Experience with the TerraShare Product in a Production Environment

TREVOR GREENING, Colorado Springs

ABSTRACT

Demand for digital satellite and aerial imagery has grown at an extraordinary rate over the past few years. The volume of image data on-line in at any time in a photogrammetric production facility may exceed several terrabytes (TB) with as much as 10TB or more of data not being unusual for a large organization. The rapid migration to Soft Copy photogrammetric production techniques has also increased the quantity of digital raster imagery. Moreover, e-commerce initiatives centering on the sale and/or hosting of digital imagery have further inflated the volume. The number of image and related metadata files may approach several tens of thousands, which when combined with the data volume poses a very significant challenge from a data management point-of-view. In the late 1990's, realizing the severity of the problem, Z/I Imaging commenced development of its TerraShare technology which is aimed directly at dealing with the management of large quantities of digital imagery and associated metadata. In the Fall of 2000, under a Performance Testing Agreement, TerraShare was deployed at the Colorado Springs office of Sanborn Colorado, LLC (formerly Analytical Surveys, Inc. – Colorado). The primary goal of the testing was to evaluate the performance of TerraShare in a high-volume Soft Copy production facility. This paper describes the operational environment into which TerraShare was deployed, continues with an assessment of the effectiveness of the software, and concludes with an overview of anticipated directions and goals for further exploitation of TerraShare at Sanborn Colorado, LLC.

1. INTRODUCTION

In recent years, many photogrammetric mapping companies have moved rapidly to adopt Soft Copy production techniques. At face value, the benefits of moving to an all digital environment seem to be obvious, however full exploitation of Soft Copy production methods can only be realized through careful planning and structured implementation. Hunt (2000) has identified many of the most important considerations, including the Information Technology (IT) infrastructure, photogrammetric production software, project data management, staff training, and technical support. In many cases, the IT infrastructure is extended in an *ad hoc* fashion, as a company migrates its operations to Soft Copy. Often, this results in an unplanned, inefficient, heterogeneous IT architecture. This outcome is illustrated in Figure 1, where the piecemeal growth of the infrastructure is evident as various disparate production groups have extended their capacity with little consideration given to the overall organization.

At Sanborn Colorado, LLC, the production infrastructure grew in just such a manner. The network topology tends to follow the geographic layout of the office space, rather than the data access, bandwidth, and functional requirements of each operation. Thus, several terabytes (TB) of on-line (disk) data storage are distributed in a haphazard fashion amongst numerous workstations. There are also many choke-points in the network, where switches are overloaded with network traffic.

The result of this *ad hoc* growth is an extraordinary waste of IT resources, in particular disk space and network bandwidth. Under this scenario photogrammetric technicians are compelled to copy image files and associated project metadata to local workstations, resulting in multiple copies of large image files, and duplicate metadata. The outcome is a considerable wastage of human resources and a loss of productivity, as well as significant problems associated with data management, in particular the maintenance of the integrity of project data. Additional problems include operational inconsistencies, such as the difficulty of imposing standardized operational procedures. Project management capabilities, such as production tracking, are also compromised under such operational circumstances. Finally, if collaborative production operations span an entire

