

InDetail



Safyr

An InDetail Paper by Bloor Research
Author : Philip Howard
Publish date : March 2013

In our opinion, every major user
of an Oracle or SAP application
should be a Safyr customer

Philip Howard

Executive summary

Safyr (pronounced as in 'sapphire' and previously known as Saphir), from Silwood Technology, provides 'ERP metadata intelligence' to organisations that have implemented one or more of SAP, SAP BW, Oracle E-Business Suite, Siebel, PeopleSoft Enterprise, PeopleSoft EnterpriseOne (formerly JD Edwards) and which need to understand those applications in the context of data-oriented projects.

Before we discuss Safyr we need to consider what this means. Large ERP environments, such as SAP R/3, can be extremely complex. A typical installation might have 90,000 tables in the database. Moreover, the customisation done by each customer can add to this complexity: it is common for users to add thousands of additional tables to the environment and we know of accounts where this customisation is measured in tens of thousands. This huge proliferation of tables makes understanding the environment very difficult. Further, SAP enforces database constraints (primary and foreign keys and so forth) not within the database itself, but within its application server. This means that traditional modelling tools cannot capture all the relevant metadata about the system, because they work only at the database level.

Other technologies, such as data profiling tools and the connectors used by data integration vendors, suffer from the same drawback or can only work at the level of the standard R/3 environment without taking account of any customisation that may (almost certainly will) have taken place. This has a direct negative impact on complementary solutions such as business intelligence, master data management, data warehousing and test data management, which may need to know about the SAP environment. So there is a crucial need to get a deep and consistent understanding of the ERP environment that captures not just standard capabilities but also the modifications introduced by users.

Traditionally there have been two approaches to capturing this deep insight (metadata) into your ERP environment. One is, in the case of SAP, to implement SAP BW. We know of companies, for example, whose primary purpose for licensing BW is precisely to gain this level of understanding prior to loading relevant data into a third party data warehouse! The second possibility is to employ SAP services personnel to help. However, these people are typically

charged out at around \$2,000 per day and a significant amount of time will be taken up by their simply learning about any changes you may have made to the base system before they can actually deliver information of value. The other drawback with the latter approach is that it is only as good as the period during which the consultancy lasts. If you have a new project next year you will need the consultant to come back again. Safyr offers a third way that is, in our opinion, significantly more cost effective than either of these other methods.

Fast facts

Safyr provides a deep understanding of ERP environments including SAP R/3, SAP BW and all of Oracle's major ERP and CRM packages. It extracts relevant metadata into its own repository and then lets you sort, filter and manipulate that information for various purposes that we will describe in due course.

Key findings

In the opinion of Bloor Research, the following represent the key facts of which prospective users should be aware:

- Safyr works with SAP R/3, SAP BW, Oracle E-Business Suite, PeopleSoft EnterpriseOne (previously JD Edwards), Siebel and PeopleSoft applications, including industry specific versions of these environments. Silwood is a certified partner of both Oracle and SAP.
- Safyr extracts all relevant metadata (including any customisation) about source systems. Since the product is Unicode compliant this can be presented in any language. This includes not just information about tables but also information about referential integrity that may be defined logically (in the application server) rather than physically (in the database). Details about all relevant constructs are captured.
- Extracted metadata is stored in Safyr's own repository so that it can be investigated without impacting on the source systems. Repository management is provided to support extracted metadata from multiple applications or multiple versions (test, development and production) of a single application.

Executive summary

- There is no automated way for Safyr to know when source metadata has been changed. Processes will need to be put in place to ensure that Safyr is updated in the event that you add newly customised facilities to your implementation.
- You can search, sort and filter against metadata to create domain-specific subject areas either starting on a table basis or from an application perspective. A major feature of the latest release (version 6.0) is that you can filter out all tables that have no values in them.
- Subject areas can be visualised directly within Safyr but can also be exported to leading data modelling tools for use when building a variety of solutions.

The bottom line

Silwood Technology has no direct competitors and Safyr offers far and away the best method for extracting metadata from SAP and Oracle environments. Unfortunately, too many people are unaware of the time, money and effort that Safyr might save them. Or they think that they can achieve the same result using SQL (they can't). Or they have too much invested in their existing systems to be prepared to admit that they might have invested in the wrong technology. Or they are quite happy to pay \$2,000 a day for consulting services.

In our view, all of these people need educating. If you are a large enterprise and you use any of the relevant SAP or Oracle applications then you will almost certainly have data marts and warehouses, integration issues, master data management, archival or data migration projects—to take just a few examples—that would benefit from the use of Safyr.

