

Managing the Digital World of Intelligence

Data-intensive GEOINT applications demand innovative storage solutions.

With the intense security measures employed by government agencies, particularly NGA, data proliferation runs rampant, requiring specialized systems.

*By Andrew De Lara, Chan & Associates
(www.chanandassoc.com), Fullerton, Calif.*

The entire world is literally at the fingertips of the U.S. Department of Defense (DoD). Long gone are 20th-century reconnaissance and aerial photography over battlefields dating back to World War I. Today federal government agencies use sub-meter imagery to build the “battlespace” where specially targeted individuals, groups or countries—from one corner of the globe to the other—are tracked. From Iran to Russia, Venezuela to China, between urban structures, in targeted zones within vast expanses of mountains and treacherous terrain and everywhere in between, there’s an imagery reconnaissance platform collecting geospatial intelligence (GEOINT) and monitoring threats to U.S. security.

The National Geospatial-Intelligence Agency (NGA), the U.S. agency within DoD responsible for GEOINT, is

Cutting Edge Networked Storage

Cutting Edge's ISS (inset) and MSS products provide scalable storage platforms that allow NGA and other agencies to store and retrieve massive amounts of current and future geospatial data.



charged with collecting, analyzing and distributing geospatial data to adjoining government agencies for a variety of military, civil and international needs. During its history, NGA (and its predecessor organizations) has evolved from map-based imagery products to a digital and soft-copy environment that takes advantage of myriad data sources, from airborne and satellite systems to geographic information system-based virtual globe programs such as Google Earth. The data are stored on servers with high-performance networked-attached storage (NAS) systems, available to users working with critical intelligence operations.

Challenges of Storing the Earth

The proliferation of imagery intelligence brings many unique challenges to the operators of the complex systems involved. One of NGA's biggest challenges is storing and retrieving such massive amounts of current and future geospatial data.

GEOINT is one of the most computational and data-intensive operations on the planet. Advancements in data gathering technologies, analysis algorithms, and 2-D, 3-D and 4-D geological and geospatial modeling, as well as real-time data processing, have resulted in a proliferation of data, with raw datasets greater than several terabytes in size. Accessing and managing historical data archives adds additional time and expense in this global competitive environment—where time is money.

With the intense security measures employed by government agencies, particularly NGA, data proliferation runs rampant, requiring specialized systems. To meet these needs, Cutting Edge Networked Storage (www.cuttedge.com) has provided commoditized and custom storage solutions to government data-intensive programs for more than 14 years. The company specializes in GIS environments, tailoring enterprise solutions for complex, mission-critical applications such as GEOINT.

"The resolution of satellite imagery has greatly increased, and government organizations such as NGA need to monitor every corner of the world in intervals of just a few minutes," says Cutting Edge CEO Michael Ehman, whose company has worked with NGA and adjoining firms and agencies.


Cutting Edge's Integrated Storage Solution (ISS) and Modular Storage Solution (MSS) provide scalable storage platforms that help users at NGA and other agencies consolidate GIS applications with vector and raster data sets for continuous and timely information access, exploration and publication. The MSS also is being deployed with ESRI and Google Earth programs, among others.

Overseeing a Sea of Data

One company that knows first hand about the unique challenges posed in the GEOINT environment is Simulyze (www.simulyze.com), which provides software products and associated integra-

tion and operations support to government and commercial customers in the areas of situational awareness, real-time data processing, visualization and analysis.

Simulyze's Flight Control, a commercial situational-awareness software product, integrates numerous planning and real- and near-real-time data sources. The product integrates geospatial foundation information such as maps and imagery with operational planning data, such as air-tasking



An interactive Google Earth globe displays real-time data processed and output by Flight Control. The display shows real-time air traffic and ship traffic over and around the United States. A user can configure customized alerts for items of interest that are displayed when the event occurs.

orders; near-real-time operational data, such as the Global Command and Control System Common Operating Picture and Predator data; and near-real-time intelligence data. Flight Control processes many disparate data sources, normalizes the data and correlates/fuses the data into an integrated decision-support system environment.

Chris Schmitt, a C4 ISR Systems Analyst and Google Earth specialist for Simulyze, points out that proper storage volume is critical for a successful GIS.

