Enterprise Photogrammetric Production Using TerraShare

DANIEL WUESCHER, Z/I Imaging, Madison
FRANK NEIDHART, Z/I Imaging, Aalen

ABSTRACT

The positive results of increased production throughput achieved with TerraShare as the workflow and geo-data manager will be presented. New technology, recently added to TerraShare, including distributed processing and pluggable components, turns TerraShare into a powerful enterprise-wide processing environment. The enabling technology will be presented, along with comparisons between traditional serial production and enterprise, distributed production.

1. INTRODUCTION

As sensor technology improves, and high resolution data becomes increasingly more affordable and popular, many production shops are encountering problems supporting the massive amounts of raster, elevation and vector geo-data being generated. The shear volume of data is creating difficult problems for users over the entire enterprise.

The Information Technology (IT) department needs an effective data management solution – a solution which will let its personnel manage terabytes of data, using all of the available enterprise assets, without affecting the performance of the rest of the shop. The production department needs an efficient data access solution – a solution that makes it easy for operators to find and use all of the geo-data necessary to perform their day-to-day activities. Supervisors need a convenient data visualization solution – a solution that makes it easy to quickly and accurately assess the quality of the products being produced. Finally, customers and partners need a faster and more efficient data distribution solution – a solution that allows authenticated users outside the production shop firewall to review, and optionally participate in, the production process.

Similarly, the ever increasing complexity of software and processes within a modern production shop combined with the competitive pressures to constantly improve throughput, lead to yet another set of difficult problems. Production supervisors need better tools to plan and manage complicated production processes. Operators need a simple, consistent environment where the data and the processing commands are easy to find and simple to use. Management needs a convenient mechanism to obtain up to date progress and status. All the while, customers and upper management are demanding increased production and lower costs.

Some or all of these problems can be found in nearly every photogrammetric production shop, from small private ortho photo producers to large government agencies. The need to efficiently manage data, and the processes that act on that data, is becoming increasingly important.

In many cases, IT may attempt to deal with the data challenges by purchasing additional disk space. Although additional disk space may be necessary, simply adding more gigabytes of storage doesn't address the data organization problems. Indeed, the lack of data organization can result in an enormous waste of existing disk space. The cost of maintaining multiple copies of large data sets on several different computers, or the time wasted in moving that data around, either via the local

network or via media such as CDs or DVDs, cannot be overcome just by adding more disk space to an enterprise. In fact, when users perceive available disk space to be essentially unlimited, the data organization problem can be exacerbated as multiple "personal" copies of datasets proliferate over the enterprise.

Increasing available disk space also does not solve the basic data organization problem of assembling disparate physical data (files and metadata) into a single logical entity. Photogrammetry data is often distributed over a large number of files, all of different formats and content: raster files, elevation files, vector files, header files, project files and/or databases, metadata files and/or databases, etc. Even though we would prefer to deal with a simple logical data structure that is consistent with our production workflow, we are forced to locate all of the various physical files and databases that contain the information we need.

In the same way, production supervisors and controllers have recognized production and workflow management concerns and have attempted to address these challenges (with varying degrees of success) using an ad-hoc set of procedures, scripts and batch files. Although their solutions are often inventive and creative, there are some inherent drawbacks in such an approach. "Home-grown" software can be expensive to write and maintain. If user programs or scripts are built on top of an application, integration with that application is often times tenuous, especially if the application doesn't support a well defined and well maintained interface or software development kit (SDK).

As for the overall production process, establishing a well defined process for the production workflow is obviously a good practice. However, controlling or enforcing that process workflow is more difficult. Operators are expected to know what the process is and there is little or no guidance provided by the system or the software. As such, operators are forced to switch to different application environments and must know what scripts and batch files to run in order to get the "right" output.

TerraShare, a software suite of powerful data and production management applications from Z/I Imaging Corporation, offers a large economic benefit to businesses by reducing the inherent production costs described above. In the next sections, we will outline various aspects of the software and highlight the benefits of each as compared to traditional or standalone workstation approaches.

