

DDN | Whitepaper

Any Threat, Anywhere, Anytime

Scalable Infrastructure to Enable
the Warfighter

ddn.com

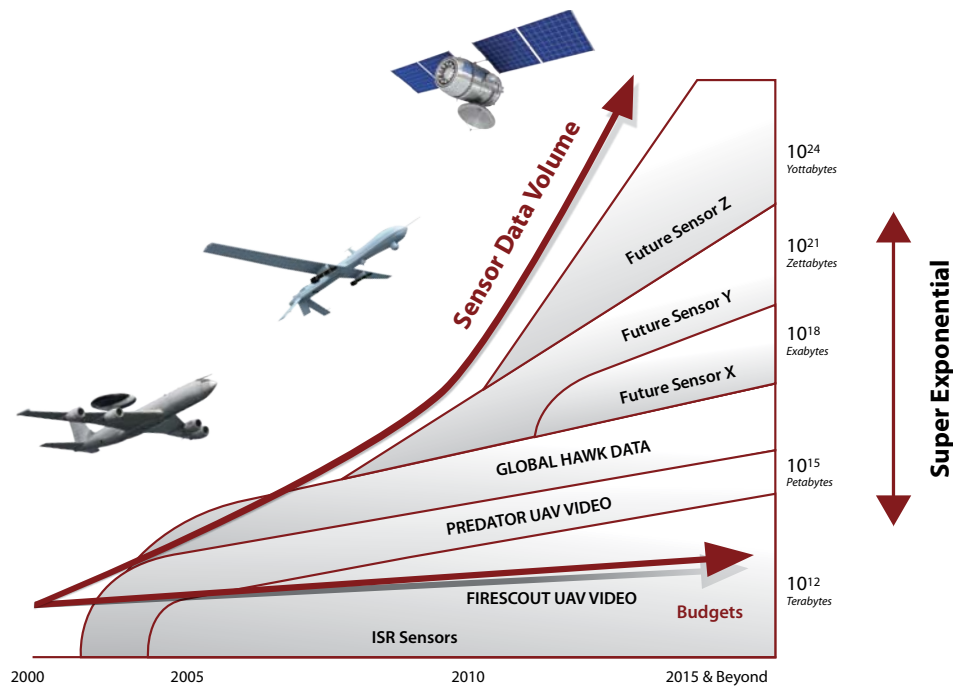
I N F O R M A T I O N I N M O T I O N [™]

Table of Contents

The Big Data Challenge and Opportunity.....	3
Challenges with Traditional Storage.....	4
Why DataDirect Networks.....	5
Intelligence Ecosystem.....	7
Game Changing Product Portfolio for Big Data.....	7
Direct Attached/Array Platforms.....	8
File Storage Platforms.....	8
Hyperscale Big Data Cloud Storage.....	8
Engineered for a Decade for Big Data.....	9
Sample DDN Federal Customers and Applications.....	10
Case Study: DDN solves Big Data challenges for Naval Research Labs.....	10
The Challenge.....	11
The Solution.....	11
Why DDN.....	11
Summary – DDN Products can Handle.....	11
Any Threat.....	11
Anywhere.....	12
Anytime.....	12

The Big Data Challenge and Opportunity

The proliferation, management and analysis of intelligence data is a fast growing concern for Federal and Defense agencies. They are challenged with exponential data growth from a variety of sources: from satellites, to Unmanned Aerial Vehicles (UAVs), to increasingly higher-definition sensors installed on surveillance aircrafts. This new data class of machine-generated data is referred to as Big Data, and reflects datasets that are so large that it is beyond the ability of conventional hardware and software technology to effectively capture, store, administer and analyze. National defense agencies today are at the center of Big Data creation trend – and are increasingly challenged to manage infrastructure growth to scale with next-generation sensor data.



