WHITE PAPER

Enabling low-latency processing, scalability and data security in the field

DAVID FREEMAN
Senior Product Manager, Sensor Systems
Mercury Systems

Placing Data Center Storage Technology at the Tactical Edge

The advancement of sensor technology such as radar, image capture and radio frequency devices has resulted in a massive increase in the amount of data collected by defense customers. This information, known as big data, has proved to be difficult to store, access or analyze in real time. Solving this storage issue is vital to ensure that the critical decision-makers of today and tomorrow can take well-informed actions and avoid limiting the capabilities of current, emerging and future sensor technology. This paper explores the data storage challenges that must be overcome and presents new technology from Mercury Systems and VAST Data Federal as a compelling solution.

COMPUTING CHALLENGES AT THE TACTICAL EDGE

Big Sensor Data and High Latency

Mission-critical defense systems now employ countless sensors to collect imaging data and to track electromagnetic communications, radar responses and physical vibrations such as sonar and voice recognition. These sensors are continually being improved, providing greater precision and expanded data streams with every upgrade or advancement. The result of these increasingly powerful sensors is an influx of big data, to the point where some individual systems must process terabytes of input per second.

Department of Defense customers are using compute-heavy applications such as artificial intelligence, signals intelligence and sensor fusion to analyze and convert this ever-increasing amount of data into useful and actionable intelligence. Additionally, this big data is also rich with development possibilities for new and sophisticated applications such as facial recognition, airframe performance and electromagnetic waveforms.

However, the storage systems deployed on most vehicles, aircraft or vessels do not provide enough available storage or bandwidth to take full advantage of the information being gathered. For example, the onboard storage system of an F-35 aircraft can only store 5-10 minutes of combined data collected by all its sensors. Instead, much of the data collected is offloaded via satellite/radio transmission to a traditional data center for processing.

These data centers, often located far from the battlefield, can easily provide the amount of storage needed. Yet the latency—or the time it takes to transmit and receive data—does not make data centers a suitable solution for the U.S. DoD's vision for a digitally connected and agile military. If connected processing and storage resources were available locally, it could eliminate this time-consuming process, and increase both the amount of accessible data and the ability to make well-informed decisions at the edge.

SWaP and Environment Considerations

Taking equipment out of a data center and placing it at the edge normally requires making technological compromises. This is because traditional storage systems aboard land, air and sea vehicles face size, weight and power (SWaP) limitations and therefore cannot capture, store and process the amount of data collected by modern sensors.

Defense and aerospace computing components also face a high likelihood of exposure to different environmental conditions, such as high vibrations, sudden physical impact, altitude, temperature and humidity, and parts must generally be able to meet the following requirements:

- MIL-STD-167 — Shipboard equipment vibration and shock
- MIL-STD-461 — Electromagnetic interference emissions
- MIL-STD-810 — Environmental field readiness and conditions such as shock and vibration
- MIL-S-901 — Installation in shock-isolated cabinets that protect against external explosives

Hard disk drives (HDD) traditionally used for large amounts of storage use electricity to power moving parts, known as platters and spindles, when reading and writing data. These moving parts are susceptible to damage or malfunction from extreme conditions, and are also heavy, noisy, and significant heat emissions generators. Solid state drives (SSDs), also known as flash storage, do not contain moving parts and are faster, noiseless, dense, lighter, and only use electricity when writing data and not when reading data. However, they have traditionally been much more cost-prohibitive than HDDs when it comes to providing a vast amount of storage.

For these reasons, most data centers today use a data tiering method, in which a hybrid mix of various HDDs and SSDs are partitioned into different sections, such as a redundant array of independent disks (RAID) for backups, and numerous other disks and partitions for archival, rarely accessed, mission-critical, time-sensitive, and system data. While effective, this separation of data across drives requires a large, physical air-conditioned space and can also create latency issues depending on the type of data being accessed.