2. TERRASHARE

TerraShare is a client/server solution that enables users to manage all of their geo-data over their entire enterprise. The key component in this solution is the virtual TerraShare file system. By creating and managing a collection of virtual folders and files, users are able to organize and view complex collections of physical geo-data as single, logical entities. Instead of directly using the physical files dictated by the computer operating system, users and their applications interact with logical files that have meaning in their production workflow.

TerraShare consists of several server and client side modules. The primary TerraShare server module contains the database and business logic used to track and manage the physical data as well as all of the logical data and metadata. The primary client module, the TerraShare Explorer Plug-in, is used to view and manipulate TerraShare files and folders. The TerraShare Explorer Plug-in is implemented as a Microsoft Windows namespace extension. Because of this, in many ways,

TerraShare files and folders look and act much like Windows files and folders. However, in addition to providing standard Windows-like functions such as cut, copy, paste, move, etc., TerraShare files and folders offer a myriad of other functionality designed to improve the production performance over the entire enterprise.

2.1. Data Management

2.1.1. Data Abstraction Layer

TerraShare provides a data abstraction layer that creates a separation of logical photogrammetry data from real physical files. TerraShare gives users the ability to group one or more physical files with metadata stored in a database to form a single logical entity, called a TerraShare file. TerraShare files can then be organized into TerraShare folders, similar to Windows' files and folders. All TerraShare files and folders are presented in the client with a Microsoft Windows file explorer namespace extension. This data abstraction layer is the key feature of TerraShare. Although, the concept itself is very simple, the organizational framework that results from this grouping is quite powerful.

One consequence of creating the data abstraction layer is that ordinary users can be separated from all of the IT issues associated with maintaining terabytes of geo-data. Users no longer need to know what machine to connect to in order to find data. They don't have to be informed when the physical data is moved or archived. Users simply access "the data" through the TerraShare namespace and the TerraShare files are always found in the same place, regardless of which machine the user is logged on to or where the physical files happen to be.

Likewise, IT personnel can do their job without worrying about how it will impact the rest of the enterprise users. Physical files can be relocated to different computers, archived to tape, restored back to online, etc., all behind the scenes without users needing to be informed or affected.

Another, more obvious benefit of creating the data abstraction layer, is that users are able to interact with logical files that fit with their business processes. Consider that a complete image data set may be composed of the following: one or more main raster files, one or more separate overview/reduced resolution data files, a world file, an associated coordinate system or design file, a separate header file, and a text or spreadsheet document containing metadata about the raster file. In a traditional setting, a user would need to be aware of all these files, and it would not be uncommon for some of them to be misplaced or lost as the "raster" file is copied, moved or archived. The data abstraction layer in TerraShare ties all of these files into a single logical TerraShare file. As a result, all users are able to access all of the data connected with the image without needing to know what or where the auxiliary information is stored. Additionally, TerraShare allows users to assign user defined metadata (which is stored in the TerraShare database) to each TerraShare file. This gives users nearly unlimited flexibility and extensibility when it comes to organizing and tracking data.

All these may seem like simple concepts and perhaps without much value. However the value comes in the form of eliminating the hidden costs of data disorganization. Consider a photogrammetry shop with six employees. If each of those employees wastes only one hour each week, either looking for the right data or data share, or copying a data set to a local drive, or even restoring data from an archive because he can't find it on disk or on a CD, than that shop is throwing away almost two man-months of work each year, simply because their data is not

organized or accessible. The amount of time wasted will only worsen as technology improves and the data sets get even larger.

2.1.2. Enterprise Data Access

As mentioned earlier, the primary client user interface for TerraShare is called the TerraShare Explorer Plug-in or Explorer Plug-in for short. As previously indicated, this user interface is implemented as a namespace extension within the Microsoft Windows Explorer. TerraShare folders and files are made to look and act (as much as possible) like standard Windows folders and files. This makes it easy for users to immediately interact with TerraShare because they are working in a familiar environment.

