of time and money, it also avoided most of the risks associated with reworking their IT system."

That made TVSN's CEO, Rob Hunt, happy. "By the time PowerForce was done with it," he said, "we brought our outsourced call centre back in-house, we gained a whole lot more control, and we cut costs to a third of what they had been – which is just amazing! In fact, our new order processing system, EasyOrder, saves us about \$1.5 million each year."

Now, PowerForce is completing the upgrade from ARev to OpenInsight, the latest generation of Revelation's application development tools. "TVSN is increasingly using Oracle as the primary database," said Amadio, "but we're migrating the code that was previously written in ARev right over to OpenInsight. So there's an upgrade path from software that's nearly 20 years old to much more modern software. And the 16-bit moved to 32-bit, so there's no need to feel as though you've been left behind."

Enterprise Mix and Match

While most of TVSN's growth has developed organically as more viewers tune in, they've also made some strategic mergers with similar companies. In 2003, TVSN absorbed three other businesses and each one of them needed to be incorporated into the core TVSN system. Amadio says that's a perilous process for a lot of technology platforms – but it's one of the ways Revelation really shines. "Business integration, if it's not done properly, can be a disaster. In this case, we had only a few months to bring in data and business processes and modify the TVSN system to accommodate these additional businesses. If we had to sit down and rewrite all this stuff in JAVA or VisualBasic or C++ or another environment, you're looking at years. With Revelation, we had a very flexible tool set and a flexible database, so we were able to make it happen quickly with complete stability. We had no problem bringing the customer database from the other businesses online in the TVSN system. It was a seamless switchover." So it is obvious Revelation provides solutions that work enterprise wide.

Untangling the Web

In addition to the call centre, customers can also do their ordering on the TVSN Web site. In the early scramble to keep up with their wild growth, the various developers working for TVSN had ended up creating separate databases for the call centre, the Web site, the IVR and one of their early business acquisitions. Once again, Powerforce stepped in and straightened it all out with Revelation software, the web interface being managed by Sprezzatura's S/Web.

"They had three separate databases with duplicated business rules," Amadio explained, "which was driving everyone mad and costing the business a fortune, because they had to do everything three times. If you wanted to introduce a new customer loyalty program, well then you had to code it three times, and then you had to synchronize them. We re-engineered the Web site and integrated it with the existing IVR database so that the business processes, the discounting structures and the loyalty programs all work the same way, using the common set of PL-SQL stored procedures. OpenInsight drives all that code. OI is the glue that makes it all work. It sits in the middle, orchestrating the Web site and IVR and interacting with the database – and it's had a dramatic impact upon the business."

For TVSN, an important feature of a Revelation-based Web site is that TVSN's staff can easily update the site. "Their Web site is 100% dynamic," said Amadio, "and every part of it is user maintainable. The entire management of the site is in their hands. As soon as they have a product available to sell, it's on the Web – completely interactive with the database. So from an end-user perspective, the system is extremely flexible and offers them complete control. And Revelation is what makes it possible."


The Future Looks Bright!

"We're 47% up on last year, so we're still growing very quickly," says Rob Hunt. "In the next two to three years, our revenues should hit \$300 million to \$400 million, which will probably drop about \$30 million, or \$40 million to the bottom line."

To reach their goal, TVSN continues to count on Revelation. "We've got five members on our board of directors and every single one of them has an IT horror story from some other company," Hunt said. "With that thought in mind, the board's philosophy is 'There's nothing more important for this business than IT.' And their mandate to us is to 'make sure that we're always scaled for growth, and make sure that it's stable.' And that's the thing about Revelation – it's completely scalable, and it's just so stable!"

"We're at the stage now," Hunt continued, "where if our call centre, or Web site were to go down for just one hour, we could well drop \$60,000 or \$70,000 worth of revenue – which any company can ill afford to do. So we don't want to be the guinea pig for any piece of software. We only want to use what we know is proven, and the Revelation platform has just been so stable and dependable for us. Part of that has to do with how we've maintained it, and being careful about the way we've built upon it – and then the other part of it is just Revelation's inherent stability."

