DOI: 10.2352/issn.2168-3204.2016.1.0.28

DoD Visual Information Storage Challenges and Lessons Learned

Paul G. Robinson, Defense Visual Information (DVI), Defense Media Activity (DMA), Fort Meade, MD

Abstract

Defense Media Activity (DMA) is the Department of Defense's (DoD) direct line of communication for news and information to U.S. forces worldwide. The agency informs DoD audiences, entertains DoD audience overseas, trains Public Affairs and Visual Information professionals, and manages the DoD's visual information. Defense Visual Information (DVI) manages DoD visual information in support of U.S. military activities and operations and conducts visual information planning, policy, procedures, guidance, management, and standards. Defense Imagery Management Operations Center (DIMOC) provides DoD enterprise-level visual information services including operational support, digitization, storage, access, records management of the Department of Defense's visual content, and accessions to the U.S. National Archives and Records Administration.

DVI has a need, like all archives, for digital storage. Hence our study of the storage industry and the solutions they develop, which vary significantly. The need for more storage conflicts with our constant budget pressure, necessitating the right balance of performance versus cost. Solutions must be reliable and fast enough to support the mission while not breaking the bank.

Research into the various solutions available began in 2006 when a new effort was started by DVI to fix the digital asset management problem.

This paper will discuss the history of storage systems used by DVI, why the first system could not be used, research with various vendors and the requirements presented to the vendors. The paper will also discuss the selection and use of very high capacity storage designed to hold video master files. The paper will also explain how we moved the storage system in 2011 to the new agency.

The paper continues by discussing the growth of the storage system in DVI, moving from vendor to vendor. The paper concludes with a discussion about current cloud storage options and what that means for DVI.

Defense Visual Information

Defense Visual Information (DVI) manages DoD visual information in support of U.S. military activities and operations and conducts visual information planning, policy, procedures, guidance, management, and standards. Defense Imagery Management Operations Center (DIMOC) provides DoD enterprise-level visual information services including operational support, digitization, storage, access, records management of the Department of Defense's visual content, and accessions to the U.S. National Archives and Records Administration.

DVI has a need, like all archives, for digital storage. Hence our study of the storage industry and the solutions they develop, which vary significantly. The need for more storage conflicts with our constant budget pressure, necessitating the right balance of performance versus cost. Our solutions must be reliable and fast enough to support the mission while not breaking the bank.

The Beginning

Research into the various solutions available began in 2006 when a new effort was started by DVI to fix the digital asset management problem. In 2000 work began to build a new system, the Visual Information Management System (VIMS). The first system fielded by DVI was discontinued in 2000 by the manufacturer with limited support afterward. However, the VIMS project proved to be too ambitious resulting in poor progress over several years. A fresh start was made in 2006 with an eye toward using as much of the hardware and software purchased for the previous effort as possible. The new system was labeled the Defense Asset Management System (DAMS) and went live less than two years after it started, albeit with only digital still images initially. By reducing the scope, the project made progress and was able to deal with the very real issues of migrating from a system built in 1996 to a much more modern system.

DIMOC began with the existing storage system and after consultation with the vendor determined:

- About $160K was needed for recertification
- The system was a WORM storage type (Write Once, Read Many)
  - WORM does not work for DIMOC's requirement because as edits are made to metadata they are written back to the asset's header. This would result in multiple copies of each asset.
  - Reclaiming drive space as previous versions were deleted was problematic, time-consuming and inefficient.
- System consisted of Gen2 and Gen3 devices, Gen4 was already shipping, so Gen2 would be obsolete soon, further increasing costs.

After careful analysis by the Digital Asset Management Systems Officer, DIMOC decided that the most effective solution was to abandon the old system and acquire a new system, with new requirements.

Establishing Requirements

Given the reset of the VIMS program DIMOC also restarted the requirements process. An element of the requirements that had not previously been addressed involved reliability and security. DoD has stringent information assurance (i.e. information security) requirements and the previous efforts did not take this sufficiently into account.

This portion of the work is extensive, and not within the scope of this paper. However, to comply with the information assurance requirements, the DIMOC design included redundant systems geographically separated, one on the east coast and the other on the west coast. Additionally, the storage would have local backup capability, providing essentially three copies of every asset.