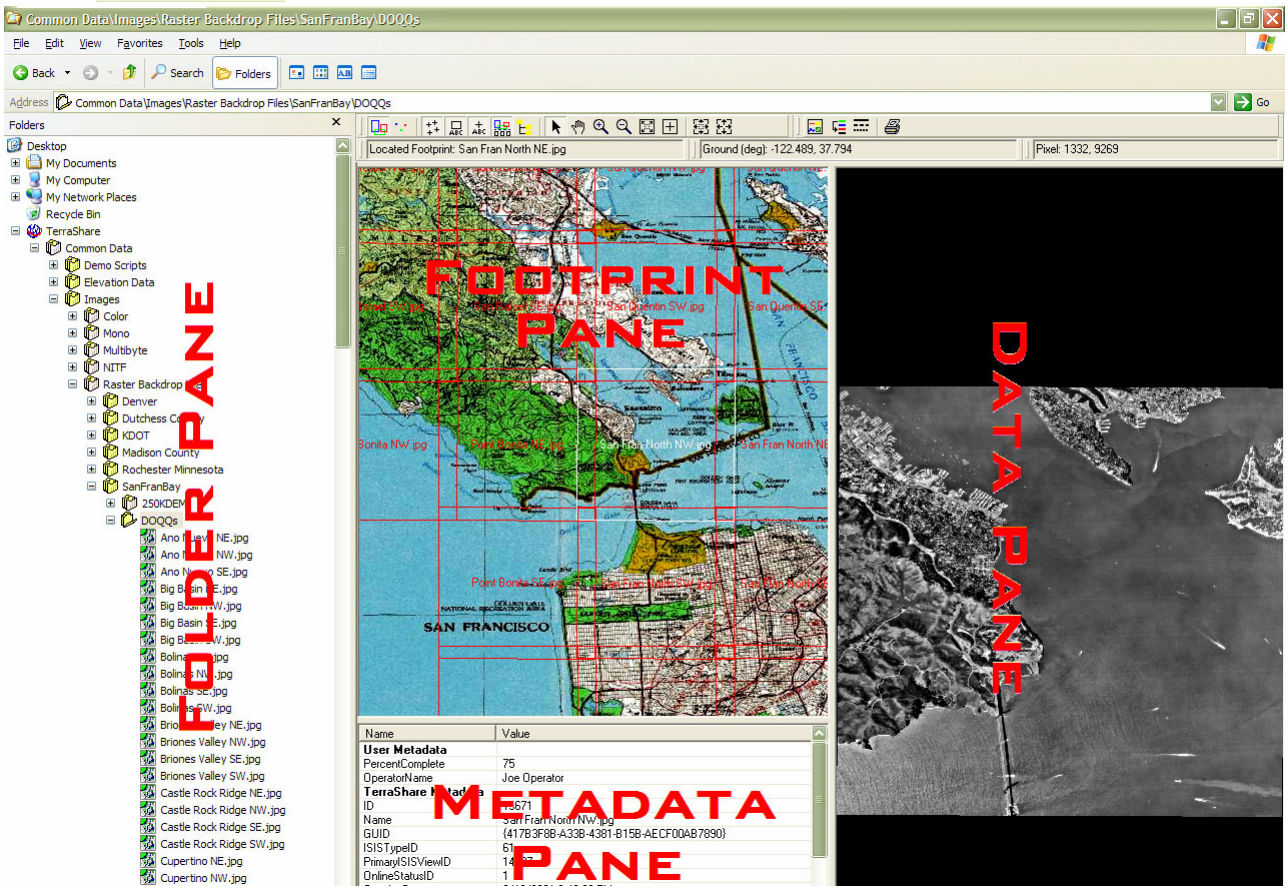

Figure 1: TerraShare Explorer Plug-in. The selected footprint, "San Fran North NW" is highlighted in white.

From an enterprise standpoint, however, the Explorer Plug-in offers many other added benefits. First, because the file system is virtual and not tied to any particular computer, the folder and file hierarchy appears the same to all users on all machines. Users don't have to remember specific computer or share names to access files. When new data is added to TerraShare, users don't need to be informed of its location. It simply appears in TerraShare and everyone has immediate access to it.

