



magnaview

visualize anything • visualize everything

Technical Information : MagnaView 4.1
01 -10-2009

Technical information 4.1

1. General	3
2. Extraction, Transformation, Loading	5
3. In-memory analysis	7
4. User interface	9
5. OS	10
6. Security	11
7. Enterprise Suite components	12
8. System administration	15
9. IT-management	17

1. General

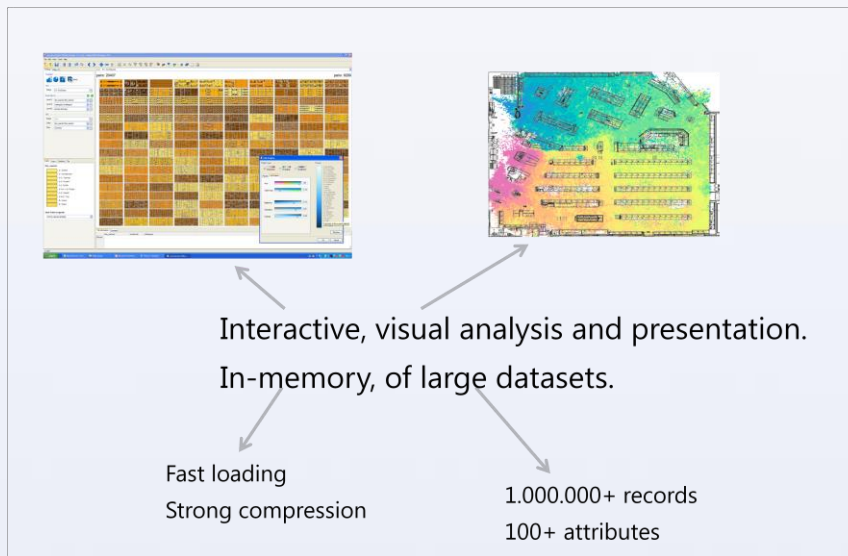
What are the most important technical objectives for MagnaView?

To start with, our technical objectives are not stand alone principles: they follow from the business needs of our customers. Our customers want to impact from their tools, and therefore real insight in their massive amounts of data, and they want it fast. Less effort and more impact.

The first technical consequence is that we are constantly improving performance in all aspects of our applications. Think about loading data, rendering visualizations, improving memory use, client-server communication, etc.

The second consequence is that we devote a lot of time to the user interface of our applications. Users should be able to achieve what they need in as little time as possible.

To achieve 'Real insight' MagnaView offers visualization and easy interaction. In turn, to be able to offer this interaction in real time, again requires careful software engineering. Algorithmics and complexity analysis are a necessary part of our work.



What programming languages are used at MagnaView?

We do our programming in Delphi. We are currently working with the latest 32-bit version of Delphi, and we're eagerly awaiting the 64-bit version, because this will give us again extra speed, and possibilities to address more internal memory, that is, the possibility to give the user access to more data. We use assembly language for specific parts of MagnaView.

What other tools are used to ensure the technical quality of the software?

For development purposes, we have built our own loggers and tracers, in order to monitor and improve performance. Besides, we use a host of applications for bug tracking, version management and prevention of memory leaks.

What's the technique for the Web version?

For our Web client, Web Live, we strongly feel users should not be bothered with add-ons or plug-ins. These require time to install, create unnecessary waiting times when loading, and add little functionality for the user. Moreover, add-ons and plug-ins will result in less performance and are harder to integrate in ways the user does not notice. They may also require system administration. Therefore, the MagnaView server component just communicates with the client by means of simple file transfer, notably .css, .html, .js and .jpg. files.

If you start MagnaView in a browser, the server component is started, and in no-time the user can start working. You can really feel the difference, in particular if you compare MagnaView to other applications that use add-ons and plug-ins that need to be started.

What browsers does MagnaView support?

MagnaView supports all major browsers, in particular IE6, 7 and 8, Firefox, Chrome and Safari.