One of the goals of the new system was obtaining Authority to Operation (ATO) within the Defense Information Assurance

DOI: 10.2352/issn.2168-3204.2016.1.0.28

Certification and Accreditation Process (DIACAP), which was achieved in 2013 [1]. The DIACAP accreditation process was modeled after the National Institute of Standards and Technology (NIST) Risk Management Framework (RMF). DoD has subsequently moved to RMF in place of DIACAP, with a new round of certification and accreditation required [2].

The New Storage System

A requirements document was developed after DIMOC consulted with the storage industry and determined the type of storage that met its' needs. In addition to DIMOC requirements (capacity, communication types, drive technology, reliability, etc.), specific DoD-focused questions were asked such as[3]:

- Does the storage system meet DoD Information Assurance (security) requirements?
- Is the storage system redundant, capable of 99.9% uptime or better?
- Is the storage system capable of automated remote replication across the continent? The DIMOC replication site was to be at the Riverside, California facility, saving the cost of leasing a dedicated facility.

The procurement was published via the Federal Business Opportunities (www.FebBizOpps.gov) website, the government's procurement web site. DIMOC research resulted in selecting a business class storage solution over an enterprise solution, meaning an ATA-drive based storage was chosen rather than SAS storage. Performance of ATA was satisfactory and was, at the time, about one tenth the cost of SAS-based storage.

Requirements included storage management capabilities and replication. DIMOC procurement was not simply to obtain storage, but to have a complete solution to ensure all of the DIMOC types of uses were met. DIMOC initial procurement capacity was for about 32 terabytes (TB) at each site with redundant front-end management systems. The system consisted of the actual storage (the drives and the housing for them), two management servers, two fiber optic switches, two Ethernet switches and a server used to manage all of the above, packaged in a half-height rack fully configured – ready to go.

The contract was awarded and delivery occurred in late 2007 [4]. Soon after that, the manufacturer of the storage device discontinued the product. Not only was the model discontinued, the manufacturer stopped making all enterprise storage, opting instead to partner with another manufacturer.

Lessons Learned:

1) Often times a purchase results in buying a product nearing the end of its' lifecycle. Vendors are not often upfront about this, and are sometimes unaware of the coming change. Government procurement can also delay the purchase.
2) Definitions of terms is important. "Redundant" meant one thing to DIMOC and something else to the selected manufacturer. For example, the storage device had two fibre channel controller cards, but each card controlled sets of drives independently, with no failover capability. DIMOC overcame this gap by:
    a) Configuring the device to use each independent set of drives as a single volume in the storage management system.
    b) Combining the separate volumes into RAID 5 using the zetabyte file system (ZFS).
    c) A card failure would then simply degrade the RAID until it was replaced.

While this configuration worked, it was later abandoned as too expensive. The storage system was losing considerable capacity due to redundant RAIDs (each set of drives was also RAID 5 to protect from drive failure at the lowest level), and during the next procurement round this weakness was corrected in DIMOC's requirements.

The storage system was specified to be capable of supporting a disaster recovery system geographically separated from the production system. This part of the system was an integration developed by an international company and had numerous advanced capabilities such as thin provisioning and snapshots. This also provided valuable lessons learned.

3) Information assurance is complex, and compliance is determined at the local level.
4) Just because a system has been installed in one DoD environment does not mean success in another. DIMOC is very different from other DoD agencies. DIMOC has a Public mission that necessitates more rigid security rules than systems accessible only from within the military network.

During market research each vendor was questioned about DoD IA compliance and DIMOC was assured the systems being evaluated were approved in other DoD agencies. However, upon installation of the selected product, various issues immediately occurred, such as the use of telnet for management. Telnet is not authorized on DoD systems, secure shell (SSH) is required instead. The international company was able to rewrite their code within 30 days to meet DIMOC IA requirements. Because of their quick response to the install, DIMOC was successful and the system performed well within the constraints. Here too, more lessons were learned.

5) Bleeding edge technologies may not align with each other.
6) The Transport Control Protocol/Internet Protocol (TCP/IP) does not work well over very long distances, in this case cross-country.

DIMOC selected the ZFS file system for a number of reasons, for example the software RAID just discussed. However, ZFS is very demanding of a storage system's availability. When one part of storage management system became unavailable, and thus failed over to the secondary devices, ZFS would shut down, taking the DIMOC system offline with it. It took time to understand what was happening, and even with a brief switchover time(less than three seconds), ZFS would shut down.