Super Exponential Growth in Big Data

An example of Big Data used for Intelligence Surveillance and, Reconnaissance (ISR) purposes is the hundreds of hours of sensor data streams per flight generated by each Unmanned Aerial Vehicle (UAV) or Wide Area Surveillance system. All this sensor data can provide invaluable information to the warfighter, but the challenge is to store and sort through it to find actionable data. For example, the Defense Advanced Research Projects Agency's (DARPA) Autonomous Real-Time Ground Ubiquitous Surveillance Imaging System (ARGUS-IS) is like 60 to 100 Predator UAVs looking at a particular area all at once, and uses 368 five MegaPixel detectors capable of generating more than 260 Giga bits per second of raw data¹. This is about 3000% more than what a standard network storage system can support today. ISR agencies are finding it more and more

¹Overview: Autonomous Real-time Ground Ubiquitous Surveillance – Imaging System (ARGUS-IS); DARPA, 2010-02-19

difficult to capture and analyze this amount of data. This problem is projected to get much worse. Modern generation drones currently record video in about 30 directions but there are requirements for as many as 65 with much higher levels of resolution and quality. This means that these sensors are expected to produce 1000x more information just in the next few years.

“ In the Air Force alone, UAVs churn out 39 video feeds, 24 hours a day – a figure that could jump to 3,000 feeds as airborne surveillance programs proliferate. ”




Not being able to store and analyze this information to identify relevant intelligence can lead to missed opportunities to find insights into new and emerging kinds of threats, and provide rapid, real-time decision support to the right people, in the right place, at the right time. It can lead to reduced situational awareness, delays in decision-making and potential lapses in security.

ISR agencies need comprehensive and scalable Big Data strategies. This paper examines the challenges of Big Data and the difficulties faced with using traditional storage solutions to solve those challenges. It then examines how ISR agencies are deploying DDN solutions to power next generation defense infrastructure to handle Any Threat, Any Where, Any Time.

Challenges with Traditional Storage

The challenge with Big Data is that the datasets are so large that it becomes difficult to capture, store, search, share and analyze the data within a reasonable period of time. According to Gartner², an industry analyst firm, Big Data challenges span three dimensions – Volume (amounts of data being created), Velocity (speed at which the data needs to go in and out) and Variety (range of data types and its sources).



Dimensions and challenges of Big Data

²“Gartner Says Solving ‘Big Data’ Challenge Involves More Than Just Managing Volumes of Data,” Press Release, STAMFORD, Conn., June 27, 2011

To appreciate the relative size, it's instructive to understand the scale of data being created from one high resolution ISR camera. If a video camera images a piece of land of dimensions 10 km by 10 km allocating one byte per square meter and captures this at 30 Frames Per Second (FPS), it generates a single stream of raw data of approximately 3000 megabytes per second of bandwidth. Adapting to a higher resolution or dimension represents a geometric progression – moving from 1 m resolution to a ½ m resolution increases bandwidth requirements by a factor of four. A traditional storage technology like a Network Attached Storage (NAS) system is not capable of handling a single stream of that bandwidth, making it incapable of ingest from these high resolution ISR cameras. Furthermore, if an agency was to accumulate a month's worth of that data, it would require over 10 Petabytes of storage, which is beyond the scaling limits of a traditional NAS system. Now multiply this requirement a thousand fold in an actual theatre of war where there may be numerous "sources" as the number and types of image sensors increase, and many "destinations" like central mission control, mobile mission stations and mobile ground forces that request access to the overhead video. This exponential increase in requirements shows the challenges faced by ISR agencies trying to solve Big Data requirements using traditional storage technologies.

Why DataDirect Networks

DDN has leveraged over a decade of experience gained from powering the world's fastest supercomputers to design a portfolio of products that can scale independently across all dimensions and handle the diverse set of Big Data challenges. Let's take the example described in the previous section and understand how DDN products solve those challenges across the Big Data dimensions.