Does MagnaView make use of instruction sets for the graphics card? That would really speed up the process of visualization...?

We've experimented with OpenGL, but we found out that it wasn't faster than our own rendering engines. Using our own engine is more flexible and scalable, and therefore, we've decided not to use OpenGL or languages alike. This also has the advantage that users don't need special hardware for running MagnaView software.

What else do you do to achieve more performance for the user?

For each view a thumbnail is created, and this is stored in a cache file. When starting a project, this cache file is accessed, and the user can immediately start working.


Also, when working in a Windows environment, loading MagnaView executables and data files from a network share takes less time because these are compressed. This reduces network traffic, which is good for IT-management, but it also reduces the time the user needs to start the application.

2. Extraction, Transformation, Loading

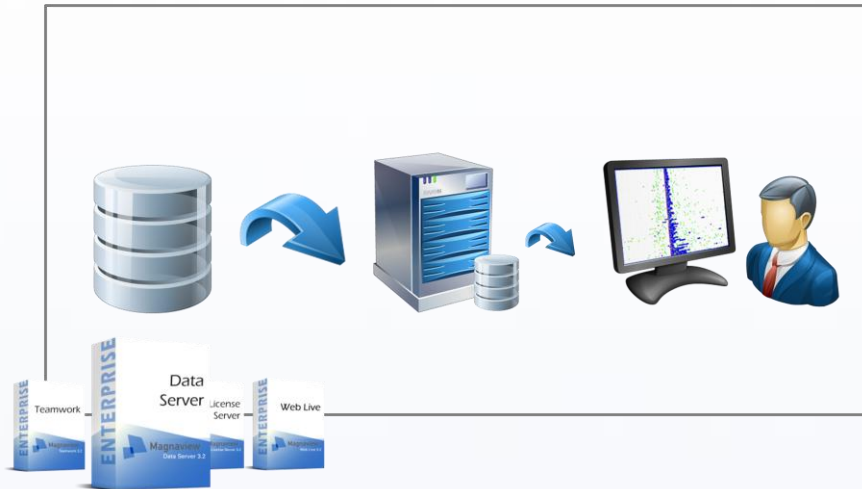


What data sources can MagnaView access?

MagnaView can access all data source types that are used in practice such as ODBC (table, query), Access, Excel, TXT and XML. Data sources can be accessed from a local drive, a network share or a url.



In an enterprise environment, we recommend using *MagnaView Data Server* for extraction of data. Data server takes data from some source, for instance by means of a query to a database, and converts it to a native data source. Scheduling the extraction process with *Data Server* compresses data and encodes it. Our native data format (.mvp) achieves better compression than zip.



How does MagnaView achieve this compression?

Just because we use internal memory as efficiently as possible, information *about* the dataset is used effectively. When a native data source is generated, a kind of dump of this efficient internal memory is made. When using a native data source, the source can therefore be loaded in internal memory in no time.

Just because the information about the dataset is used, an mvn data source compresses data more efficiently than for instance a compression algorithm like zip achieves. The more regular the input, the more efficient.

Did you do any tests to compare MagnaView to other applications?

For most of the datasets we see, if we create zip files, we mostly find MagnaView native sources to be less than half the size of the zipfile which contains exactly the same data.

We also ran tests to find out how fast our extraction is. We made a comparison with a number of other treemapping tools, and with Excel, and found that MagnaView was the fastest by far when loading data. In comparison to other treemapping tools, MagnaView is at least 5 times faster, and in comparison to Excel 2 times faster.

How does MagnaView support the Transformation of data?

MagnaView is a strong tool for visually and interactively exploring data. The user filters data, selects attributes, adds computed values etc. Since this is *precisely* what a user does when transforming data, we decided to have MagnaView itself be used for ETL. A MagnaView project – or more precisely the visible attributes in a view in a project – can be used as input for another MagnaView project.

This way, the user does not have to learn using an extra ETL- tool. Using MagnaView's intuitive user interface, which the user knows for exploring and analyzing data, is all that is needed.

