

File Economy: When Less is More

An objective white paper on LizardTech's DjVu

By Gordon E.J. Hoke and Bernard M. Chester, IMERGE Consulting,
www.imergeconsult.com

Copyright 2001, IMERGE Consulting, all rights reserved

The 2002 world of technology is stoked with contrasts. On one hand, leading operating system vendors offer updated versions of their software that are bloated with inelegant code, relying on hardware developers to expand memory and storage capabilities to compensate for their cumbersome constructions. On the other hand, the Internet is choked with huge files, and a rapidly expanding array of remote, often mobile, devices scream for smaller files that fit their appealing, diminutive size.

Into this world of too much data comes DjVu (pronounced "DAY-zha voo"), software that dramatically shrinks file sizes with minimal compromises, while adding multiple features for web-based publishing. Reductions to as little as two-tenths of one percent of the original are possible.

While few IT workers have heard of it, DjVu is new only in its current, most sophisticated version. AT&T Labs developed it several years ago, and it has surfaced occasionally for specific operating systems or applications.

In March, 2000, however, LizardTech [www.lizardtech.com], which started at Los Alamos National Laboratory but now calls Seattle home, turned the technology into a product. LizardTech makes DjVu widely available in a format that could become the industry standard. DjVu avoids many of the limitations of Adobe's PDF, currently the dominant document format in the marketplace.

"DjVu is specifically built for Web-based publishing," states John Grizz Deal, LizardTech founder. "We're not competing with PDF, we are taking on the *mis-use* of PDF."

LizardTech has gone to significant effort to make DjVu available and reliable. A DjVu viewer is available as a free download on the Web. Free modules integrate DjVu with Microsoft Office and Kofax Ascent Capture. LizardTech offers an "open source" reference library that lets corporate and government customers feel secure about the future of DjVu. A solo version of the encoder (currently \$299 for commercial use; free for non-commercial use), a workgroup version (\$5,000 per CPU) and an enterprise version (\$7,000 per CPU) are now shipping.

All of this means that anyone who struggles because a) their files are too bulky to fit their storage, or b) transmission times are unacceptably long, or c) they require a single source document from web to color print-out will find a worthy, economical solution in LizardTech's DjVu.

What Does It Do?

While DjVu is a boon to organizations overwhelmed with storage demands, its greatest strength may be its scan-to-Web capabilities. DjVu takes large units of unstructured data – including books, magazines, newspapers, and catalogs – and enables them to be posted on a Web site with surprising speed. Readers download these files with identical alacrity. One example is a 110-page, full color corporate report that consumes 2.5 gigabytes in TIFF but reduces to less than three megabytes in DjVu (the same file in PDF is over 150MB!). Readers need that free (600KB) DjVu viewer available on the Web. Optional, bundled Optical Character Recognition (OCR) technology from Expervision permits full-text retrieval.

DjVu separates a scanned document into two layers, sending them sequentially. The first contains text and line drawings, and it arrives with lightning speed. The second holds graphics, photos and color backgrounds. When users download files, they get the text almost immediately. Graphic information appears as pixels arrive.

DjVu is notably different from standard document formats. It doesn't drop text and line information to shrink the file size, hence that part is dubbed "lossless". While backgrounds and graphics do sustain information loss, DjVu's advanced algorithms and wavelet techniques retain noticeably better images than competing strategies. A DjVu user who receives an image can zoom, pan and print documents of nearly the same clarity as the original. This provides obvious advantages for maps, diagrams, photographs, and artwork.

Who is LizardTech?

LizardTech, founded in 1992, is a privately-held company dedicated to making research technologies available for mass use. Investors include Oak Ventures, Mitsubishi, Maveron Ventures, Encompass and Seapoint Ventures, as well as smaller firms and individuals. "Our company was founded to commercialize technologies related to raster-based imaging," explains Deal. "Over the last nine years, we have assembled a comprehensive team of mathematicians, engineers, computer scientists, programmers, and marketers from companies like Los Alamos, Xerox PARC and Xerox, Adobe, Microsoft, Apple, Minolta, Intel, AT&T Labs, and Altamira."

Deal reports the company has leveraged over \$50 million in laboratory research. New investments are designed to market products to show a profit by 2002. Deal notes the company has always been revenue driven and has always had an income stream. Deal predicts sales for the current fiscal year will be near \$15 million.

