


DATA ON-DEMAND WITH VCINITY & DELL: ACCESS REDEFINED



Pratyush Panda

Specialist 2, Inside Product
Dell Technologies

Anjali Jayanti

Specialist 2, Inside Product
Dell Technologies

The Dell Technologies Proven Professional Certification program validates a wide range of skills and competencies across multiple technologies and products.

From Associate, entry-level courses to Expert-level, experience-based exams, all professionals in or looking to begin a career in IT benefit from industry-leading training and certification paths from one of the world's most trusted technology partners.

Proven Professional certifications include:

- Cloud
- Converged and Hyperconverged Infrastructure
- Data Protection
- Data Science
- Networking
- Security
- Servers
- Storage

Courses are offered to meet different learning styles and schedules, including self-paced On Demand, remote-based Virtual Instructor-Led and in-person classrooms.

Whether you are an experienced IT professional or just getting started, Dell Technologies Proven Professional certifications are designed to clearly signal proficiency to colleagues and employers.

Learn more at www.dell.com/certification

Table of Contents

1. Abstract.....	5
2. Introduction	6
3. Internet	6
3.1. TCP/IP.....	6
3.2. TCP over UDP	8
3.3. How is TCP affected by Throughput and Latency?	9
3.3.1. Effect of Throughput	9
3.3.2. Effect of Latency	10
4. Vcinity.....	11
4.1 Understanding the traditional issues and roadblocks.....	12
4.2 How does it work?	13
4.3 Remote Direct Memory Access.....	14
4.4 Implementing RDMA in VTP?	14
4.5 Packet tracking and Retransmission	15
4.6 Role of FPGA.....	15
5. Vcinity Data Access Platform.....	16
5.1 VDAP Application and Working.....	16
6. Security Aspects: Maintaining Data Integrity over Transmission	17
7. Product offerings by Vcinity – Ultimate X® (ULT X) and Radical X (RAD X) family.....	18
8. NetApp’s Flexcache vs Vcinity’s VDAP.....	19
9. Exploring the edge, hybrid, and multi-cloud solutions of Vcinity and Dell: use cases and benefits.....	20
9.1 Edge Deployment.....	20
9.2 Hybrid-Cloud Deployment.....	21
9.3 Multi-Cloud Deployment.....	21
10. Vcinity & Dell: Edge Access Solution (EAS)	22
11. Conclusion.....	23
12. Bibliography	24

Abstract

Data is increasing at an inconceivable rate than ever imagined. As per IDC reports, by 2025 we can expect 175 zettabytes of data worldwide, which not only must be stored and managed but also has to be processed to derive useful insights. Imagine edge data this huge, will need a lot of computing resources for it to be processed, in turn amplifying the need for the data to always reside near the compute, so that it can be processed at much higher speeds. Today most of our data usually happens to be present on-prem or in the cloud. This would give rise to the pervasive issues of duplication of data at on-prem, cloud, or colo environments again near the environment where the compute resides, in turn aggravating the complications, inefficiencies, and capital expenditure of scaling storage infrastructure while managing to keep a track of storage ecosystem namespace.

However, data nevertheless will be generated and will go and reside where it has to. But these zettabytes of data are again duplicated at various intervals for processing or ease of access just leads to multifold more data and inefficiencies as we discussed. To avoid these various problems and limitations of our modern-day infrastructure customers are looking for a resolution that they usually satisfy by implementing various caching mechanisms or WAN Optimization.

Vcinity produced Vcinity Data Access Platform (VDAP) and in partnership will Dell Technologies provides customers with revolutionary solutions with their Edge, hybrid, multi-cloud, and the recently announced Edge Access Solution. These solutions help connect any compute anywhere, to any data existing anywhere in the world.

While TCP is heavily affected by latency and the speeds drop exponentially in the very first few seconds of data transfer over noticeable latencies, Vcinity Transmission protocol (VTP) is built upon HPC concepts of RDMA which maintains unimaginably stable speeds even at high latencies of 1,800ms while sending data, which would be same as going around the globe a few times. Hence, there is no requirement to copy the entire file locally, instead, the client just asks for the file, and they just get it. It simply acts as a regular NAS mounted to the server, allowing you to simplify the process of bringing in the data you need and when you need it.