The product

Safyr is currently in version 6.0, released in March 2013. It supports SAP R/3, SAP BW, Siebel CRM, PeopleSoft, PeopleSoft EnterpriseOne (previously JD Edwards) and Oracle E-Business Suite. Specialist products such as SAP Banking or PeopleSoft Finance are also supported. Safyr integrates with all leading data modelling tools including CA ERwin, Embarcadero ER Studio, Sybase (SAP) PowerDesigner, IBM Systems Architect, Silverrun, and so on.

The product runs on Windows platforms and its repository may be based on any of the leading databases: Microsoft SQL Server, Oracle or IBM DB2. Alternatively, the product ships with SQL Lite and you can use this if you prefer.

Rather than describe the features of Safyr per se, in the following sections we will explain how you would use the product in practice—the steps you would go through—and we will highlight various features of the product as we go along. In general, we will refer to SAP systems in the following discussions but the same principles apply to other ERP and CRM systems. Where there are any significant differences we will note these.

Extracting metadata

The first thing that you would need to do is to set up the database that will be used to host your repository. This will typically be done by a database administrator but, once accomplished, everything else within Safyr can be achieved by relevant users, typically data modellers, business analysts or information architects.

Safyr does not come with a reference model or template. In most environments this would be a bad thing but in this case the reverse is true. ERP environments are so heavily customised that a template-based approach has no value. Instead, once the connection to the source system is made it is literally a question of pressing the relevant button and going to have a cup of coffee while the software reads the data dictionary (in the case of SAP) in which the metadata resides. As previously noted, this is where SAP stores all its referential integrity information. Needless to say, Safyr uses ABAP code to integrate with the data dictionary. In Oracle environments PL/SQL is used. Note that the fact that the metadata is extracted into the Safyr repository means that there is no subsequent load on the source system.

One question that potential users of Safyr might ask is what happens if the meta-metadata (which is what Safyr reads) behind the application environment changes? There are answers to this. Firstly, Safyr will be updated appropriately. Secondly, this is extremely rare: for example, SAP has only made one change at this level since Safyr was first introduced in 2002.

Search, sort and filtering

Once the metadata is within the Safyr repository the next thing you will typically do is to search, sort and filter the information you have retrieved. For example, suppose that you want to know which of the more than 1,100 tables related to customers is the customer master? The way you would do this using Safyr would be to filter the data so that you were only dealing with tables that related to customers and then to sort by the number of child relationships that each table had, as you

The product

would expect the customer master to have more child relations than any other. The results of such a filter and sort are shown in Figure 1.

Tbl Name	Short Desc	Tbl Type	No. of Child Tbl	No. of Parent Tbl	Row Count
KNA1	General Data in Customer Master	Transparent	847	51	98
TS82A	Infotypes: Customer-Specific Settings	Transparent	123	2	3964
TBRF100	BRF: Application Class (Customer-Specific Features)	Transparent	110	5	17
T151	Customers: Customer Groups	Transparent	66	0	28
T171	Customers: Sales Districts	Transparent	46	0	40
T188	Conditions: Groups for Customer Classifications	Transparent	40	0	10
TINC	Customers: Incoterms	Transparent	40	0	93
TZONE	Customers: Regional zones	Transparent	30	2	338
T077D	Customer account groups	Transparent	28	2	140
THIT	Customer Hierarchy Types	Transparent	22	2	5
THV_APP_C	Customer Tab:Details for Hierarchy Visualization applica	Transparent	19	0	2
TPRIO	Customers: Delivery Priorities	Transparent	19	0	10
THVOBJCAT_C	Customer Tab: Object Categories for Application	Transparent	17	1	4
TSKD	Customer Taxes	Transparent	16	2	260
KNB1	Customer Master (Company Code)	Transparent	15	25	84
/DSD/H_RACVHD	DSD Route Settlement: Customer Visit - Header	Transparent	15	5	0
TWKT	Customers: Account Assignment Groups	Transparent	14	0	18

Figure 1: SAP tables referencing customer, sorted by number of child tables, with row count

You can, of course, drill down from this screen to each individual table, simply by clicking appropriately (Figure 2)