That's why Revelation products have been around for so long. And that's why Revelation will be the choice of growing enterprises for years to come.



OpenInsight for Workgroups - a repository-based applications development environment that enables development teams to work collaboratively to design, develop, deploy and scale high-performance business solutions. OpenInsight's complete tool set (including a Form Designer, Editor/Debugger, UI Workspace, XML Workspace and more), advanced BASIC+ scripting language and powerful client/server development combine to deliver a complete environment for deploying and evolving network applications. OpenInsight is delivered with a flexible Multi-Valued, variable length, filing system, or has the ability to use SQL Server, Oracle, Lotus Notes, Universe/UniData, D3, or any ODBC compliant data back-end. Delivered with both XML and Web-enabled tools, OpenInsight is a tremendously powerful development tool for Windows and Linux-based applications.

JOI - an integrated environment for managing the entire lifecycle of Java component based enterprise applications. JOI provides a single centre and integrated toolset for automating the design, development, assembly, deployment, customization, extension, and maintenance of Java component-based enterprise applications, tailored to the specific needs of the Revelation developer.

Revelation Software – Delivering feature rich MultiValue application development tools that empower both developers and users alike.

Experience the toolset at www.revsoft.co.uk



With more than 1.6 million users, at over 80,000 sites, on all seven continents, Revelation Software is in nearly every industry imaginable.

Revelation Software specialise in the provision of database development tools that permit applications to be developed in a fraction of the time required by traditional development environments. Revelation products may be scaled from single-user to large enterprise wide networks, from the desktop to the World Wide Web - for a fraction of the cost of other enterprise software solutions.

Revelation has a proud heritage, built on the provision of PC Database Software since 1982. During this time, Revelation has embraced the benefits of "post-relational" data structures that are now being more readily embraced by the mainstream.

Revelation applications are characterized by the enduring nature of the solutions implemented using its products. Revelation continues to invest in technologies that enable systems implemented in its earliest development tools to operate successfully on modern 32-bit client and network operating systems, and be extended to the Web - 20 years on, Revelation-based systems are still working hard for their owners.

Revelation continues to provide new tools that natively exploit the capabilities of the latest operating systems, network computing platforms, and programming languages and techniques. The combination of this focus and a commitment to continuous improvement and innovation has enabled Revelation to thrive and prosper in the challenging and ever-changing market of application development tools and deployment technologies.

"Social Software Limited has relied on Revelation for over 18 years. Throughout numerous changes in operating environments as well as unexpected changes in application requirements there has been no substitute tool which has offered the same development and maintenance benefits.

Thus we have developed an extremely mature software library including reporting tools which significantly extend the Revelation's environment in order to achieve a tight fit with our marketplace. As application developers we have complete confidence that Revelation will be a key part of our capability to satisfy any kind of application demand and challenge which typically exceeds customer expectations in our chosen domains."

Marc Radley, Managing Director – Social Software Limited

Revelation Software USA

99 Kinderkamack Road
Suite 109
Westwood, NJ 07675
United States of America

Toll Free: (800) 262-4747
Phone: (201) 594-1422
Fax: (201) 722-9815
Email: info@revelation.com

Revelation Software UK

45 St Marys Road
Ealing
London, W5 5RQ
United Kingdom

Phone: +44 (0) 20 8912 1000
Fax: +44 (0) 20 8912 1001
Email: info@revsoft.co.uk

Revelation Software Australia

Unit 2
21 Old Pitwatter Rd.
P.O. Box 300
Brookvale, NSW 2100
Australia

Phone: +61-2-9939-6399
Fax: +61-2-9939-6366
Email: info@revelationsoftware.com.au

- www.revsoft.co.uk -