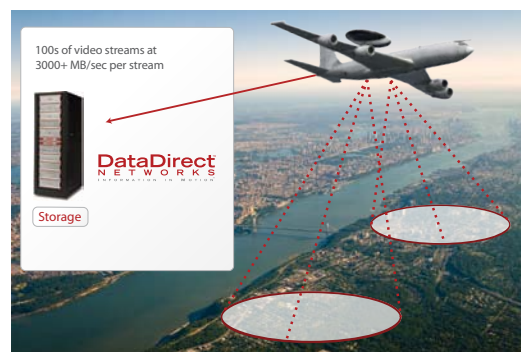
Why DDN for Big Data

With Traditional Network Attached Storage



Surveillance Area and Camera Resolution Limited by bandwidth of backend storage

With DDN Solutions for Big Data



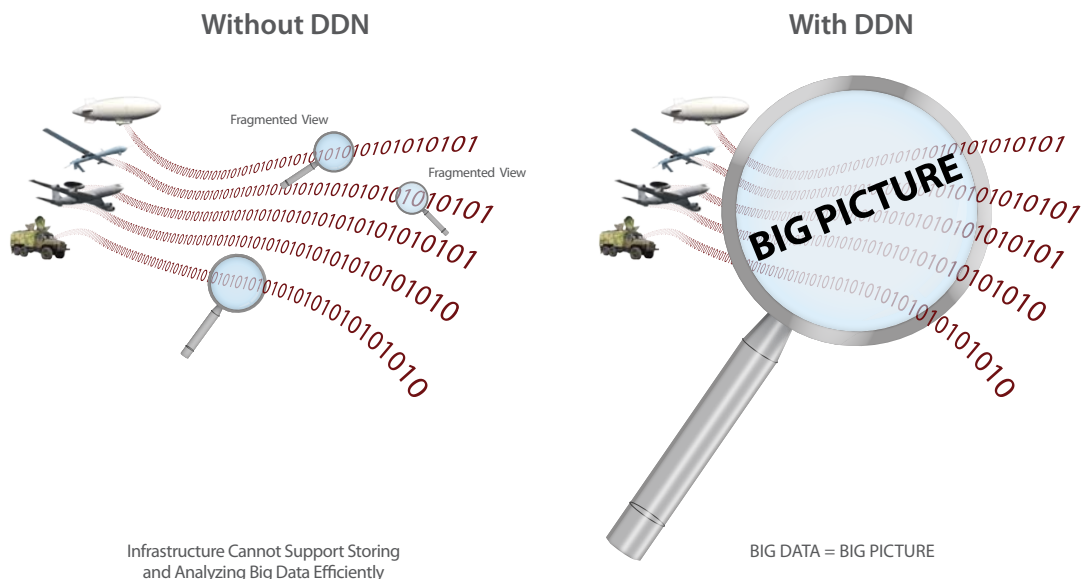
100x greater area Surveyed with streaming Optimized Big Data solutions from DDN

- **Velocity** – DDN portfolio of products includes high performance object and file system that utilize high levels of parallelism with low latency, high speed interconnects to handle the enormous rates of content creation in a theatre of war. For example, a single DDN solution can simultaneously ingest multiple streams of surveillance video, at a rate of

3000 MB/sec per stream that is generated by next generation high resolution ISR cameras. This is an order of magnitude higher than what a traditional storage solution like a NAS can support and this lets agencies survey 100x larger areas of land by supporting higher resolution cameras and much higher number of cameras per storage solution.

- **Volume** – DDN’s elegant building block architecture scales seamlessly on demand to hundreds of Petabyte in size and billions of data objects in number, where each object can be an asset like an individual image or video that is deemed valuable. This scale gives ISR agencies the ability to deploy a single large system that provides them with a consolidated view of all their intelligence data. Compare this with a traditional storage solution that can only scale to a fraction of this size, forcing agencies to deploy multiple small storage islands with added overhead and complications.
- **Variety** – DDN solutions are optimized to handle a wide range of data and workload types – from highly transactional to highly sequential workloads, small to extremely large files, to the ability to tag data with rich contextual information to simplify manageability and search. For example, DDN object based solutions can be used to tag data based on other properties of the data beyond its file name, like geospatial metadata, co-ordinates, altitude, bearing etc. This gives agencies the ability to manage billions of types of data and to capture co-relation to other intelligence information. Traditional storage solutions that rely on locating data based on a simple file path name break down when the number of files exceeds a few millions, as they frequently do in the theatre of war. Compare that with DDN solutions that can handle hundreds of billions of types of objects.