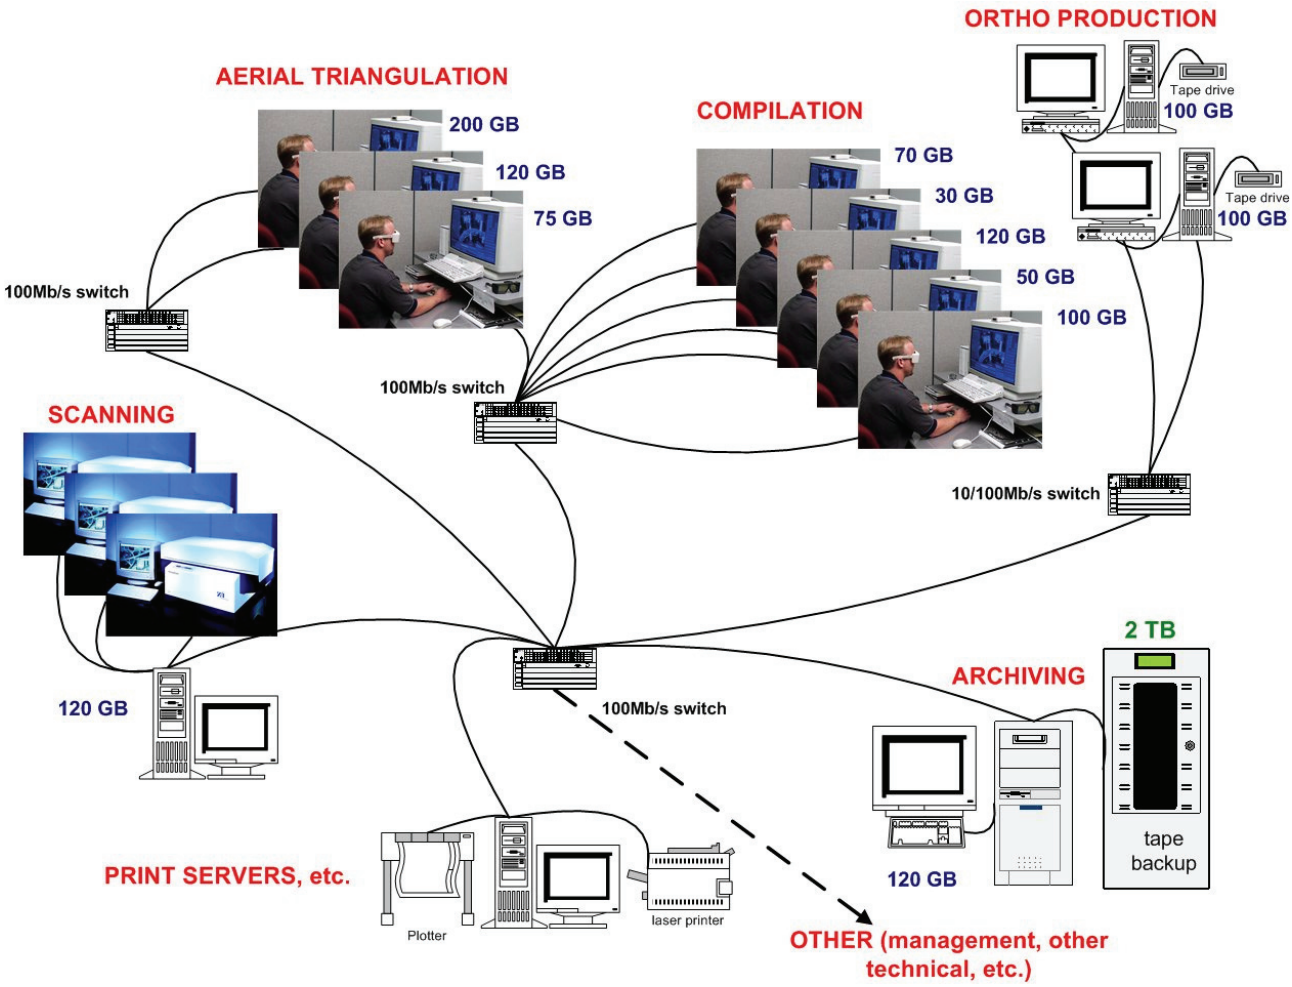

Figure 1: Piecemeal IT infrastructure growth (typical)

enterprise in several geographic locations, then the production workflow is often interrupted by the need to backup intermediate data and to transfer it to remote locations, usually via mail or courier service.

Some level of relief from these difficulties can be obtained by re-engineering the IT infrastructure. For example, it is possible to redesign the topology to better conform with the production organization, to upgrade to a high bandwidth (e.g., Gigabit) infrastructure, to add a storage system such as a Network Attached Storage (NAS) appliance, and/or a Storage Area Network (SAN) – see Bima (2001), and to add bandwidth to the Internet connection. However, many organizational and operational problems still remain, even after a significant investment has been made in IT enhancements.