An Evolving U.S. Defense Strategy

The Pentagon's Joint All-Domain Command and Control (JADC2) is a concept and directive meant to provide U.S. armed services and allied partners with an improved ability to detect, make sense of and act on information across all domains. The vision to create a wholly connected military in which information flows quickly, securely and accurately across land, air, sea, space and cyber is set to transform the nation's defense capabilities — and also increases the need to place large amounts of secure, rugged and accessible storage at the edge. For example, JADC2 will require the use of processing- and data-intensive technology such as:

- Artificial intelligence (AI) and machine learning (ML)
- Sensor fusion processing
- Signals intelligence (SIGINT) electronic warfare (EW)
- Big data query engines
- Video/image capture
- High-performance computing simulations and deep learning training
- Log analytics, online backups, and archiving

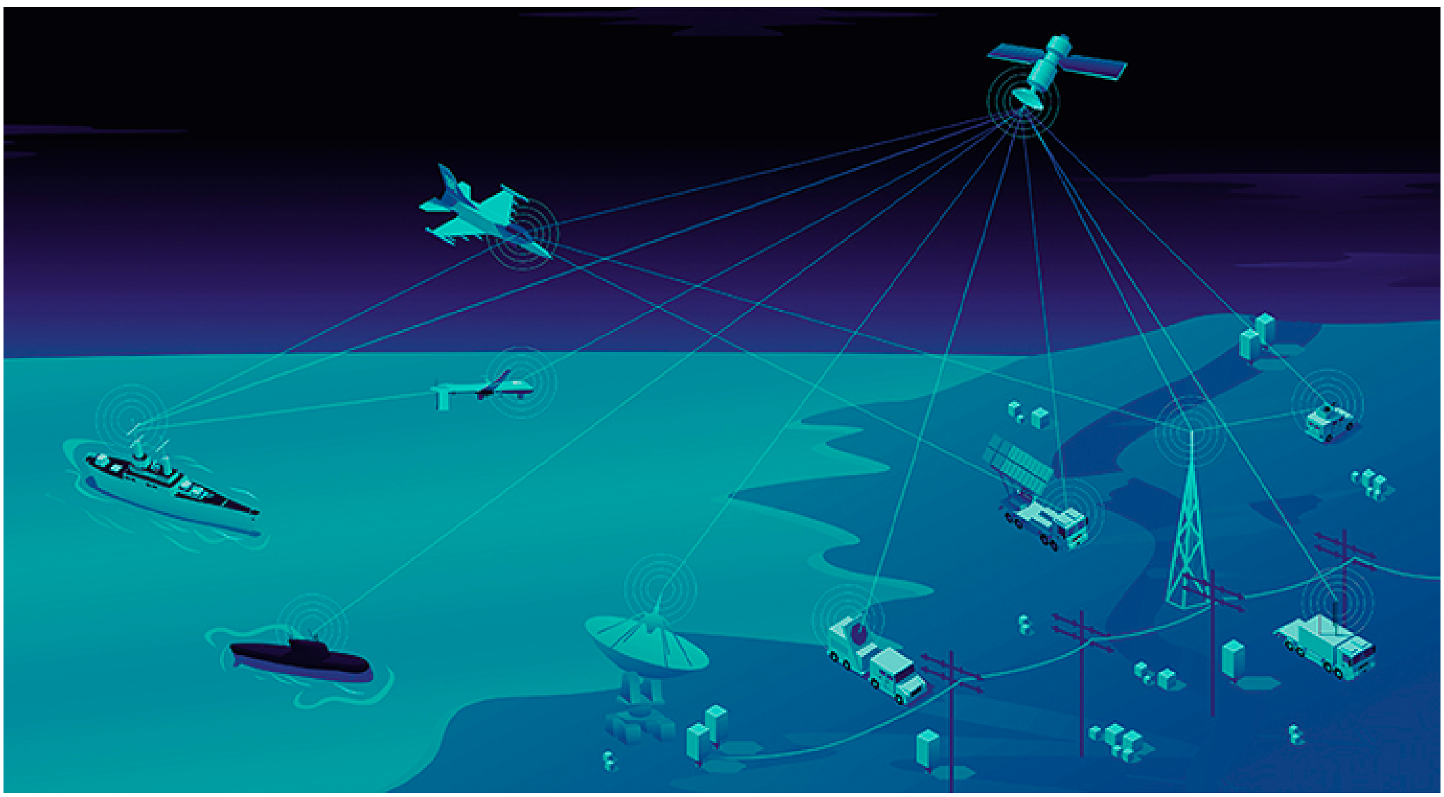

POTENTIAL DATA STORAGE PLATFORM BENEFACTORS

- Future Vertical Lift (FVL) is the U.S. Army's plan to create a family of five different-size military helicopter platforms. Ultimately meant to replace helicopters such as the UH-60 Black Hawk, AH-64 Apache and CH-47 Chinook, the FVL aircraft are being designed to share common parts such as sensors, avionics and engines. Determining the onboard sensors and computing and processing electronics needed for FVL requires the collection of large amounts of data from whole flight and specific, complex scenarios such as landing in fog, brownouts and blackouts.
- The U.S. Navy's Boeing P-8 Poseidon is a multiple-mission patrol aircraft. It performs anti-submarine, anti-surface, intelligence, surveillance and reconnaissance (ISR) missions in addition to search and rescue. It is also armed with torpedoes and anti-ship missiles, is capable of deploying and monitoring sonobuoys, and can coordinate with drones. The P-8, also equipped with the Navy's most advanced sensors and radars, could greatly benefit from low-latency access to petabytes of storage.

RUGGED DATA STORAGE (RDS) SYSTEM AS A SOLUTION