Upgrades and Moving the System

As in all systems, DIMOC is faced with upgrades. The DAMS system was put in operation in mid-2007 to import existing assets, and then dual import paths were setup to keep it current with incoming assets. In April 2008, the DAMS was put in service and the AP Preserver was decommissioned. While DIMOC was fielding this new system, the organization was also preparing to move from Alexandria, Virginia to Ft. Meade, Maryland.

In 2011 the storage system's capacity was increased to approximately 72TB at each location, replacing the core storage provided with storage from a new manufacturer partnered with a system designer [5]. This new model was awarded by sole source

DOI: 10.2352/issn.2168-3204.2016.1.0.28

contract to ensure a seamless and uninterrupted upgrade process [6].

An EOL (end of life) for the servers used in the DIMOC system was published in 2010. These servers were the management servers, managing not the files but the storage configuration, drive assignments, and so on. Replacement of these servers had to be planned, again taking into consideration the complex nature of their role. The replacement servers also used a different operating system, necessitating more work to certify them.

Moving the storage system, along with the applications (everything had to be moved), to Ft. Meade was a major challenge. Rather than move the Production system or purchase all new equipment to establish a new production system, DIMOC shifted systems around the continent.

First to move to Ft. Meade was the Disaster Recovery (DR) system, an exact replica of the Production system. This meant accepting some risk since DIMOC would have to operate for a period of time without a disaster recovery system, but that risk was far out-weighed by the assurance of continued operation.

The DR system was disconnected and shipped to Ft. Meade where the team then worked to bring it on line in the new location and to update settings, etc.

Essentially the DR system became the Production system. With the DR system installed and operating at the new location, replication was established between the two in order to synchronize them, then a date was selected for the switch over.

This is a simplified description of the move process, which was complex and fraught with issues since the system was being installed onto a completely new network. The problems were compounded because the new network was implemented with more stringent controls than the old network and the support staff was stretched thin. The staff was supporting two locations along with large groups of personnel moving from Texas and other geographically separated offices in addition to the headquarters location in Alexandria, Virginia. The video storage system was not moved as part of this effort for reasons discussed later in this paper.

Sale and Architecture Changes

With the system in the new location, DIMOC upgraded the servers to replace those no longer supported. It was during this period DIMOC also achieved DoD certification and accreditation. Shortly after the upgrade to new servers system design company (the international company), was bought by a U.S.-based company with a much larger presence in the storage industry [9]. DIMOC now faced a period of uncertainty while the new company decided the future of the old product line.

Not surprisingly, technology marched on and DIMOC's original solution was overcome by new designs. The original architecture involved management servers connected to the storage shelves that contained the drives. Four years later new designs incorporated the server into the storage shelves, an integrated solution that presented new challenges in securing the new operating system (now an abbreviated version of Linux) used in the DAMS design.

With the size of the existing DIMOC storage system, wholesale replacement is difficult and time-consuming. DIMOC operations are twenty-four-hours-a-day/seven-days-a-week (24/7) with combat camera personnel assigned around the globe who need DIMOC's systems to be available 24/7, three hundred and sixty-five days-a-year. Taking the system offline to move data to a new system was problematic. And since DIMOC had purchased new servers just prior to the sale, there was concern about the fiscal impact.

The new company's management decided to sell off certain elements of the international company's business line. However, this was complicated by the intricacies surrounding Intellectual Property (IP) ownership and in the end the new owner could not make a sale. This situation hung over DIMOC for almost eight months! When the sale was called off, the new company made their customers whole with full credit for recently purchased systems to enable upgrading to their solution, with a fully formed and tested migration path.

7) Lessons learned:Buyers are at the mercy of vendors in ways that can complicate and impact organizational plans.
8) Close coordination with the vendor's staff during company transitions is critical.
9) Ensure the vendor is kept fully aware of your organization's concerns.
10) Listen to the vendor's point of contact to gain valuable insight into company thinking, which may help determine what direction to take.
11) Do not panic when companies change hands, but document all communication.

Replication

Once DIMOC completed the move to the new location in 2011, replication to the west coast site had to be re-established. The earlier replication to California from DIMOC's east coast site was working without incident (although video was not being replicated). The east coast site had a digital signal level three (DS-3) communication line to the west coast location, while the new site was upgraded to an optical carrier level three (OC-3), about three times faster than DS-3. Despite the higher speed, DIMOC was unable to successfully achieve replication. No amount of testing was able to resolve the problem for almost two years, so DIMOC sought a new solution.