At this point, it is worthwhile examining the complexity of a large-scale photogrammetric production environment, as exemplified by Sanborn Colorado, LLC's operation. The magnitude of the difficulties associated with managing large image datasets and thousands of data files can be inferred by examining Table 1.

Project	Number of images	Individual Image size (uncompressed tiled tiffs with overviews)	Total project size (uncompressed tiled tiffs with overviews)	Total project size (JPEG[1] compressed tiff with overviews)	Number of ortho-image sheets	Ortho-image sheet size	Total ortho-image size	Total ortho-image database (IDB[2]) size
-	-	[MB]	[GB]	[GB]	-	[MB]	[GB]	[GB]
1	1,084	112	119	30	160	48	8	15
2	1,790	375	656	179	2,600	56	142	200
3	636	375	233	64	n/a	n/a	n/a	n/a
4	96	470	44	8	100	72	7	10
5	400	375	146	32	n/a	n/a	n/a	n/a
6	1,110	380	412	100	1,360	25	33	50
7	1,724	375	631	103	n/a	n/a	n/a	n/a
8	328	375	120	26	364	75	27	40
9	780	380	289	48	560	100	55	90
10	800	375	293	48	515	75	38	75
11	5,180	375	1,897	528	3,075	75	225	320
12	1,334	310	404	100	2,192	75	161	250
13	937	375	343	75	463	25	11	20
14	240	375	88	19	1,701	6	10	20
15	163	375	60	13	80	24	2	5
16	379	375	139	30	316	24	7	15
Totals:	**16,981**		**5,874**	**1,403**	**13,486**		**726**	**1,110**

Storage Total = 8.9 TB

[1]Quality Factor = 13.
[2]Sanborn's proprietary multi-resolution image database.

Table 1: Subset of projects in production at Sanborn Colorado, LLC (July, 2001)

The 16 projects listed, involve approximately 16,500 high resolution scanned aerial photographs (mostly scanned at 14μm), and almost 13,500 orthophotos of varying resolutions. Intermediate products such as compressed image files and final deliverable products such as multi-resolution image databases raise the total number of images to more than 30,000. As indicated the total size of these products is around 9TB. Although not all of the images are on-line at any one time, it is clear that their management poses a very significant challenge.
It is here that TerraShare's capabilities come into play. Z/I Imaging's TerraShare technology goes a long way towards rationalizing and optimizing photogrammetric production by overcoming many of the difficulties mentioned above. In practice, however, TerraShare goes a great deal farther than the simple support of digital photogrammetric production. It provides an enterprise-wide environment for managing geo-imagery and derived products, including associated metadata. Furthermore, it opens up opportunities for e-commerce activities such as the warehousing of geo-imagery on behalf of clients who do not wish to manage and update their image databases.

2. TERRASHARE

A technical overview of TerraShare is provided by Wuescher (2000). As Wuescher points out, TerraShare has two primary functions. Firstly, rather than representing data in units that are logical from the point of view of the operating system (OS) – 'physical' files and folders (directories), TerraShare represents data in entities that have a logical meaning from the perspective of the user or application. For example, in TerraShare, a photogrammetric project is presented to the user in logical units: project, image, model, etc., rather than as a multitude of physical data files. Clearly there can be a many-to-one relationship between physical files and logical entities, such as in the case of a stereo model which is defined be two image files, and several metadata files containing

information about the interior orientations (IO), relative and absolute orientations (RO, AO), camera calibrations, etc.
Secondly, TerraShare manages large numbers of physical files on an enterprise-wide basis. In effect, the user is shielded from having to track the names and locations of the actual data files which can be distributed throughout the enterprise.
The logic of the first function, and the 'transparency' afforded by the second contribute very significantly to productivity, data integrity and product quality. Additional descriptive information regarding TerraShare may be found in Rogers (2000). The following section describes Sanborn's initial experience with TerraShare.