In essence, DDN products allow ISR agencies to look at the entire dataset efficiently rather than sampling it and getting fragmented views. Getting the big picture allows them to develop deeper insights than what could be gained from looking at only a portion of the data. It gives them better accuracy in their results because the results are derived from the entire data set. The final benefit is speed in the analysis of big data, allowing them to identify threats and make better decisions much faster.



ddn.com

I N F O R M A T I O N I N M O T I O N™

Intelligence Ecosystem

Partner focused, DDN is aligned with the world's leading defense System Integrators and Vendors to provide optimized solutions. Sample partners include



Game Changing Product Portfolio for Big Data

DDN's portfolio of products includes the world's highest performing block storage arrays designed for data-intensive environments. Coupled with these array platforms are the world's most scalable file storage appliances to scale-out for performance and scale-deep for capacity. Each file storage appliance is designed to fit a specific Big Data requirement, ensuring that the customer gets the right solution to their problem. Finally, DDN offers an object based storage technology that radically simplifies and improves how intelligence information is stored, distributed, and accessed across multiple geographically dispersed sites. The next section examines these platforms in further detail.

Direct Attached/Array Platforms

S2A9900

Recipient of multiple awards, including Frost & Sullivan's Digital Media Storage Innovation of the Year and HPCWire's Best Product and Technology of the Year, the S2A9900 is capable of high speed ingest and delivery of concurrent data streams with real-time performance guarantees and zero performance penalty on a disk or enclosure failure. This ensures that critical surveillance applications will get the required bandwidth at all times with no loss of video streams or information even during disk or enclosure failures.



SFA12K/SFA12K-20E

The SFA12K is the world's fastest system for Big Data. It is capable of providing peak performance of up to 40GB/sec from a single system, greatly accelerating ISR applications and giving agencies the ability to make split second decisions. It is designed to offer unrivaled balanced performance for highly transactional (small files or random access) as well as high-bandwidth (large files or streaming access) workloads. This allows agencies to consolidate Big Data applications onto a single platform to lower TCO and simplify manageability.

The SFA12K-20E is an embedded storage appliance that uses our exclusive, next generation In-Storage Processing™ technology to enable the file system to live right inside the storage controllers. In-Storage Processing eliminates latency, storage gateways, hardware, complex networking and makes it extremely simple to manage - reducing the power, cooling and footprint requirements for Big Data ISR applications.



File Storage Platforms

Built upon the Storage Array Platforms are the world's most scalable File System appliances. Each appliance is specially designed to solve a particular Big Data challenge.

- **NAS Scaler** is highly scalable clustered Network Attached Storage (NAS) solution that supports standard protocols and fits seamlessly into existing datacenters to handle the diverse and dynamic demands of Big Data.
- **xSTREAMScaler** is a high performance file system with extensive data management features designed for defense organizations looking to create Full Motion Video (FMV) solutions over both high-speed Fiber Channel as well as Local Area Networks (LAN).
- **GRIDScaler** is a hybrid SAN and LAN file System with robust support for a variety of operating systems concurrently accessing the file system over parallel file system as well as NAS protocols.
- **EXAScaler** is an Open Source parallel file system that provides Linux-based applications near-wire-speed data transfer capabilities and unbounded capacity.