RDS is an all-flash network-attached storage (NAS) system from Mercury Systems that is ruggedized for mission-critical aerospace, defense and commercial edge applications. Designed for speed and versatility, it puts low-latency processing, scalability and data security at the edge. Designed in collaboration with VAST Data Federal, it provides end users with a universal storage and flash-based cloud for all their data, eliminating the bottlenecks and complex storage tiering that is traditionally required with hard disk drives and other media. This approach, combined with a modular and disaggregated design, gives customers the ability to deploy multiple petabytes of data storage in the harshest of environments and to securely access big data and artificial intelligence (AI) and machine learning (ML) applications.

RUGGED DATA STORAGE (RDS) SYSTEM INFOGRAPHIC

RDS comprises three major rugged subcomponents (C-Box, D-Box and Switch Box) and is deployable starting in standard 5U 19-inch rack mount fixtures. To save space, the D-Box in RDS is 2U high, taller than the 1U high D-Boxes used in data centers but just 22 inches deep as compared to the 38-inch depth of its data center cousin used in commercial, non-rugged applications. The use of ruler-based E.1L flash storage, a slim rectangular and front-loadable memory stick, also saves space as compared to traditional 2.5-inch U.2 SSDs.

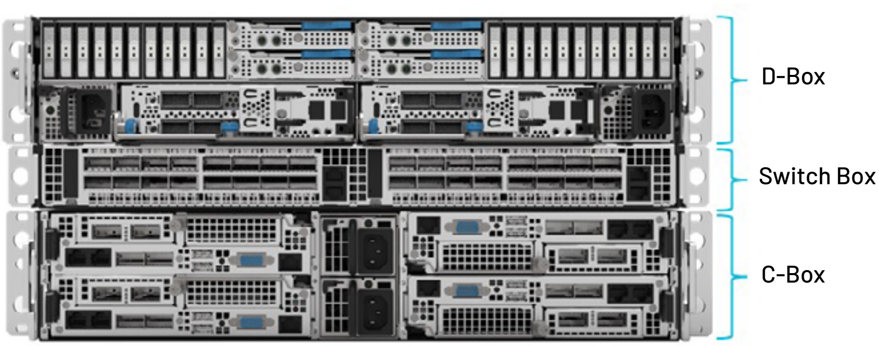

Container/Computer Server Box (C-Box)

- High-performance computing power from four nodes, each containing dual-socket Intel® Xeon® processors for a total of 128 cores and 1 terabyte of processor memory.
- C-node protocol servers handle all file system-related logic, such as listening for requests, interpreting them and then accessing drives over high-speed NVMe.

Data Storage Box (D-Box)

- Array of 22 removable high-speed flash/ solid state drive (SSD) memory E1.L rulers can be ejected and replaced individually.
- Four integrated NVIDIA Bluefield data processing units (DPUs) move and parse data via use of programmable acceleration algorithms for data compression and security.
- NVMe-to-Ethernet connections achieve read speeds up to 40 GB/s and write speeds up to 5 GB/s.