Another added benefit of the Explorer Plug-in is its native viewing capabilities. In addition to the standard Windows viewing modes, list, icons, details, etc., the Explorer Plug-in provides the user the ability to view files and projects in a geospatial footprint mode. When in footprint mode, the traditional two pane view (folders and files) is replaced with a four pane view, (see Figure 1). The leftmost pane, the folder pane, displays the list of folders just as in the standard two pane view. The

top middle pane, the footprint pane, displays footprints of all the georeferenced files in the selected folder or project. A raster backdrop may also be displayed behind the footprints in order to add context to the footprint data. The bottom middle pane, the metadata pane, and the rightmost pane, the data pane, show data associated with whichever footprint is selected in the footprint pane. The metadata pane shows all the metadata associated with the selected TerraShare file, and the data pane shows the primary view (in general, a raster file) of the selected TerraShare file. In this way, every user has instantaneous access to every file in TerraShare. No special viewers or applications are necessary.

Since the data and metadata are stored in a relational database,[1] any or all of the data can be queried and sorted. Operators, managers or even authenticated partners and customers can reliably search through the enterprise's entire data set of images, elevation files, vector files, etc., to find all of the files that meet their needs.

This type of functionality is difficult to quantify as far as measuring improvement in production efficiency. However, we contend that a simple to use, consistent, enterprise-wide file system makes it easier for operators to find and use the data they need each and every day. As we saw in the previous section, even a small increase in the day-to-day efficiency of an operator, results in a very large return in terms of recovered time. Also, the ability to access all "live" data and metadata from any desktop in the enterprise will prove to be indispensable to managers and supervisors that need to make decisions everyday. TerraShare enables a supervisor or manager to sit at his or her desk and immediately review the quality or quantity of the work being done over the entire enterprise. The manager's personal efficiency is improved because he or she can access the necessary information faster. The overall production efficiency is improved because managers have access to more information and that information is more up to date. This results in improved project management and quicker reaction time to production anomalies which might cause the schedule to slip.

TerraShare also provides the ability to make queries against the entire enterprise data set, which gives users a very powerful tool to quickly find the files they need. This can be especially valuable in time sensitive "crisis" situations. For example, it is not uncommon for requests to be made for all available image and/or elevation coverage over a certain area. Often, these requests come with other constraining factors such as, only data with one meter or better resolution, only color imagery, or perhaps only data collected or flown within a certain date range. Unless a site has its data organized within a relational database, these types of requests are either impossible to reliably fill, or can take hours or days to look up the necessary information. With a TerraShare query, site operators can find all of the relevant data within seconds, and can provide answers or data when turnaround time is critical.

2.2. Workflow/Production Management

So far, we have discussed data management and organization and shown that by implementing an advanced geo-data management system, users can improve production efficiency by eliminating data management waste and overhead and increasing data access capabilities. Now we will show similar production efficiency increases can be made by integrating a workflow and production management solution.[2]

[1] Microsoft SQL Server 7.0 and greater and Oracle 8i and greater are supported.
[2] Note that some of the features described in these sections are still in the development stage and are not available in the released version of TerraShare.

2.2.1. TerraShare Services

The TerraShare services layer enables application programmers to build on top of the TerraShare enterprise framework. By meeting a published interface new commands can be inserted directly into the TerraShare framework. By virtue of being TerraShare services, these commands inherit the enterprise features already built into TerraShare. Among these built in features are multi-user file protection, a raster data service and an elevation data service.

Multi-user file protection is necessary in any enterprise solution. Services can essentially check out a file for processing and TerraShare prevents other TerraShare users from changing that file until the service completes its own processing and checks the file back in.

The raster and elevation data services are utilities that other TerraShare services can use to request raster or elevation data. Requests for data can be file, folder or area based. These services are integrated with the TerraShare footprint viewer so that area requests are simple to implement. TerraShare extracts the data from all the TerraShare files that fall within the request criteria and then returns the data in the format specified by the caller.

All of these services make it much easier to convert existing single workstation based programs into enterprise applications while at the same time facilitating a tighter integration of all production processes and applications into the same environment.