In addition to DjVu, LizardTech develops and distributes MrSID, technology originally conceived at Los Alamos. Both products "optimize bandwidth and provide immediate access to high-resolution, multipurpose digital images and multimedia content." LizardTech products integrate with over 200 applications, addressing a potential market in the tens of billions.

LizardTech's corporate strategy has two fronts. One is to acquire groundbreaking, complementary technology and bring it to market. For example, the company recently acquired Altamira's Genuine Fractals technology. These products expand the market opportunities for LizardTech dramatically.

At the same time, Altamira's success in OEM bundling will help LizardTech reach its second goal. Although DjVu and MrSID are standalone products, their success is directly related to their integration with other products. Toward this end, LizardTech is working with Mitsubishi in Japan to bundle DjVu and MrSID into computer peripherals and other digital devices.

"For example," says Barry Weiss, Senior Director of Business Development, "Palmtop devices and cell phones have less capacity than computers and seriously need DjVu for viewing. If I am an insurance worker in the field, and I need to see an insurance form and send it to my farmer client, it takes a long time to send via current Palm technology. In DjVu, however, it goes really quickly."

"Our customers take our technology and put it into their solutions. We help their companies compete better." Deal states.

"We have a history of working with partners, and now we are of a size where we can support our partners with great technology, market support, vertical support, etc. We support our customers, and we support our customers' customers," he adds.

Since late 2000, LizardTech has balanced its drive for technology with a marketing initiative that bodes well for near term success. Its omnipresence at the 2001 AIIM trade show was a watershed event in terms of name recognition.

IMERGE Consulting considers LizardTech a stable, innovative company with the financial, technological, and human resources and strategic marketing plan to be successful in the foreseeable future.

The Technology to Meet the Needs: Modern Compression

Modern compression technology solves problems. As research has 1) improved telecommunication speeds and 2) permitted quality video and audio over the Internet, now it is finding ways to encode documents more effectively by reducing the size of a document image without losing any of the content or usability.

Compound document encoding is not the same as file compression. File compression looks at the patterns in the data and attempts to reduce the space necessary to hold it, but leaves the user with the task of re-expanding it with every use. That creates added overhead on both the initial storage and each retrieval of the document. File compression, by its approach, has limits on how small the result can be. It expects that the numerical patterns within the data are going to be repeated multiple times. Document compression extracts the structure of the document image and encodes it in an efficient way. By understanding the nature of the document being imaged, intelligent

techniques can be used to encode the image. The image is not handled as a set of pixels, but instead is dissected into visual and conceptual aspects, each of which is encoded using the optimal technique for that class of information.

Modern document compressions use very sophisticated mathematical algorithms to divide the image into pictorial, color, text, and straight edge factors. Each factor is designed to compensate for inefficiencies of the others.

- Picture compression techniques work by encoding tiles of a picture using mathematical approximations, but, as a result of this approximation, lose acuity and blur the sharp edges of small text and forms.
- Solid color areas are poor candidates for picture compression approaches, since they will produce larger and less visually clean results. However, they can be easily represented as geometric areas and a color definition, and permit the picture compression to be even smaller.
- Raster and vector formats are excellent for producing a clean image of text and forms, but require larger data streams than picture techniques. They are very inefficient for representing photographs. By using them sparingly, the size of the resulting file can be significantly smaller while maintaining the edges needed to make the document readable after compression.
- Indexing tools, such as Microsoft Index Server or Verity, provide the ability to do content-based retrievals from a document library and use the text contained in a document. They do not understand images, so the text is extracted using OCR and efficiently stored as a character stream.

A corollary to this is that the more complex the documents, the greater the benefits from these “divide and conquer” compression technologies. The magic is in taking a document image and extracting these characteristics.

Needs of the Marketplace

There several current “myths” that would make one doubt that image and document compression technology are still needed. Nonetheless the problems of today’s Web content managers and operations managers cry out for these technologies to help them address some of their largest (no pun intended) problems.

Myth Number 1: Disk storage is cheap

There is no denying that the cost to store a million characters has dropped significantly and continues to drop. However, this has only encouraged computer users’ worst pack-rat habits. IT managers have generally abandoned earlier attempts to be “corporate mommies” and get users to manage their disk space. It is analogous to getting teenagers to clean their rooms.