Switch Box

- Dual Mellanox 16-port 100G Ethernet switches configured to support communications between the D-Boxes and C-Boxes and access by end users.

Rugged

- RDS is designed to meet U.S. military standards, such as MIL-STD-167, MIL-STD-461, MIL-STD-810 and MIL-S-901 when used in shock-isolated cabinets.

Features

- Hyperscalable all-flash data infrastructure at the cost of disk storage
- Low-latency access to all data (up to 100x faster than HDD)
- Disaggregated commodity Ethernet fabric
- Multi-petabyte-scale NAS and object storage and composable flash cloud
- Artificial intelligence/machine learning-ready infrastructure
- Advanced data compression and short-depth chassis

Flash-Based and Universal Storage

RDS ensures low latency and speed via use of fast consumer-grade flash memory known as QLC (quad-level cell) for storage. As compared to HDDs, this dense and nonmoving commercial flash reduces power consumption by 66% and processing power use by 25%. The low-latency QLC flash memory in the D-Box is also networked to multiple nonvolatile memory express (NVMe) connections in the Switch Box, with each supporting sustainable high-performance speeds of up to 40 GB/s read and 5 GB/s write.

The VAST Data software used in RDS was designed from the ground up to use QLC flash memory and to provide a single-tier flash cloud known as universal storage to users. This approach means all data and the entire storage architecture itself are contained across one single file system. This universal storage approach, combined with the fast NVMe connections, makes all stored data quickly accessible and also eliminates the need for RAID backup.

RDS software also lessens the amount of physical storage needed with similarity-based data reduction algorithms that compress and de-duplicate data from 2x to 10x, depending on the type of data. For example, with just a 3:1 compression ratio, end users can store 1.7 petabytes of data on just 676 terabytes of raw flash storage. VAST software also assigns algorithmic signatures to all data for later accessing — using a similar concept behind a popular cellphone app that assigns signatures to song data and allows users to identify tracks — to further ensure quick access to data.

Upgradable, Maintenance and Reliable Design

With a modular and disaggregated design, RDS end users can scale up processing power (C-Boxes), flash storage (D-Boxes) or NVMe switches (Switch Boxes) independently and as needed to meet their mission's unique requirements. A front-loadable design allows users to upgrade to more advanced storage and processing hardware as it becomes available, and to pull and replace any malfunctioning parts. Built-in redundancy in both the hardware design and the VAST software also ensures no information is lost when flash rulers are replaced. Instead, the RDS system will rebuild the data on the new rulers.

VAST software also uses erasure codes to protect data from drive failures. This is accomplished by writing data across multiple drives so that if one drive fails, it can be rebuilt into the remaining free space of the storage system or onto a replacement drive. To help avoid such failures, software-driven write cycles extend the useful life of the QLC flash to 10-plus years.

IN A FLASH: THE RISE OF VAST DATA

Since its first shipments in 2018, VAST Data has established itself as the fastest-selling storage infrastructure company in history.

This success is due to a new approach to storage called universal storage that the company developed and pioneered. The concept uses all-flash storage and revolutionary algorithms to provide exabyte levels of data that is accessible at the speed of flash memory. This solution is being deployed into data centers around the world and across industries, rapidly eliminating the economic arguments for data tiering using disk storage.

Vast Federal, a VAST Data subsidiary focused on modernizing government agency data, is also providing a simple and scalable all-flash cloud to numerous U.S. government organizations, including the Department of Defense, Department of Energy, NASA, Department of Health and Human Services and National Oceanic and Atmospheric Administration.

The company's development partnership on the RDS system combines the power of the VAST universal storage platform with Mercury System's proven rugged servers to deliver a scalable, readily accessible and reliable flash cloud to the edge of defense.