3. In-memory analysis

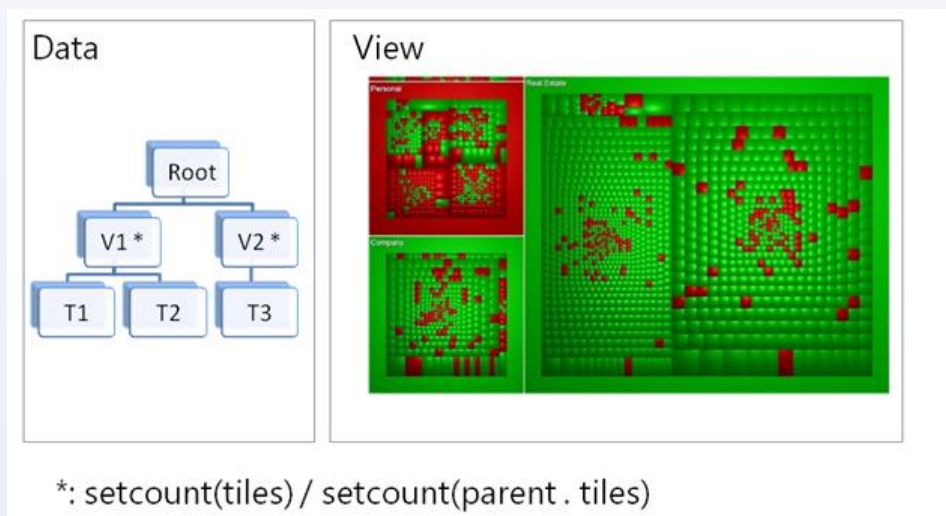
In-memory analysis is a big phrase in the market. What is its role for MagnaView?

It's crucial, and we set up the architecture of our applications to reflect this. In general, more and more internal memory is available to end users. Users know in-memory analysis speeds up their work, and therefore, they demand their analysis to take place in-memory. No more waiting for querying data, please.

In MagnaView, users can easily do analyses with hundreds of thousands, up to millions of records, with hundreds of attributes per record.

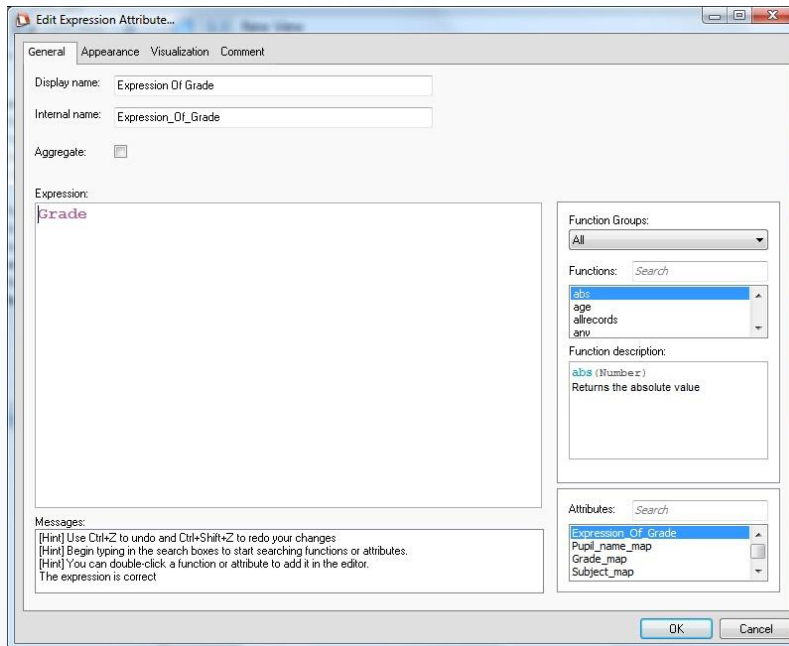
Does in-memory stop with the data being available as such?