3. TERRASHARE DEPLOYMENT AT SANBORN – PRODUCTION EXPERIENCE

As part of a performance testing agreement between Z/I Imaging and Sanborn, TerraShare was deployed at Sanborn's Colorado Springs production facility. TerraShare is now the production environment for the majority of Sanborn Colorado's photogrammetric mapping operations, with the majority of the projects listed in Table 1 currently (as of July 2001) being managed via TerraShare. Figure 2, below, illustrates Sanborn's photogrammetric production workflow, and the present extent of TerraShare usage.
As indicated, orthophoto production and data delivery are not integrated into the TerraShare environment. Sanborn uses proprietary orthophoto production software ('METRO'), which has not yet been made 'TerraShare Aware', i.e., METRO cannot read and write TerraShare files. Although the figure indicates that the bundle block adjustments take place outside TerraShare, Sanborn has very recently commenced using the block adjustment utility imbedded in Z/I Imaging's ISAT

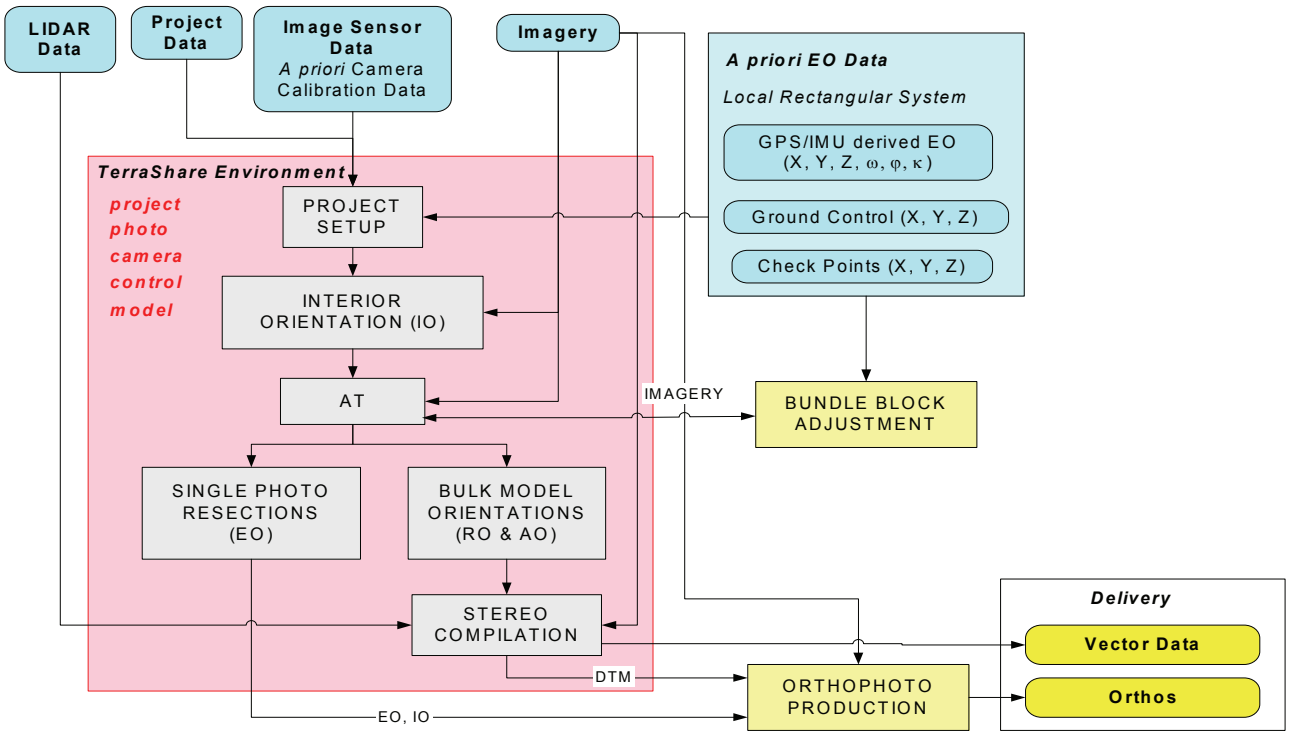

Figure 2: TerraShare usage at Sanborn Colorado, LLC (July, 2001)