Fieldname	Posit	Fieldtext
MANDT	1	Client
KUNNR	2	Customer
LAND1	3	Country
NAME1	4	Name
NAME2	5	Name 2
ORT01	6	City
PSTLZ	7	Postal Code
REGIO	8	Region
SORTL	9	Search term
STRAS	10	Street
TELF1	11	Telephone 1
TELFX	12	Fax Number
XCPDK	13	One-time account
ADRNR	14	Address
MCCD1	15	Name
MCCD2	16	Name 2
MCCD3	17	City
ANRED	19	Title
AUFSD	20	Central order block
BAHME	21	Express station
BAHNS	22	Train station
EBBNR	23	Int. location no. 1
EBBNR	24	Int. location no. 2
BEGRU	25	Authorization Group
BRSCH	26	Industry
BUBKZ	27	Check digit
DATLT	28	Data line

Field Name:	PSTLZ
Short Desc:	Postal Code
Posit:	7 Key Flag: N Mandatory: N
Data Element:	Data Element Name: PSTLZ Short Desc: Postal Code
Domain Info:	Domain Name: PSTLZ Short Desc: Postal code Data Type / ERP: String / CHAR Length: 10 Decimals: 0
Descriptions:	Postal code This field contains the postal (zip) code for the house address (street and city). This entry is checked against the country-specific entry in table 005 (indicator PLZCTL).

Figure 2: drill-down showing primary and foreign keys

There are a number of additional points to make about this screenshot. For example, note the different coloured key symbols. These indicate where there is a primary or foreign key or both. This sort of information is really difficult to capture using modelling tools because referential integrity is not physically enforced but only logically. In addition, note the capture of short and long descriptions. As a comparison, the connectors typically provided by data integration vendors can similarly extract descriptions of tables but they cannot recognise primary or foreign keys so they are not much use when it comes to joins. Also, within Safyr you have the option to display these text descriptions in the language of your choice, which will be important for multi-nationals or overseas deployments.

The product

While base tables are the usual metadata type that you will work with, there are other options. In particular, you can work with database views or with data elements (such as account balance) and domains (such as money), where data elements are a subset of domains from which they inherit their characteristics. You can search against any of these metadata types independently. In some environments, notably JD Edwards (which has especially poor navigational facilities), Safyr will generate relevant metadata types even though these do not exist in the source, simply to aid in exploration.

Eliminating zeros

If you go back to Figure 1 you will see that there were over 1,100 tables listed. That is an awful lot. What Silwood has done in the latest release of Safyr is to introduce a filter that allows you to exclude all tables with no values in them, based on the row count. The way that this works is that Safyr reads the database catalogue from Oracle, DB2 or SQL Server to see how many rows there are in the table and then you can select a filter of "value = zero". You could, of course, select some different value to exclude also. The company is currently looking at how it could extend this to filtering at the column level, thereby not displaying attributes that are unused.

Create subject area

While you will likely use search, sort and filtering features for discovery purposes, most commonly you will want to work with a specific sub-domain of the overall data. In order to support this Safyr provides the ability to define a 'subject area' that you can work with in detail. To do this you select the relevant tables you want to work with and then there is an advanced search capability that allows you to choose associated tables. This includes extended filtering capabilities, such as searching for "fields with ..." or searching for particular field patterns. You can also opt to search either within the existing subject area or you can exclude the subject area from the search.

One issue that may arise is with respect to empty tables. That is, tables with no data in them. This is quite common in large ERP implementations. Clearly one would prefer to exclude empty tables from your subject area if possible. Safyr does provide the facility to drill into specific tables provided that appropriate permissions have been defined.

The product

Modelling

Safyr includes its own diagramming capabilities—see Figure 3—which you can use to visualise the subject area you have defined.