No, for many applications in the market, in-memory means that just *data* is available in-memory, and some basic possibilities for user interaction. For MagnaView in-memory implies that complex computations should be possible on the data, real-time. For this purpose, we have a declarative expression language.



On one hand, this expression language contains regular arithmetic, statistical etc. functions. On the other hand, an analyst can access the full hierarchical structure of the data underlying a view.

These expressions can be evaluated in an ELT-step, but just as easily in-memory, in real time, while the user is working on his analysis. We have examples with amazing complexity, which run in real time.



Isn't working with the expression language difficult for analysts.

The functions in the language are easy to understand. Analysts will recognize functions for variance, the left part of strings, etc. These are just like the primitives in Excel or reporting tools. Only in case of complex computations is access to the hierarchical structure necessary.

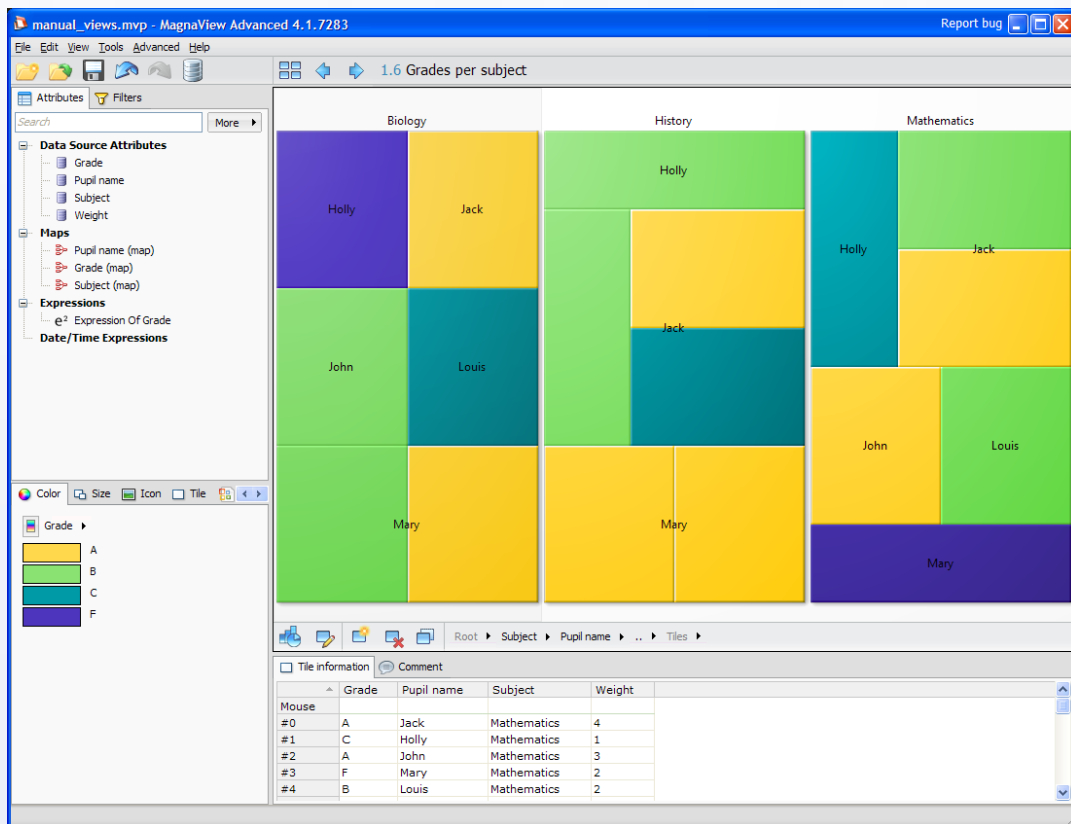
Besides, computations in spreadsheets often lead to errors because data is referred to as groups of cells. You can easily select or specify the wrong group of cells with your mouse. Trying to find such errors with complicated '='-references costs massive amounts of time for an analyst using a spreadsheet application. In MagnaView, the data is referred to explicitly. It's easily possible to track even the most complicated computation.