A solution involving the User Datagram Protocol (UDP) had been discussed and was demonstrably faster than other methods, so it was chosen as the new transmission solution. In 2013 DIMOC completely replaced the replication capability built into the system with a UDP application using the fast and secure protocol (FASP) [10].

Using UDP results in speeds about 100 times faster than TCP/IP. Because the synchronization software is not as robust as replication technology, DIMOC still builds snapshots within the storage system for local backups. Note that the FASP technology is patented [7] and the speeds obtained are not possible without that technology.

With the change in replication methods DIMOC reduced the costs for support, since their replication solution was no longer required. However, the management capabilities and local backup (snapshots) are still only available using proprietary software (separate from the replication software).

Lessons Learned:

12) TCP/IP is not a fast protocol, and is very sensitive to configuration.
13) UDP is very fast and just as reliable as TCP/IP when implemented using the FASP protocol.

DOI: 10.2352/issn.2168-3204.2016.1.0.28

The Video Solution

Video requires far more storage than still imagery, to include the additional streaming requirement. Video imagery stores three versions of the same file, just as still imagery does, but with a far larger file size.

Storing video in the same way as still imagery (always on, spinning hard drives) is expensive, so DIMOC investigated alternative solutions.

One solution is to use tape systems. The Linear Tape-Open (LTO) format is commonly used for very large file backup due to its speed. However, tape storage is a sequential storage technology meaning that each file must be read or written to tape before the next file can start. A solution to speed up file delivery involves using algorithms to build caches of files on spinning disks. However, this solution requires diligent management to keep them from becoming overwhelmed and ensuring proper sizing of the spinning disk storage to avoid having unused spinning disk capacity, or too little capacity.

Another technology was developed in 2005 called Massive Array of Inexpensive Disks (MAID) [8], and was available in 2007. This design involves massively dense storage on spinning disks with management of the disks to reduce energy consumption. The system shipping in 2007 could hold almost a petabyte of data in about ten square feet. Today that same footprint can hold six petabytes, the only limit being the individual drive capacity.

The disks are managed so that only about twenty-five percent of them are powered on at any one time. When requests for files arrive, and the files are on disks that are idle, the system will spin that set of disks up (turn them on) and turn off unused disks that are running. Time to delivery of the first file bits is typically less than ten seconds. Because the files are on disks random file access is used rather than the sequential access tape uses. A study was commissioned by DIMOC to compare MAID against LT05 and costs were comparable [9].

Given DIMOC's physical facility limitations at the time, the MAID solution was very attractive. A search was conducted for agency-leased facilities in the area that could provide more space, as well as power, but to no avail. The MAID solution, if fully populated, used less than 30 Amps of 240V electrical power, and imposed a very small temperature increase. This when combined with the rapid and random access to the files, lead DIMOC to select the MAID solution.

Lessons Learned:

14) The MAID system did not include a file system, so more expense was incurred to purchase a file system.
15) Highly reliable
16) The company producing the MAID solution was a startup funded by VC, which proved extremely troublesome when the company was dissolved.
17) The sale to a new, strong, publicly held company saved the technology.
18) Beware of companies seeking to jump on the latest and greatest technology. Other storage vendors claimed MAID technology in their systems, but close examination demonstrated otherwise.

Moving the MAID system

The MAID is a large rack about 3 feet by 3.5 feet, covering 10 square feet. It has no feet in order to spread the weight evenly, a fully loaded system can exceed 300 pounds per square foot. In other words, it is massive. Rather than ship such a large system across the country, DIMOC moved the Alexandria system to Ft. Meade. This also saved money, as DIMOC did not have to pay for reconfiguration.

Because MAID is a separate system, DIMOC could configure the DAMS to ignore its absence during the transfer and remain online with notices to users that video was temporarily unavailable. The physical move was executed in one day, and the system was available again within the first day in the new location.

DIMOC also halted ingestion of video for that period; queuing video received for later ingestion, so no synchronization was required and no data was lost.

Why not use the Disaster Recovery site?

The original move plan involved switching to the Disaster Recovery site while the Production system was moved and reconfigured. However, there were numerous actions taking place all at the same time such as application development and organizational growth. And DIMOC's parent organization, Defense Media Activity (DMA), was not only moving to a new facility, the entire organization was being re-organized, missions were being evaluated, and there was a lot of disruption to daily business.