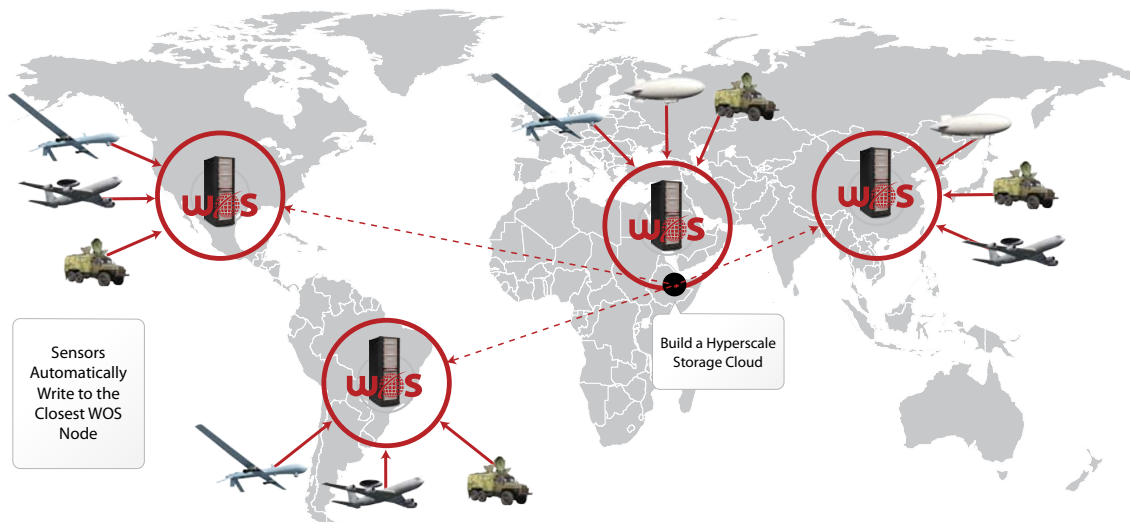


Hyperscale Big Data Cloud Storage

Web Object Storage (WOS) is a scale-out object based storage system designed for the most content-intensive federal environments in the world. WOS can be used to build a geographically distributed surveillance platform that supports millions of sensors and digital video cameras. Data is automatically written directly to the closest WOS node to ensure lowest latency and most optimal performance, while policy-based replication mirrors the data automatically to multiple sites, ensuring uninterrupted access to surveillance data even in the event of a whole location outage.

One unique feature of this object based solution is the ability to tag pieces of intelligence data with information like location, speed, geographic co-ordinates etc for easy search and discovery. These tags provide rich contextual information to the decision maker to reduce the cycle time of the “kill chain” and enhance situational awareness by integrating intelligence data across disparate sources.

Hyperscale Big Data Cloud Storage Capture and Collaboration



Engineered for a Decade for Big Data







	DataDirect™ NETWORKS	Other Vendors
Broadest Portfolio of Products to Handle Big Data challenges	✓	✗
Concurrently handle hundreds of thousands of high resolution sensor and FMV surveillance streams	✓	✗
Uninterrupted data capture and performance during system correction	✓	✗
Supports clustered supercomputers with 10,000s-100,000s of CPUs	✓	✗
Self-Healing & autonomous cloud storage at massive scale	✓	✗

ddn.com

INFORMATION IN MOTION™

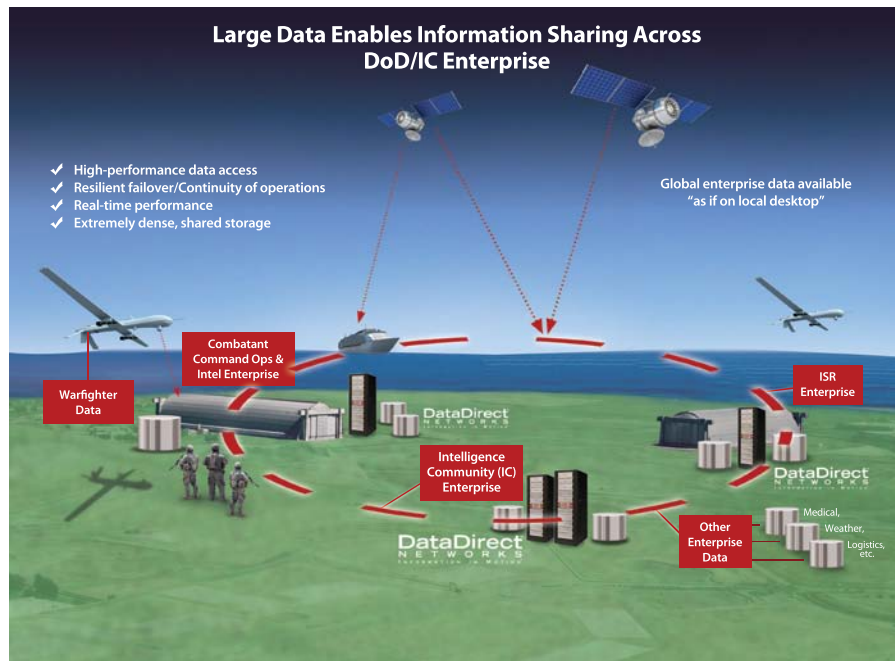
Sample DDN Federal Customers and Applications

Following are sample federal agencies that have deployed DDN systems to solve their Big Data challenges.