Resolving storage quandaries by “throwing more memory” at them is too close to an industry standard. Lobbying for, or even demanding, conservative use of disk space makes IT managers targets for undeserved frustration and aggression. Industry and regulatory demands are requiring that electronic information be maintained for longer periods of time. The Federal Electronic Signature Act of 2000 mandated that users

maintain copies of all electronic agreements, and recent judicial decisions have encouraged maintaining copies of web sites for defense against lawsuits. The use of data mining tools requires quantities of historic data to support them. Passage of the Federal Electronic Signature Act made many organizations realize that they need to archive web transactions and the state of their web sites at the time the transactions occurred.

The unfortunate IT manager is therefore constantly faced with systems that are running out of storage, and the challenge of adding more capacity to existing systems -- be they magnetic hard drives or optical jukeboxes -- is daunting. Disk space may be cheap, but it isn't free, and once installed it must be backed up and maintained.

Myth Number 2: Everyone has high speed Internet connections

The majority of Internet users are still using dial-up modems, so their connections are still running about 28K bits per second, or, if they are lucky and are close to the telephone switch, 50 Kbps. The recent failures of DSL ventures clearly demonstrate that not everyone is rushing to buy a fast connection to the Internet.

Even intranet users don't always have a fast connection. The corporate LAN may be high speed, but it is increasingly burdened with traffic. As a result, the effective speed for sending a large file may be no better than a dial-up line. Telecommuters and mobile users find themselves limited by modem speeds.

Finally, wireless devices users can't ignore file size. Any technique that reduces the size of any document that might be sent is deliciously attractive when transmission speeds are reduced. Hence, trying to deliver document images, particularly color documents, while using web technology still presents a challenge. Sending the images as TIFF or PDF will tie up browsers for a significant time. Web designers still need to create applications with consideration for the "slowest common denominator".

Myth Number 3: Document compression is only useful for the Internet

Offline information distribution still dominates, and compression is even more useful there. For example, CD-ROM is a convenient way to distribute or store data today. The disks can be read on any machine, and the cost of drives to write them is insignificant. However, whether distributing or archiving documents on CD-ROMs today, their 650 Megabytes can fill up quickly when storing a large number of documents. A document compression format can squeeze more onto each disk.

Even with electronic messages, file attachments swell the size of storage requirements. Imagine how much smaller mailboxes would be if there was a way to reduce the size of attachments; imagine the cost-of-storage savings and the pleasure of IT managers!

The Products that Apply the Technology

Several new document compression technologies currently in vogue -- PDF, TIFF-FX, JPEG2000 -- but DjVu has some real advantages over these and others.

PDF is considered by many to be the de facto standard for storing documents. Why does DjVu excel in many areas?

The Technology: Superiority of DjVu Document Compression

A key factor in evaluating document formats and compression technologies is the quality of the result. After all, one may choose from many alternatives. Three factors must be considered:

- The quality of the resulting image
Document compression is not worth using if the resulting document is not readable. DjVu uses “lossy” techniques on the picture portions of the document where the loss is not visible on a computer screen, and “lossless” techniques to keep the sharp edged information, like text, clear and legible. It intelligently meshes the best of raster approaches like TIFF with the compression of yet-to-be-released wavelet schemes like JPEG 2000.
- The utility of the new format
DjVu files provide some capabilities unavailable in other formats. By separating the picture from the text, DjVu improves the user’s ability to read the text and do manual indexing. The optional OCR capability used with standard indexing engines -- such as Microsoft Index Server, Verity K2 and Fulcrum SearchServer -- permits the user to index DjVu files by content.

Records managers and corporate counsels have clear concerns about the immutability of electronic documents. Unlike other formats such as TIFF and PDF, DjVu files cannot be altered. However DjVu does support hyperlinks and annotations, so the format can be used in document review and dynamic publishing processes.