Also, DIMOC had only one staff member with the required skills and knowledge responsible for the move. And while DIMOC had a staff of nine information technology contractors, they did not have the authority to make decisions that involved spending government funds. So staff levels were, and continue to be, an issue.

DIMOC was extremely busy and the parent organization, who provides DIMOC's network, was also extremely busy. As a result, DIMOC suffered network issues and the staff could not keep up with application changes to the backup system. All too often they could not access the backup system as a result of faulty network configurations. Concern grew that DIMOC would not have the capability to use the DR system for an extended period. In fact DIMOC was unable to test the system until just a few weeks prior to the move, and that was only a weekend test.

The test revealed several issues, mostly with new and vital applications, as well as video capability at the remote site. DIMOC downtime for the move was unknown, and there was pressure to bring the production system online quickly in the new environment.

By moving the DR systems to the new site DIMOC had the ability to take more time, with more care, in setting up the system in the new environment. This proved very valuable as it took about four months to bring the new system online. The old system was kept running during that period, relieving some of the pressure to get the new system up and running.

Cloud Storage

As noted earlier, new technologies are always being developed, and Cloud capabilities is one of those. Future solutions involve Cloud storage and services. However, DIMOC must follow the rules set by the Defense Information Services Agency (DISA), and the rules do not always fit our unique requirements [14].

DIMOC is faced with significant restrictions as the rules are currently written and is working to overcome these issues, but the rules are evolving very quickly making it even more challenging for DMA to get ahead of the curve.

DOI: 10.2352/issn.2168-3204.2016.1.0.28

Adding to the complexity is the DIMOC initiative to use a commercial provider for the DoD VI asset system. DIMOC now hosts almost all of its assets within a commercial system. That system uses cloud storage, as well as some cloud services such as transcoding. But DISA rules address using cloud services with traffic originating within the Non-classified Internet Protocol Network (NIPRNET). Since the DIMOC commercial system cannot reside on the NIPRNET, another solution must be found.

Future Storage Requirements

One might think the storage challenges have been addressed given the use of Cloud storage, but not quite yet.

DIMOC still operates certain applications, such as a Single Sign On server, that rely on storage and database systems. Until those are moved to Cloud services DIOMC will require some level of storage, including a disaster recovery site.

Additionally, DIMOC has a mission requirement for a classified system. Cloud storage certified to hold material classified up to secret is not yet available, but DIMOC will likely be a user when it is..

The End Game

DIMOC's goal is to have storage become a simple commodity. By employing Cloud storage, DIMOC can off-load the work needed to specify, obtain and maintain the storage system. That will allow the small staff to focus more on the management, organization and use of the VI archive's content, bringing more value to the collection.

The value of an asset increases with use, the more it is used the more valuable it becomes. DIMOC is committed to making all of our assets more valuable.

References:

[1] Defense Media Activity Authority to Operate Letter, 2013.

[2] DoD Instruction 8510.01, Risk Management Framework (RMF) for DoD Information Technology (IT), 12 March 2014

[3] Defense Visual Information, Statement of Work Storage Solution, Amended 12 June 2007

[4] DISA/DITCO-SCOTT Contract: HC1013-07-P-2314, awarded 31 August 2007.

[5] Storage Contract Upgrade Contract: HQ0516-11-P-0042, awarded 20 September 2011.

[6] Storage Contract Upgrade SOW, 12 September 2011

[7] Method and system for aggregate bandwith control US 20090063698 A1: http://www.google.com/patents/US20090063698 Last accessed on 14 January 2016

[8] What is MAID (massive array of idle disks) – Definition from Whatis.com http://searchstorage.techtarget.com/definition/MAID Last accessed on 14 January 2016

[9] MAID Study - Master Book Proposal DIMOC Site Survey Report 4 June 08.pdf

[14] Defense Information Systems Agency, Department of Defense, Cloud Computing Security Requirements Guide, http://iase.disa.mil/cloud_security/Documents/u-cloud_computing_srg_v1r1_final.pdf Last accessed on 14 January 2016

Author Biography:

Paul Robinson is the DIMOC Systems Officer, the senior Information Technology officer for DVI/DIMOC. He is a retired USAF Combat Camera photographer. He has more than twenty years of experience on the leading edge of digital photography. He is a Certified Information Systems Security Professional (CISSP). Paul is an avid motorcyclist.