Agency						
Application	Sensor Data Assimilation & Archiving	Signal Intelligence	Media Exploitation	High Performance Computing	Modeling & Simulation	Satellite Ingest

Case Study: DDN and Silicon Graphics solve Big Data challenges for Naval Research Labs

One of the many large-scale defense deployments of DDN infrastructure to solve Big Data challenges is at the Naval Research Laboratory's (NRL) Large Data Joint Capabilities Technology Demonstration (JCTD) program. DDN and Silicon Graphics (SGI) delivered an advanced image storage, retrieval, and processing infrastructure to NRL that allowed it to process and share real time intelligence in the battlefield.



ddn.com

INFORMATION IN MOTION™

The Challenge

The Challenge was to gather, share and analyze battlefield imagery from satellites, UAVs, and first responders to allow analysts to operate on it in real-time from remote locations as if the information were local to their desktops.

The Solution

DDN supplied each Large Data JCTD site S2A storage appliances that were connected with low latency, high speed InfiniBand based interconnect to provide shared access to a Petabyte-scale storage pool, while SGI's Altix servers enabled real-time visualization and manipulation of the incoming images. DDN and SGI worked together to integrate their respective systems with WAN encryption and replication technology supplied by additional vendors.

Why DDN

- **Scale** – DDN was able to provide a scalable storage system that could handle hundreds to thousands of sensors streaming data in an active theatre, with the amount of raw data that needed to be stored, processed and shared running into Petabytes and even Exabytes.
- **Real-Time** – DDN systems delivered Real-Time performance and ensured no loss of critical data even with hardware failures.
- **High Performance** – DDN systems were able to mirror the raw data in real time over low latency, high speed InfiniBand networks and handle large files and high bit-rate streaming.
- **Total Cost of Ownership** – DDN systems provided world leading density and energy efficiency that substantially reduced the power, cooling and footprint requirements for NRL.

Summary – DDN Products can Handle

Any Threat

DDN storage systems enable today's defense organizations to capture, store and analyze massive amounts of Big Data to gain real-time insight into any potential threat. With DDN storage systems, these organizations are guaranteed the scale, performance and reliability required to capture and analyze content at the highest resolutions and provide rapid, real-time decision support to the right people, in the right place, at the right time.

Anywhere

DDN's comprehensive portfolio is purpose-built to be deployed anywhere to meet the diverse needs of defense organizations. The portfolio of products includes file and block based solutions for centralized or distributed sites, ruggedized solutions for the battlefield, embedded solutions for simplicity and data center efficiency and object based hosted solutions to geo-distribute intelligence data and provide synchronized, global access.

Anytime

Defense Organizations expect the highest levels of availability, reliability and durability from their infrastructure. DDN is the only organization capable of providing guaranteed quality of service and access to critical surveillance data, regardless of misbehaving components or media. Always at full-speed, there is no performance or data loss even when drives, drive channels or even entire disk enclosures go offline.

Please visit www.ddn.com or call 1-800-837-2298 to learn more about how DDN's technology and products can help you.

DDN | About Us

DataDirect Networks (DDN) is the world's largest privately held information storage company.

We are the leading provider of data storage and processing solutions and services, that enable content-rich and high growth IT environments to achieve the highest levels of systems scalability, efficiency and simplicity. DDN enables enterprises to extract value and deliver results from their information. Our customers include the world's leading online content and social networking providers, high performance cloud and grid computing, life sciences, media production organizations and security & intelligence organizations. Deployed in thousands of mission critical environments worldwide, DDN's solutions have been designed, engineered and proven in the world's most scalable data centers, to ensure competitive business advantage for today's information powered enterprise.

For more information, go to www.ddn.com or call +1.800.837.2298.

Version 10/11

ddn.com

I N F O R M A T I O N I N M O T I O N [™]