- The size reduction
The primary reason for using a document compression technique is reduce the size of a document file. This makes everything from storage to transmission easier to use and much more efficient, especially with the advent of wireless communications. DjVu does an excellent job of cutting files down. Here are two examples of how well DjVu works in the real world:

Consultant Bernard Chester took a 7-page CCITT Group 4 Facsimile TIFF file of a contract he had signed as a test. The TIFF file was scanned in at 300 dpi by 300 dpi in black and white, and took 330,735 bytes. Using DjVu Solo 3.1, he converted it into a bundled DjVu file of only 39,536 bytes, which is only 12% of the original, while maintaining the quality of the original. For contrast, ZIP compression of the TIFF was 95% of the original, and Acrobat Distiller made it larger, at 377,072 bytes!

Chester then took his favorite photograph, which is 5.3 x 8 inches, and made a 96 pixel/inch true-color JPEG from it. The JPEG file was 116,880 bytes. Using DjVu Solo again, he converted it into a DjVu of 35,461 bytes, which is 30% of the original. ZIP could not provide any reduction

in the storage size of the photo, and a PDF of the document was 56,536 bytes.

While both of these cases show excellent real world results, other situations and applications demonstrate even more dramatic file size reduction ratios.

Open Technology

Because of LizardTech's key role in prompting and advancing the DjVu technology, many have the mistaken impression that DjVu is a proprietary technology. Quite the contrary! DjVu is less proprietary than the Portable Document Format (PDF) of Adobe Acrobat, and certainly less proprietary than the Microsoft Windows™ or Apple Macintosh™ OS code commonly run on desktops.

In fact, LizardTech has published the DjVu format, to make it available to the public. Unlike Adobe, LizardTech makes its viewer developer's kit freely available, and provides a reference library under an open source agreement. Again, in contrast to Adobe, LizardTech has provided sample encoders in the developer's kit and not just viewing hooks. As a result, investments in DjVu are not dependent on LizardTech's ability to produce tools, and users are assured that they will always be able to decode their files.

LizardTech does market a number of quality tools for generating DjVu, and it actively supports improvements to the technology. The encoding tools apply LizardTech's proprietary techniques for analyzing a document image and generating the highest quality files. These tools are designed and priced to address different levels of need – casual interactive; small production runs, and large production batches. In addition, LizardTech has a number of utilities to help with conversion from other compression formats.

Viewers for all Common Platforms

Obviously, a technology with broad application should be available to all potential users. DjVu viewers are free and available for all the common platforms, such as Windows 32 and Macintosh, and they are localized into foreign languages. Viewers with support for additional platforms are available from third parties. LizardTech releases viewers for additional platforms regularly.

Ease of Use

In addition to the production utilities LizardTech markets for producing DjVu from imaging files, the company recently released two tools for desktop creation of DjVu files. The first is an MS Windows virtual printer which permits any desktop product (such as MS Office, web browsers, and graphics tools) to generate DjVu directly. If the program can print, it can publish directly to DjVu. The second tool permits Adobe Acrobat files to be converted to DjVu, simplifying the conversion of legacy documents.

Effective Compression that Integrates with Major EDM Systems

While a number of industry vendors are discussing integrating DjVu as an option in their products, introducing DjVu into existing document and image management systems is not difficult. Indeed, in Chester's opinion, it is simple to apply in several different ways. By using the integration interfaces of products like FileNET Panagon, Documentum, and Hyland OnBase, one can introduce document compression into operations within a few hours. Chester's suggestions:

Post Process Scans

One easy way to convert imaging applications to take advantage of DjVu is by modifying the scan and index application. After the indexing process is complete, convert a TIFF file into DjVu and then commit the DjVu file into an imaging system instead of the TIFF. Most imaging and all EDM systems will permit this, and the user still has the document indexed under the same scheme as older documents. LizardTech has a sample script for Kofax Ascent that illustrates how to do this.

Users do not need to change the scan and index work process, and yet they have significantly enhanced their system. Those with an IDM system, in addition to the significant reduction in space required, can use the DjVu OCR feature to permit content indexing of imaged documents, a feature that TIFF does not permit.

Now, anyone accessing a new document -- particularly through a performance limited interface such as a web portal -- will experience faster response without any loss in quality.

Pre-Internet Transform

Those who choose not to convert document stores into DjVu can still reduce the problems they have trying to share TIFFs and PDFs over the Internet by adding some simple logic to their portals. When the browser asks for a document, users can fetch it back to the web server, convert it to DjVu on the fly and then deliver the DjVu. The responsiveness of a web site will improve greatly, without needing to make changes inside the repository system.