2.2.2. Distributed and Queued Processing

Perhaps the most powerful feature in TerraShare, in terms of directly improving enterprise production throughput, is the Command Processing Queue, (CPQ). CPQ supplies the framework that makes it possible for applications to queue processes to run at a later date or time, run processes on other computers, and most importantly, distribute multiple processes to run in parallel on multiple computers throughout the enterprise.

The increase in production is fairly obvious in this case. Processing a project containing hundreds or even thousands of files might take a day or more to run serially on workstation based applications. TerraShare with CPQ automatically distributes processing tasks to as many computers as possible, and a job that might normally take a day or more to run can be completed in just a few hours.

2.2.3. Data Visualization

Production managers and supervisors need to make daily decisions concerning product planning, scheduling, tracking, etc. Having access to the data to base these decisions on is important. However, the ability to view that data in a meaningful way also has a large bearing on how that data is used or understood. TerraShare provides the tools so that each user within the enterprise can access and view the data in whatever way make sense.

The TerraShare Footprint View Manager gives users the ability to tie the graphic attributes in the footprint view to the metadata assigned to a TerraShare file. So, for instance, a footprint theme could be created that ties “Operator Name” to footprint color and “Percent Complete” to the footprint fill pattern. By using this footprint theme, a supervisor could quickly get a graphical view of which operators were assigned to which parts of a project and how far along each piece was.

Similarly, another theme could be created that tied hours (or dollars) spent on each particular project, so that management could quickly see where time and money were being allocated.

TerraShare also supports a pluggable viewer interface. This open architecture allows users to add their own data viewers directly into the data pane of the TerraShare footprint viewer. As better or more feature laden viewers for raster, elevation or vector data are created, they can be integrated directly into TerraShare. This empowers users over the entire enterprise to access and view any and all data in whatever methods best fits their needs.

3. GEOIMAGING INFORMATION: USE IT OR LOSE IT

In the previous sections it was shown how TerraShare can support the production workflow throughout all departments. In that, the crucial points were, how the IT-administration can effectively solve a difficult data management problem, how production can improve efficiency in both accessing and processing terabytes of data, and how management can improve control over the production workflow with better data management and visualization. In order to achieve that, it was shown that the data should be the focus, and that it must be accessible at all times – independent from its centralized or decentralized physical location. But is this only true for production?

After the production of derived data such as vector maps, ortho photos, or DTMs, there is the end user in the workflow chain, typically using a GIS. GIS applications have become more powerful over the last few years such that they not only support 2-D-vector data, but almost any kind of raster data as well. Additionally, many GIS applications now support 3-D coordinate spaces. The GIS user doesn't really care about how the data he is using was generated or stored, as long as it fulfills his needs. He wants to select data based on how accurate the data is, how up-to-date the data is, and of course how much the data costs.

Good information sources are aerial and satellite imagery, LiDAR, SAR, as well as the derived products previously noted. Z/I Imaging has addressed these needs using TerraShare for data management and creating plug-ins for image display in GIS applications. In other words, Z/I Imaging has broadened its portfolio as a photogrammetric vendor, by addressing the needs of customers in handling all sorts of raster data related to Geomatics applications. Because of that, a new word was needed to address all these different kind of raster data that are used in GIS, data that was not acquired by photogrammetric means. Hence the phrases "GeoImaging" and "EarthImaging" have been introduced to the market, with photogrammetry playing only a part of it.

3.1. Why is it important to have enterprise wide access to GeoImaging data?

An image itself has nearly no value, not even as an ortho photo. The same is true for a DTM. Without the ability to derive or extract information by means of (spatial) analysis using a GIS an image would have no fiscal value. Data only has real value, if it can be accessed by data consumers when they need it, regardless if the data is in its original form or if it is a derived product of some sort. So, how can an end user access a centralized data pool and why is it so important that this is done? Likewise, how can a user be made aware that the data even exists, especially if it is sitting in a vast enterprise wide data pool, or originating from other departments?