product. Project setup, IO, AT mensuration, single photo resections for final EO, model setups (RO, AO), and stereo compilation are all performed seamlessly within TerraShare.
Despite the fact that the initial TerraShare deployment involved pre-release versions of the software, migration of projects to TerraShare proceeded very smoothly with little disruption of production, and with fairly minor technical problems. Some early resistance was encountered on the part of several Compilers, however, in retrospect this could be attributed to inadequate training and a consequent lack of understanding of the concepts.
Productivity gains were estimated by interviewing production managers. Fairly soon after the initial deployment managers reported estimated productivity gains of ~20% in aerial triangulation (AT), and ~15% in compilation. Other tangible benefits include improved quality control and a better working environment for compilers and AT technicians.

4. FUTURE DIRECTIONS

Based on experience with TerraShare, our vision for the short- and medium-term future of a production and distribution system might be described by a series of goals such as those prioritized below:

(a) Re-engineering of the IT infrastructure. This effort would involve a redesign of the network topology, bandwidth upgrades (e.g., Gigabit network components), the installation of a storage system such as a NAS and/or SAN, and the integration of a high capacity tape backup device and associated hierarchical storage management system (HSM).
(b) Migration of existing proprietary software to full 'TerraShare Awareness'. This would include orthophoto production software and ancillary tools for color balancing and auto-seaming. TerraShare's software development kit (SDK) can be used for this purpose.
(c) Delivery of geo-imagery products (specifically orthophotos) to clients via TerraShare and the Internet.
(d) Extension of the Z/I Soft Copy software via the development of proprietary tools to incorporate flight planning, semi-automatic project setup, and the seamless integration or ingestion of airborne GPS and INS derived EO data. For example, it is possible to set up the basics of a project automatically during flight planning. After photo acquisition, it is envisaged that project files could be automatically populated by extracting data from the scanned imagery (e.g., date, time, serial numbers, project name, exposure number, etc.) and by correlating this information with data from other sources such as a database of camera calibration values. This idea is illustrated in Figure 3.
(e) Exploration of collaborative production opportunities afforded by the Internet and TerraShare. Success in this area would significantly enhance communication with business partners.
(f) Hosting of data on behalf of clients. Under this scenario, possibilities exist for hosting and maintaining and/or updating data on a subscription basis.

5. CONCLUSIONS

TerraShare has already proven to be a valuable means of enhancing productivity, data management and quality control. The author believes that many of TerraShare's capabilities are yet to be exploited. Ultimate production goals center on a seamless production environment extending from project planning to data delivery. Productivity enhancements as high as 40% to 50% are not inconceivable. Significant business related opportunities are anticipated, as suggested in 4(e) and (f), above.