Business Realities: *Here Today*

Approaches like TIFF-FX and JPEG2000 are much heralded but not widely available at this writing, especially when considering that LizardTech users number almost half of the entire installed base of QuickTime. DjVu technology is readily available today and field-tested. The technology is now in its third generation, and continues to advance, thanks to the strong backing of LizardTech and AT&T.

Are There Places When DjVu Isn't The Answer?

There are some situations where DjVu is not the correct technology.

1. When any loss of clarity in an image is unacceptable, DjVu (as well as PDF and the various JPEG formats) is not an appropriate technology. For example, diagnostic-level X-Ray images should not be saved in DjVu.
2. DjVu is designed to focus on compressing documents with mixed text and pictures. In these cases it delivers the best compression ratios – upwards of 400 to 1. With pure pictures or pure text its gains are less impressive, although quite real.
3. In addition, as with Acrobat, LizardTech's DjVu encoders give their best results with some guidance as to the style of document. Unless original documents can be categorized, it can be difficult to know what setting to use for the best results. In a production environment with mixed input, this can be impossible to determine, and users may have to settle for a less than perfect conversion. One wishes that the tools would use their knowledge of the document to select the optimal settings.

DjVu in Action

When DjVu is used appropriately, it makes a dramatic difference. Here are some examples:

1) In Cobb County, Georgia, Clerk of Superior Court Jay C. Stephenson struggled with real estate records. At least 20 workers put in 66-hour weeks without closing the 46-day gap between document arrival and document posting. Too little space and the weight of too much paper forced the county to plan a \$10 million storage facility. Retrieval of documents was a length and costly process.

A Florida integrator solved the problem with a scanning system that turned all the paper into electronic files. That solved the paper overload and the labor nightmare. Retrieval was another matter. Retrieving multi-page documents with drawings and photographs can be daunting; the thought of providing these over the Internet seemed overly ambitious.

For retrieval, the integrator chose DjVu. Computers at the courthouse contain the DjVu viewer. However, anyone who wants to see these public records can download the viewer for free and get access to real estate records over the Internet. Delivery is measured in seconds.

At the courthouse, elapsed time from arrival time to posting is a mere three hours. Cost savings are significant. Stephenson says the best part is the easy, cost-free, timely access for whomever needs to see the records.

Uses for DjVu in criminal justice and legal systems are boundless. While some jurisdictions currently scan records – police’ incident and accident reports, for example – the speed of distributing these to the public, attorneys and insurance companies is an issue that DjVu can address.

2) Investors often want to see the annual reports of companies in which they consider investing. Since corporations invariably want to look successful, annual reports are often lengthy tomes, replete with photographs, charts and diagrams. Printing and distributing these works are costly, time-consuming operations. Imagine the savings in time and money (to say nothing of trees) if annual reports were available in an easily downloadable format.

That is the *raison d’être* for Corporate Reports (www.corpreports.co.uk), a British firm that is a clearinghouse for annual reports. Thousands of companies from many different stock exchanges submit their annual reports to Corporate Reports. While the highlights are available for free, any Web user can pay a small fee for unlimited rights to download reports. They all download in DjVu.

Since the files are less than one-fifth the size they would be in PDF, the downloading process is rapid. The availability of the information is even better because the text downloads almost immediately while the graphics follow. A subscriber begins reading his or her real interest, the financial report, while the product pictures follow.

3) Insurance companies use such a mix of data, information and images that DjVu holds the potential to ease multiple problems. Most insurance companies are diffuse, that is, they rely on field offices or agents to reach the consumer. Forms including diagrams, signatures, and drawings travel over significant distances. Many insurance companies equip their field staff with laptop or handheld devices, but transmission times can make these expensive devices ineffective. Transferring files at DjVu speed maximizes their use, benefiting the corporation and the consumer alike.

4) CampusView delivers virtual tours of college campuses for prospective students. “Paper viewbooks contain irreplaceable photography,” notes a founder, “but most bulk mail goes straight to the trash without even being opened. We find this offensive and unnecessary.”

CampusView’s initial efforts compiled viewbooks on CD-ROM. After forging an agreement with LizardTech, however, downloads became practical, and the company added CampusView.com. Its online library lets prospective students skim through the available reports, downloading those of most interest. Since few of CampusView’s customers have high speed Internet connections, this service would have been impractical without DjVu.