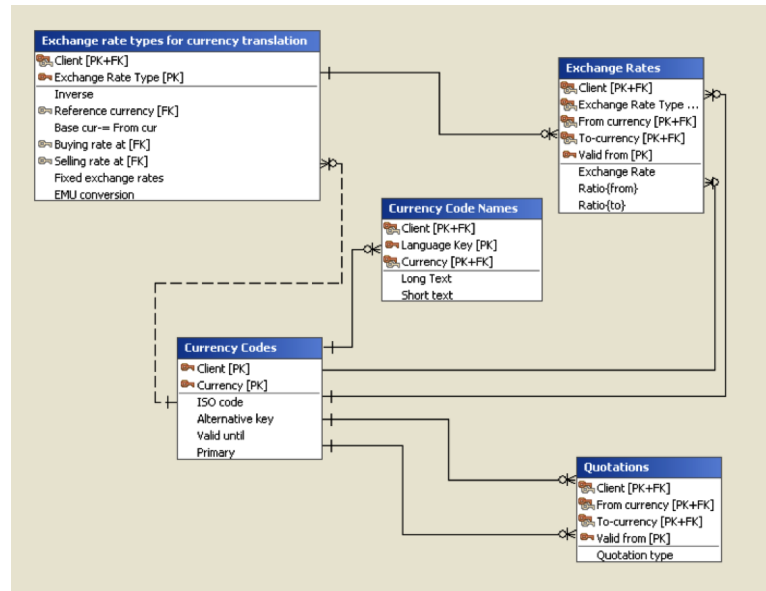


Figure 3: Subject area model within Safyr

Once you are happy with the logical and physical description of your subject area you can export the result into any of the leading modelling tools, using a wizard-driven approach. Of course, once in ERwin or some other tool then you can use this in turn to load data into, for example, a suitable data integration tool or as a rapid starter for constructing the schema of a relevant data mart. You can also export subject area details to Microsoft Excel.

Alternative approach

There is a different way of using Safyr. Rather than starting with tables you can start at the program level. That is, the tables in context to where they are used. In the case of SAP this means the SAP program tree. For Siebel it means working with applications, screens, views and applets at the user interface level and for other Oracle environments it also involves working with trees, though in this case each tree is owned by a specific application. Note that this approach works regardless of whether this is a general-purpose implementation or one that is industry specific, such as SAP Banking.

Using this method results are returned, in the case of SAP, by development class with function groups, programs and transactions separately colour coded, as illustrated in Figure 4. You can search against these based on either table views, by object class or by transaction codes. These last are the shortcuts used to do things such as create a sales order and from there you can track back up the hierarchy or down to base tables, which you can then use to create your subject area.

The product

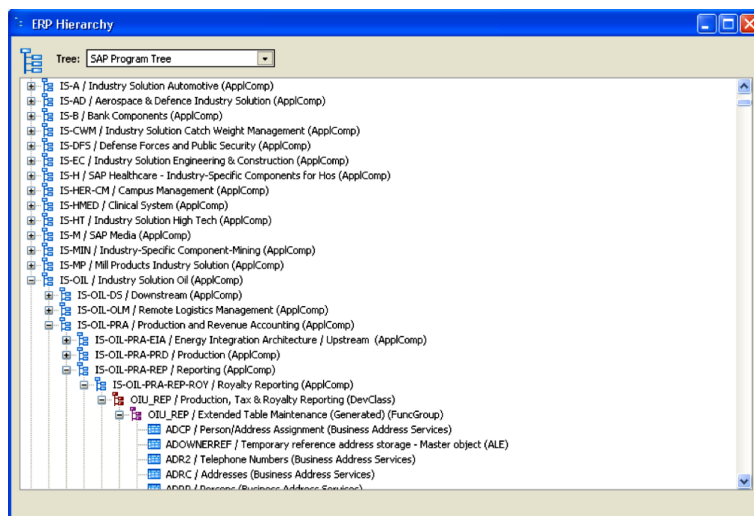


Figure 4: SAP program tree – using an alternative approach

In the case of Siebel, this system uses business components, which Safyr treats as views. However, Siebel does not allow any further drilldown to the table level so you cannot use this method for building subject areas though this technique may be useful for other exploratory purposes.

Repository management

Safyr includes a repository manager that includes support for multiple repositories, which you may have either because you have multiple ERP or CRM implementations that you want to examine or because you have distinct development, testing and production environments to analyse. One notable feature is that you can compare subject areas and differences will be detailed in a text-based report. More generally, you can access the repository using third party tools for reporting purposes. You can also export the contents of the repository if required.

The vendor

Silwood Technology is a British company that is privately owned and funded. It was founded in 1992 and was originally a reseller for Logic Works and its ERwin data modelling tool. Indeed, Silwood built the UK and European market for ERwin. However, after Logic Works went public in 1995 and, more particularly, after it was acquired by Platinum Technologies in 1998 (and later by CA), Silwood Technology sought to establish its own unique position in the enterprise modelling space. As a result, in the late '90s the company was looking for new markets. It identified integration with major packages such as SAP and Oracle as a significant issue and started work on what ultimately became Saphir and which has recently been renamed Safyr (the company wanted to trademark the name and it turned out that Saphir was not available, hence the name change). This was first released in 2002.