4. User interface

Why doesn't MagnaView use the Windows ribbon style in its GUI?

For version 4.0 and beyond, we have put quite some effort in the user interface. We've experimented with the ribbon style, and found out that the disadvantages of the ribbon interface also hold for MagnaView. It's not intuitive for the MagnaView user where he can find a specific option, and, even when the user knows where to look, too many mouse clicks are necessary.

That's why we've redesigned the interface completely. Options the user needs often are readily available, and context menu's and pull-down menu's give access to other functions.



5. OS

Why Windows?

For in-memory analysis on work stations, the data needs to be stored in the system memory of the clients. Nearly all clients are Windows-based, and therefore we only support Windows at this moment. The clients of the web version, Web Live, are of course independent of the operating system.

Is a Linux version of MagnaView available?

See the previous question. MagnaView can run under Linux when Wine is used. We intend to support this officially in the future.

How about virtual servers and appliances?

The process of a MagnaView Enterprise Suite installation only takes a small amount of time, and, all components are readily available. Therefore, bringing MagnaView Enterprise Suite as a virtual server or appliance to our customers didn't add a lot. Moreover, this route may have additional costs for our customers which could be prevented.

When our customers require MagnaView as a virtual server, the installation takes place in a virtual server which our customers make available.

In the near future, if other applications would be integrated with MagnaView, an appliance is an interesting option

How about Terminal Server?

MagnaView can run on a Terminal Server; some of our customers do. This requires a fast connection to the Terminal Server, because in a visualization tool a lot of visual information is sent to the users. Besides this, care the server needs to be able to carry the load of running MagnaView sessions. The load depends on the number of sessions and the size of the data.

6. Security

Can MagnaView authorize users for specific parts of data or for specific projects?

In general, we strongly advise our users to make use of the authorization they have in place in their network. Duplication of authorization in both the network and specific applications, like MagnaView will lead to less control and more involvement of system administrators.

OK, but is it possible to have MagnaView depend on who's using it?

Yes, for instance in benchmarking applications, users need to see their own data and compare this to the averages of other users in the benchmark without seeing their actual data. Various mechanisms are possible for this purpose such as call backs to authorization servers.

How does MagnaView protect data from unauthorized access in general?

In a Native data source data is always compressed and encoded. You need MagnaView to open native data sources. Besides, it is possible to encrypt a data file, that is, to specify a pass phrase that is used to encrypt the data, and which the user needs to open the data file.

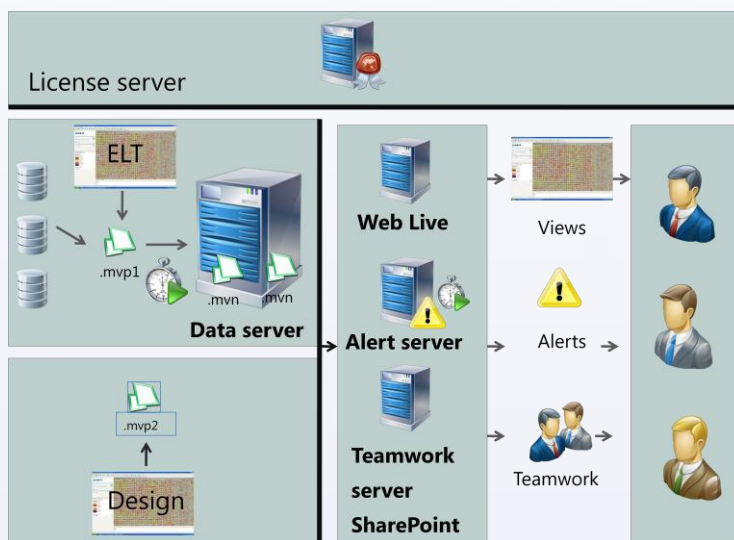
Parts of the data may be present in Project files. Project files are compressed and encoded and again, MagnaView is needed to open project files.