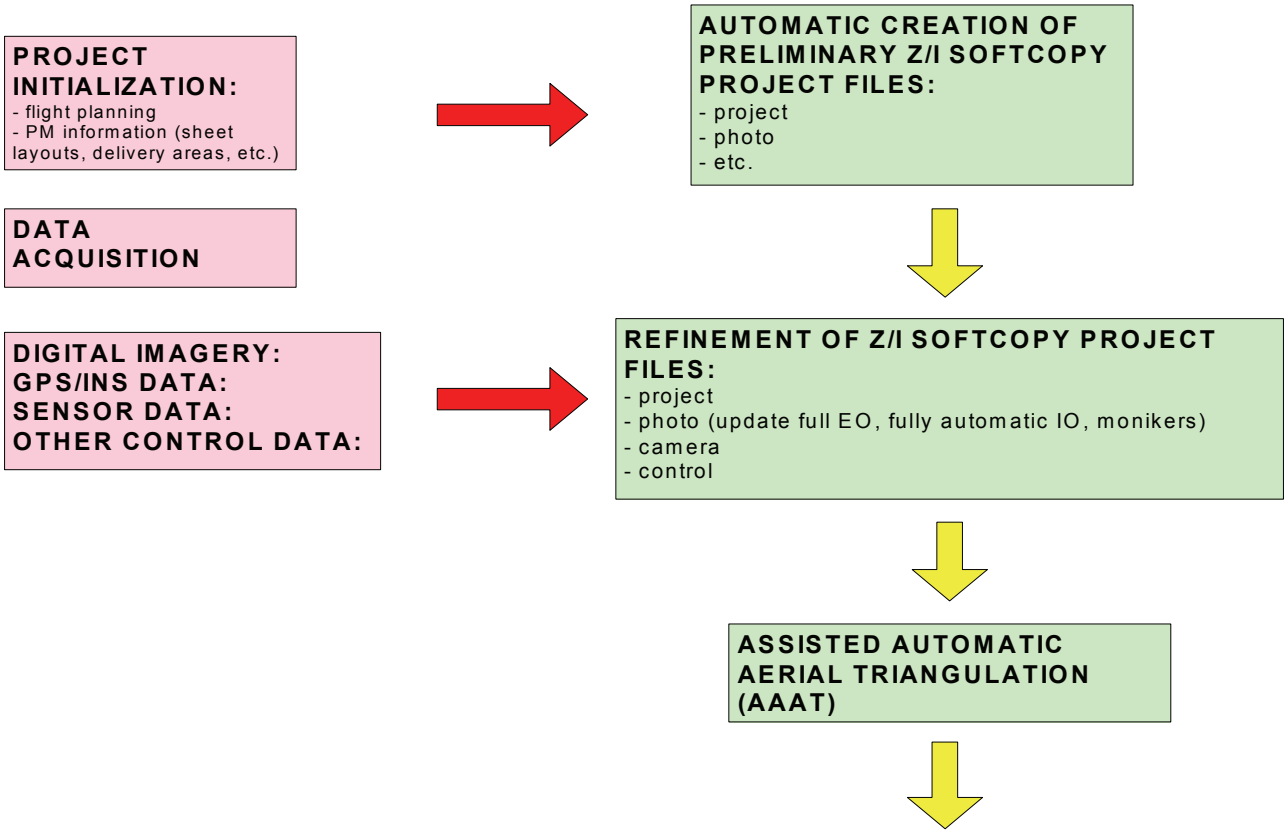

Figure 3: Illustration of the possibilities for a seamless, highly automated production workflow

6. ACKNOWLEDGEMENTS

The deployment and testing of TerraShare was a collaborative effort by Sanborn production personnel and Z/I Imaging's staff. In particular, gratitude is owed to Joe Bima, TerraShare's Product Manager for his thoroughness and dedication in migrating Sanborn Colorado's Soft Copy production to TerraShare, with minimal interference to day-to-day operations. On behalf of Sanborn, the author would like to thank Z/I Imaging for considering Sanborn's Colorado Springs office as a performance testing site. Finally, Sanborn's production staff are owed a vote of thanks for their patience during the deployment and initial testing phase.

7. REFERENCES

Bima, W.J. (2001): Data Management Solutions. Geoworld, Vol. 14, No. 7, July, pp. 38-41.

Hunt, R. (2000): Soft Copy All the Way. Unpublished Presentation. Z/I Imaging Annual Users' Group Meeting, Vail, Colorado.

Rogers, R., (2000): Exploding Imagery Market Places New Demands on Image Management Systems. Archives of Earth Observation Magazine, Nov. (available online at http://www.eomonline.com/).

Wuescher, D. (2000): A Technical Introduction to TerraShare. Internal publication: Z/I Imaging, version 01.03, June.