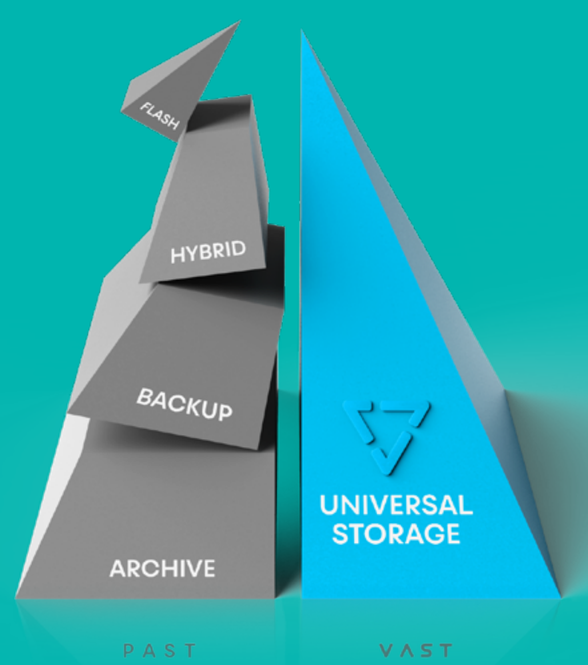

ELIMINATING MISCONCEPTIONS ABOUT DEPLOYABLE STORAGE

Misconception 1: Multi-petabyte databases require a traditional data center.
Reality: Enormous amounts of flash storage can be deployed at the edge for highly data-driven applications, without compromising performance, capacity or scalability.

Misconception 2: Storage tiering is more cost-effective than flash at the petabyte level.
Reality: Fast flash is affordable when it is used for all data and treated as one universal pool. This also eliminates the need for traditional hard disks, archival cassettes and other slow, complex mixed media.

Misconception 3: Affordable QLC flash can't provide the write-cycle longevity needed for long-term defense storage requirements.
Reality: Advanced striping algorithms allow QLC flash storage to support 10-plus years of high-performance reliability and endurance in deployed systems.

SUMMARY: FUTURE-FORWARD DEFENSE STORAGE

With RDS, Mercury Systems and VAST Data are making cutting-edge data center technology available for defense, industrial and other harsh-environment applications. Their developmental partnership on RDS connects complementary technology such as proven rugged components and a universal storage architecture to provide access to petabytes of network-attached storage at the edge and helps our armed services maintain a tactical advantage today and tomorrow. Engage with the Mercury team to explore how the RDS can meet the most demanding storage requirements, expand the effectiveness of current applications and enable completely new applications.

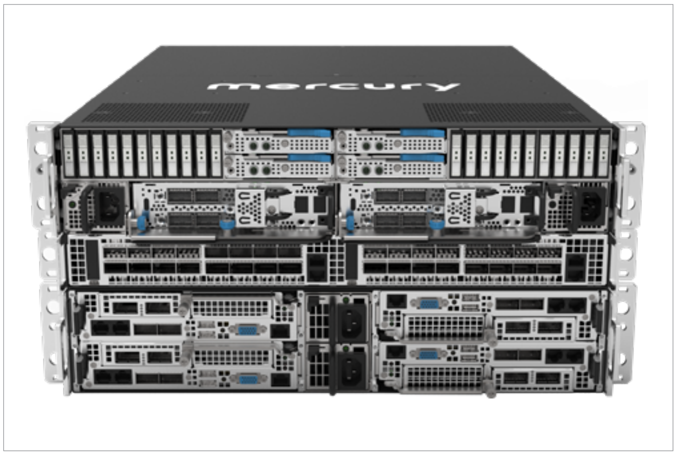

Rugged Data Storage (RDS) System

Corporate Headquarters
50 Minuteman Road
Andover, MA 01810 USA
+1 978.967.1401 tel
+1 866.627.6951 tel
+1 978.256.3599 fax

International Headquarters
Mercury International
Avenue Eugène-Lance, 38
PO Box 584
CH-1212 Grand-Lancy 1
Geneva, Switzerland
+41 22 884 51 00 tel

Learn more
Visit: mrcy.com/storagesystems
Contact: sales@mrcy.com

The Mercury Systems logo and the following are trademarks or registered trademarks of Mercury Systems, Inc.: Mercury Systems and Innovation That Matters. Other marks used herein may be trademarks or registered trademarks of their respective holders. Mercury believes this information is accurate as of its publication date and is not responsible for any inadvertent errors. The information contained herein is subject to change without notice. Intel and the Intel logo are trademarks of Intel Corporation or its subsidiaries

MADE IN USA

© 2022 Mercury Systems, Inc. 8198.00E-0123-wp-RDS