The reason this is important, of course, is to have better and more cost effective use of common data. The common data resources of any enterprise department should be available to all other

departments within the enterprise.

In previous sections we explained how user defined metadata can be stored with the data of interest as a single entity, a TerraShare file. This data can then be queried and searched for to continue with further processing. It was shown how this can be done in the Windows explorer environment. But is this sufficient for all queries? What if end users want to find their data by spatial filters, for instance, along a railway track? Tabular results and TerraShare file footprints might not be sufficient at all times for GIS specialists.

3.1.1. Query and Search from a Desktop GIS

TerraShare can easily be connected to a GIS. Since all metadata are stored in a database, they can be queried just as any other data. For instance, Intergraph has created a TerraShare data server module which allows a TerraShare database to act as a data warehouse connection to Intergraph's GIS, GeoMedia. Using this warehouse connection, GeoMedia can access the TerraShare database just as it would any of its native data warehouses. As such, users can not only query for TerraShare files directly in the Explorer Plug-In, but also search for data in vast data sets to find data related to their area of interest using all the power of a GIS.
Let's consider a situation, where an end-user wants to find all images and DTMs within the enterprise, which are in the vicinity of a highway. A buffer around the highway depicts the area of interest. This buffer can now be used as the basis for a spatial query such as "find all images that are within the buffer, were created after 1999, have a GSD of at least 30 cm, and are in color". Then, "find all the DTMs within the buffer that were created after 1999". The results can then be displayed either as tables for further selection or directly in the GeoMedia map window as vectors for direct interaction. It is then a simple matter to select an image footprint and turn its display on if needed. This is the point where additional value is put to the image! This is also the point where end-users can extract the information they need from the data they want. Today's GIS usually have easy to use GUIs where these queries can be placed in a simple matter so that, for example, the sales department without GIS specialists can query and search this data.

This works not only for individual images, but also for DTMs, raster maps or even entire photogrammetric projects. Future versions of the software will support searches for stereo models along with the ability to view those models directly within the GIS through a stereo viewer such as ISSV from Z/I Imaging.

3.1.2. Query and Search from a Web Based GIS

Is the enterprise the limit? What if there are only limited departments within the enterprise or the data is not of interest to most users within the enterprise? The World Wide Web makes it possible to share the enterprise's geospatial data to anyone with a web browser. We can integrate the data framework provided by TerraShare with a web based GIS, such as GeoMedia WebMap, just as we did with the desktop GIS solution described above. In this way we can present data and GIS searching and viewing capabilities to nearly anyone in the world.

4. CONCLUSIONS

We have shown that a data centric enterprise wide approach to data and workflow management can improve the efficiency of a production shop. A system such as TerraShare provides production sites with a suite of tools such that IT is able to manage terabytes of data, operators are able to

access and process that data efficiently, and management is able to plan and/or view up to date project data from anywhere in the enterprise. With these tools, production shops can both reduce the overhead associated with managing their data, and speed up their existing workflow processes. In addition, TerraShare provides the framework for transforming an existing workstation node based production workflow to a highly efficient, highly accessible enterprise system.

5. REFERENCES

Rosengarten, H. (2001): "TerraShare -- Automatic Management and Distribution of Geo Image Data Inside the Production Process", 21st Scientific and Technical Annual Meeting of the DGPF (Three-Countries Meeting), September 4-7, 2001 in Konstanz, Germany.

Bima, W.J. (2001): "Image Management and Distribution" – Proceedings 5th Aerial Cameras and Digital Photogrammetry Workshop, Steamboat Springs, CO, USA.

Graham, L. (2001): "Distributed Processing" – Proceedings 5th Aerial Cameras and Digital Photogrammetry Workshop, Steamboat Springs, CO, USA.

Greening, T. (2001): "TerraShare, A User's Perspective" -- Proceedings 5th Aerial Cameras and Digital Photogrammetry Workshop, Steamboat Springs, CO, USA.

Rogers, R. (2000): "Managing Imagery in the Enterprise" - GEOEurope, Ed. 9, Issue 11, pp. 28-29, U.K