Silwood uses both direct and indirect routes to market. In the latter category the company has a number of blue chip partners, both resellers and OEMs, notably CA, IBM, ASG (Allen Systems Group), Adaptive and Pitney Bowes (in the Sagent ETL tool) as well as go-to-market partners such as Sybase, Informatica, Embarcadero, Grid-Tools, IPL and Visual Metrics. Also, as previously mentioned, Silwood is a certified partner of both SAP and Oracle and has achieved the Informatica Marketplace Seal of Approval. It also has overseas partners in both Canada and the United States. In total, the company has approaching 500 customers worldwide.

Web site: www.silwoodtechnology.com

Summary

It is not often that one can give an unequivocal recommendation for a product or technology. There are nearly always some sort of mitigating or offsetting factors. However, in this case there are none. Silwood Technology is unique in having recognised, and acted on, the need for a better understanding of enterprise applications, notably the dominant ERP and CRM solutions provided by SAP and Oracle. Having this depth of understanding is critical for many enterprise projects, ranging from business intelligence to data governance, from data integration to master data management, and from data migration to application development. Moreover, the company's discovery capabilities are not stand-alone but act as complementary technology to the established tools and products that are typically used to provide such solutions. Working in combination in this way should enable faster and easier implementations and will help organisations to gain more return on their investment in their enterprise applications.

In our opinion, every major user of an Oracle or SAP application should be a Safyr customer.

Further Information

Further information about this subject is available from
<http://www.BloorResearch.com/update/2159>

Bloor Research overview

Bloor Research is one of Europe's leading IT research, analysis and consultancy organisations. We explain how to bring greater Agility to corporate IT systems through the effective governance, management and leverage of Information. We have built a reputation for 'telling the right story' with independent, intelligent, well-articulated communications content and publications on all aspects of the ICT industry. We believe the objective of telling the right story is to:

- Describe the technology in context to its business value and the other systems and processes it interacts with.
- Understand how new and innovative technologies fit in with existing ICT investments.
- Look at the whole market and explain all the solutions available and how they can be more effectively evaluated.
- Filter "noise" and make it easier to find the additional information or news that supports both investment and implementation.
- Ensure all our content is available through the most appropriate channel.

Founded in 1989, we have spent over two decades distributing research and analysis to IT user and vendor organisations throughout the world via online subscriptions, tailored research services, events and consultancy projects. We are committed to turning our knowledge into business value for you.

About the author

Philip Howard
Research Director - Data Management



Philip started in the computer industry way back in 1973 and has variously worked as a systems analyst, programmer and salesperson, as well as in marketing and product management, for a variety of companies including GEC Marconi, GPT, Philips Data Systems, Raytheon and NCR.

After a quarter of a century of not being his own boss Philip set up his own company in 1992 and his first client was Bloor Research (then ButlerBloor), with Philip working for the company as an associate analyst. His relationship with Bloor Research has continued since that time and he is now Research Director focused on Data Management.

Data management refers to the management, movement, governance and storage of data and involves diverse technologies that include (but are not limited to) databases and data warehousing, data integration (including ETL, data migration and data federation), data quality, master data management, metadata management and log and event management. Philip also tracks spreadsheet management and complex event processing.

In addition to the numerous reports Philip has written on behalf of Bloor Research, Philip also contributes regularly to IT-Director.com and IT-Analysis.com and was previously editor of both "Application Development News" and "Operating System News" on behalf of Cambridge Market Intelligence (CMI). He has also contributed to various magazines and written a number of reports published by companies such as CMI and The Financial Times. Philip speaks regularly at conferences and other events throughout Europe and North America.

Away from work, Philip's primary leisure activities are canal boats, skiing, playing Bridge (at which he is a Life Master), dining out and walking Benji the dog.

Copyright & disclaimer

This document is copyright © 2013 Bloor Research. No part of this publication may be reproduced by any method whatsoever without the prior consent of Bloor Research.

Due to the nature of this material, numerous hardware and software products have been mentioned by name. In the majority, if not all, of the cases, these product names are claimed as trademarks by the companies that manufacture the products. It is not Bloor Research's intent to claim these names or trademarks as our own. Likewise, company logos, graphics or screen shots have been reproduced with the consent of the owner and are subject to that owner's copyright.

Whilst every care has been taken in the preparation of this document to ensure that the information is correct, the publishers cannot accept responsibility for any errors or omissions.



2nd Floor,
145-157 St John Street
LONDON,
EC1V 4PY, United Kingdom

Tel: +44 (0)207 043 9750
Fax: +44 (0)207 043 9748
Web: www.BloorResearch.com
email: info@BloorResearch.com