"About two and half years ago, a government organization tasked Simulyze to design and implement an architecture to process, store and display real-time and other types of data into a variety of visualization and analytic tools," says Schmitt. "One component of this overarching architecture involved Simulyze integrating its real-time Flight Control Enterprise Servers and Flight Control Desktops with Google Earth. This capability would allow users to visualize the large volumes of data processed by Simulyze in Google Earth."

An 8-terabyte Cutting Edge direct-attached storage (DAS) unit was purchased initially to serve as the system's primary repository, providing adequate storage to satisfy the client's initial requirements. But facing rapid user growth and increased availability of high-resolution imagery, the system filled the 8 terabytes within a couple of years, sparking demand for substantially increased storage capacity. Cutting Edge then provided Simulyze and its government client with two 48-bay, 36-terabyte NAS units.

"We literally filled all 8 terabytes of the original DAS and needed to ensure that the future solution could handle the massive data involved," explains Schmitt. "The NAS provided us the ability to seamlessly create significantly larger and more refined globes."

Powered by its specialized EdgeWare 64-bit storage management software, the Cutting

The Google Earth Fusion server literally builds the globe, while the Google Earth server ultimately publishes the completed globe for users to view. The Flight Control desktop allows a user the capability to customize the filtering and alerting of the data displayed in Google Earth. Finally, the Cutting Edge NAS, which is mounted on the Fusion server, is essentially seen as a drive inside the Fusion server.

Ultimately, Google Earth servers used in GIS operations transform a jumbled mess of data into a coherent, explorable, interactive globe—no small feat. Building a Google Earth globe entails building an imagery layer, a terrain layer and a vector layer. The Google system stores the raw files—the raw imagery, raw terrain files and raw vector files—and creates "Google assets," which are files with a Google extension that Google's servers recognize and are able to read.

These assets then are pointed toward the raw file locations—usually a storage device such as a Cutting Edge NAS or DAS. From there, the Fusion server ingests the files and pre-processes them for the organization to begin building a globe.

Once the files are processed and the globe is built, the globe is published to the Google Earth/Map server, which is a separate server. The Google Earth server subsequently renders the finished map for a user to view. Users then can connect to the Map server through an IP address that points to its published database folder—much like a Web server. From there, users may literally monitor the world.

Not surprisingly, that world consists of incredibly immense terrain and imagery files. For example, a 20-terabyte Google Enterprise globe is currently in use. But even with the large scale of raw files and storage challenges, the much refined workflow allows for an efficient use of space. After Fusion processes the 20 terabytes of data, the finished product may be reduced to a 5-terabyte globe.

One known globe contains approximately 10 terabytes of raw resources, with a 36-terabyte Cutting Edge NAS on the Fusion server and an 8-terabyte DAS on the Map server. With that, clients are confident that they will be capable of accommodating their anticipated growth levels for the next couple of years. "Our staff recently collaborated with Simulyze on an optimized iSCSI protocol solution that will enhance performance of major GIS globe applications by 150 percent," says Ehman. "Cutting Edge and Simulyze will continue to work together in the future to provide storage solutions optimized for the challenges of the GIS marketplace."

The flexibility, security and specialized performance of such data storage solutions for GEOINT entities, along with the efficiency of the overall server system, ensures seamless operations running now and in times to come. EJ

Ultimately, Google Earth servers used in GIS operations transform a jumbled mess of data into a coherent, explorable, interactive globe—no small feat.

Edge solution hit the ground running with the Google Earth system. With complete NAS and iSCSI-IPSAN functionality, the utmost in data throughput capabilities and a plethora of features to ensure that the stored data remain available and resistant to failures, the company's experience with the GIS community was leveraged by this particular government organization to boost its GEOINT operations.

"The demands of the future are definitely on the minds of current system engineers and managers," says Schmitt. "With the Cutting Edge NAS, we are able to store all the resources that we need and build a globe that contains all the imagery and associated GIS data of the world—all critical to our clients and their overall missions."

Google Earth and the GIS Process

Firms such as Simulyze always make certain that current and future raw file requirements are addressed in the architectural design phase, as they're critical to ensure that storage needs are adequately addressed. This GIS architecture, which Simulyze implemented at one of its clients' sites, is composed of five specific pieces of hardware/software.

The Flight Control enterprise server processes and synchronizes the real-time and planning data from a multitude of available sources to create a single integrated common operating picture.