7. Enterprise Suite components

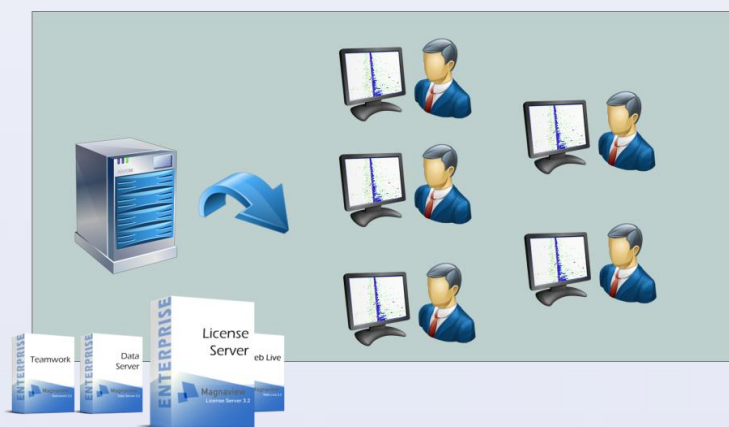
Suite. What Enterprise Suite components are available?

For Enterprise use, the following components are available:

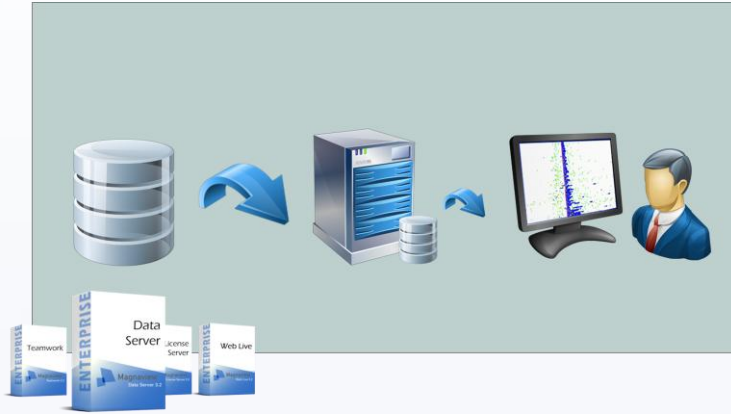
- License server
- Data server
- Web Live
- Teamwork server
- Alert Server (to be released 2010)




- In general a *License Server* gives users access to the MagnaView software. This can either be one of the Server components like Web Live, or one of the Windows Client products like Designer Pro. The advantage is that we can license for instance 20 concurrent users to use either Web Live or the Windows viewer.



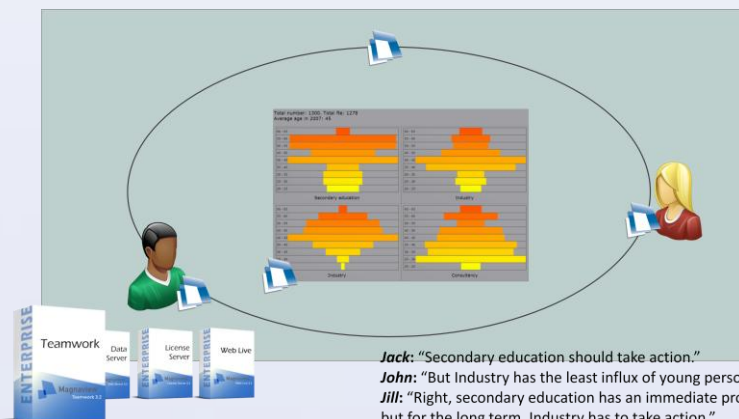
- *Data server* extracts data from some data source, and converts it to other formats like Native. Data server processes can be scheduled to run complex queries for instance outside office hours, and store the result in our native format to make sure users have fast access to data.



- The next server-based component is *Web Live*. Web Live is the server component which communicates with Web clients.