5) DjVu loads images to Web sites as well as it downloads them. This function can be critical to a retail company seeking to put a catalog on its Website. The Sharper Image is

one of the best known firms using DjVu for Web postings. Since catalog contents change regularly, the benefits of the technology are ongoing.

6) E-books, that is, publishing books electronically instead of on paper is, as of this writing, a concept that is discussed more than it is used. Standards are not yet established and reading hardware is not yet refined. Individual books are available at Web sites however, and some can be downloaded using DjVu. The techniques are still evolving.

7) While uses for DjVu are limited mostly by human imagination, several are already being rolled out.

- Online auction bidders can get larger, more detailed photos of sale items.
- Similarly, online shoppers need no longer be content with thumbnail photos of products. DjVu's zoom function reveals details not previously accessible.
- It becomes practical to offer and present collections of graphic arts and photographs on the Web. Museums, galleries, and government institutions can make their collections available wherever the Web reaches.
- Higher education uses DjVu to deliver a variety of resources. Libraries deliver journal articles, theses, research, and articles wherever and whenever they are needed. DjVu makes historical manuscripts, genealogical records, ancient maps, and registries available anywhere.
- Archives of periodicals are available for rapid download.
- Maps, financial records, and engineering documents and drawings gain helpful mobility and portability.

The Analysts' Perspective

LizardTech, born in 1992, has been growing and exploring for years, but it finally seems to have moved beyond an inconsistent adolescence to growing maturity in an energetic, focused young adulthood. The company and its products have seemingly limitless opportunities for growth. At this writing, leadership under John Grizz Deal is intelligently choosing amongst the myriad alternatives, finding the directions and markets that lead to rapid growth and profitability.

It is the analysts' opinion that the need for file/document/content compression is unlikely to diminish. Even though bandwidth is likely to increase and storage media are likely to get more dense, the flow and storage of data is expected to increase at least as fast. There will always be a premium on storing more information with fewer resources. In addition, the continuing growth of Internet distribution of information will make DjVu's special advantages for web delivery more attractive. Hence IMERGE Consulting expects LizardTech's products, including DjVu, to grow in importance and value to the worldwide marketplace.

The technology is proven and powerful, and it is apparently without equal. However, the race does not always go to the swift, and a better mousetrap is only as profitable as its marketing is effective. LizardTech needs integration with software and hardware, partnerships with companies and integrators, and effective, long-term marketing to live

up to its potential. It also needs to offer rapid, consistent and reliable technical support and customer service.

LizardTech has more-than-adequate resources to fill these needs. Its customer-base is steadily growing. Its leadership is intelligent, experienced and focused. Although as a privately-held company its financials are not available, there are reliable indicators that the company is economically stable.

The authors recommend LizardTech to organizations that need document compression. We expect DjVu to significantly challenge to become the next industry standard.

The components are all in place for LizardTech's success with DjVu. The need for compressed file size expands with the burgeoning number of small-size devices. Use of the Internet grows exponentially as the number of users skyrockets and the use per capita also expands. As a file transfer system from desktop to Web and back again, DjVu excels. The demand for the product is huge and growing.

At the same time, LizardTech seems ready to capitalize on that demand. They have assembled a blue-ribbon team to keep DjVu on the cutting edge. LizardTech's other major product, MrSID, is a fitting complement. Part of LizardTech's thrust is to forge partnerships with hardware and software manufacturers. If successful, this will lead to the inclusion of DjVu's viewer on multiple mobile devices and appliances. It can also be bundled with software to speed file transfer, whether the files are scanned images or ebooks.

While the high tech world offers no guarantees, the integration of excellent technology, dedication leadership, adequate financing, and a sensible business plan implies success. LizardTech and DjVu are well poised for the future.

■ 30

Gordon E.J. Hoke, principal writer and editor for this paper, is an analyst and consultant with over 150 publishing credits. He is a Contributing Editor to *Transform Magazine*. He can be reached at ghoke@mindspring.com and (507) 534-2293.

Bernard Chester, technical writer for this paper, is a Principal with IMERGE Consulting specializing in the integration of various technologies with document and image systems. He has been the technical editor of several document industry standards, and writes and speaks to the industry. He may be reached at bchester@imergeconsult.com or 206-979-7389.