- The Teamwork server is the component which is used to allow users to communicate about views. All comments on views, or discussions about views are stored in the Teamwork database, and are available for users.

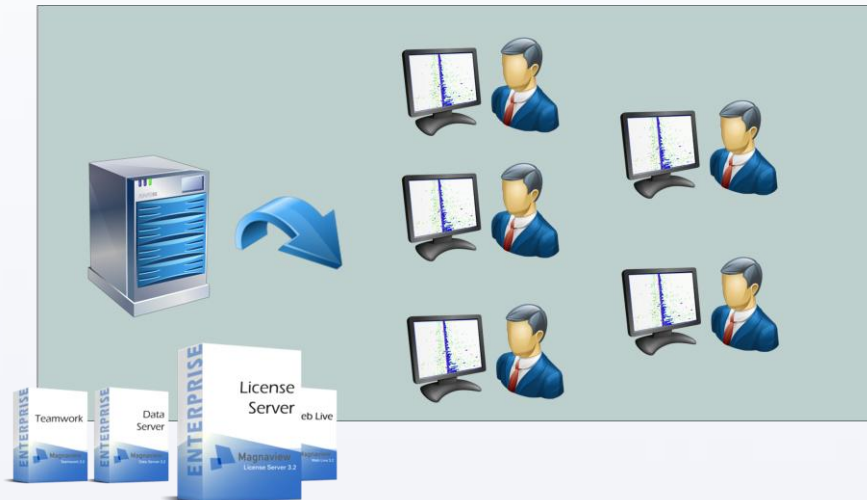


- The Alert server is under construction, but it will be a server-based component which alerts users in cases defined in projects.

8. System administration

How can Enterprise Suite be licensed?

Licensing can be either on a named user licenses or on the basis of concurrent use. Moreover, licensing can be set up in such a way that it doesn't matter whether concurrent use is made of Web Live or the Windows Viewer. For heterogeneous use, like a combination of a Windows and a web-based client, this reduces the number of licenses that is required.



Each installation of Enterprise Suite comes with License Server. For each installation, MagnaView delivers a .lic file, which contains information on licenses for software and specific products for this installation.

How easy is installation?

Installation of the complete Enterprise Suite is possible in less than half an hour, provided the necessary rights are available.

How does upgrading of software take place?

Upgrading is a matter of running an installer, and only takes a couple of minutes.

How does upgrading of special products take place?

Again, upgrading is a matter of running an installer. Some care is necessary if the customer has made changes to project files that are part of the previous installation, but careful system administration, e.g. moving customer specific files to customer specific parts of the network, will prevent problems.

Does MagnaView offer Managed Services?

If a customer wishes to use MagnaView without the burden of administration, MagnaView offers a Managed Service. The MagnaView software runs on a MagnaView server, and connections to the server park of the customer are set up to retrieve data. For end users, there is no difference between a local installation or managed service.

How do I set user rights for projects and data?

We strongly recommend using the authorization mechanisms that are in place in the network already. If a user has access to specific data, then that authorization should best be used.

In more complex cases, various possibilities are available.

9. IT-management

What do you do to reduce network traffic?

Compression is used whenever this has a positive trade-off for the user. In practice, this means anything is compressed and encoded. When using a Native data source, data is compressed, most of the times with a larger compression rate than zip. All MagnaView projects are compressed. Executables are compressed as well.

Moreover, while compressing data and projects, we also encode the content of the data source and projects.

Compression leads to more speed, better user experience, and less network traffic, and adds to security.

Traffic between the Web Live server component and the Web Live client may also contain information that should not be made public. Therefore, this is encoded as well. Besides, it is of course possible to use https-connections.

What does a server park needed to run MagnaView Enterprise Suite look like?

This depends on the (i) sizes of datasets/use of internal memory (ii) intensity of use (iii) number of users. Servers that are common in the market place dd autumn 2009 are generally sufficient to accommodate 20 concurrent users/sessions. That's also the maximum number of users for one registered version of Web Live.