This article focuses on how Vcinity redefines data access by making the WAN a global LAN and the technology behind it. The key offerings by Vcinity & DELL such as edge, hybrid, and multi-cloud and their use cases. Case study on edge data transfer from the US east coast to the US west coast and the striking results. The latest and hottest offering – the Edge Access Solution (EAS) which provides universal access and global namespace for Dell PowerScale and ECS storage systems allowing them to unleash the extensive capabilities of UDS solutions. Real-time processing of AI/ML workloads via efficient utilization of compute power with live data and cloud-bursting techniques. The Ultimate X product offerings and specifications. Competitive study on NetApp FlexCache and how Dell's new-found partnership with Vcinity provides customers an edge over the "any data anywhere" approach compared to the caching techniques of FlexCache.

2. INTRODUCTION

Computing infrastructure has massively improved over the years, a key result of innovation in Science and Engineering. Each day there is something new, something better than the last. The same goes for storage systems and networks. Storage systems have improved so vastly that a gazillion of data is being generated from various platforms such as social media, media & entertainment, R&D, IoT devices, education and healthcare institutions, banking firms, etc. With this ever-growing list and at the same time requirement of this data to be duplicated at multiple levels so that it can be accessed by end-users (e.g., Netflix movies at different parts of the world), processed to derive valuable insights (e.g., raw data from IoT devices) or the need for it to be archived for future use/reference. This has led to the thought that this must be stopped and managed more effectively. The list goes on and on regarding how data can be utilized for various things and what necessary steps we must take to secure it, but have we ever thought about how could this data be stored efficiently or made available anywhere in a split of second? The fact that it has been growing exponentially and in about a few years it might become unmanageable or simply extremely inept to keep spending for the capital expenditure (capex) on storing and managing this data as it is being perpetually generated and on top of it duplicated at various intervals.

Storage systems have improved and evolved a lot over time. From a time when we just had punch cards in the 1880s and magnetic tapes in the 1950s, to today 3rd decade of the 21st century where we are using NVMe M.2 SSDs, which are extremely fast. Today the cloud and cloud providers are providing the largest storage capabilities actively manage the services while you enjoy unlimited storage on a pay-per-use basis. Cloud, which is actively being used by consumers and business firms today stores zettabytes of data. While indeed storage systems have improved, despite all the innovation on the network side, it has still been lagging. Not in terms of capabilities or infrastructure but innovation to transport extremely vast amounts of data wirelessly to various parts of the world at real-time speeds, which TCP/IP has been unable to do. Even with 5G in the picture with speeds up to 10-20 gigabits per second, it faces considerable blockade during long-distance transmission of data over TCP. With Vcinity VDAP (Vcinity Data access platform) the traditional paradigm of moving the data before using it is being changed. Now if you want to use the data, simply mount it when and where you need it and you will be able to process data in real time without copying even a byte of the information. To understand this better, let us go through the various challenges in current infrastructure and its limitations to see how the effective solution was bought up.

3. The INTERNET

Let us revise some basics. The Internet is simply a massive network of interconnected devices. It can be computers, smartphones, IoT gadgets, etc. It is formed by the connection of all the LAN (Local Area Network), MAN (Metropolitan area network), and WAN (Wide area network). While an enterprise might have a combination of LANs and MANs leading to a private WAN. To put it together, the internet is simply a public WAN. Though the Internet is said to have its birthday in 1983, the work on the vision of connected communication of devices started back with ARPANET (Advanced research projects agency network) which then led to FTP and TCP/IP in 1970. Hence, basically, the protocol exists way before the internet. From then on, TCP/IP has been upgraded significantly with various versions, but it still lacks to produce data at real-time speeds across the world. Let us see what TCP/IP is, and what it has to offer with its drawbacks.

3.1 TCP/IP

You might hear TCP and IP going hand in hand multiple times as TCP/IP. Do you ever wonder what is the difference between TCP and IP or are they the same? First things first, TCP stands for Transmission Control Protocol whereas IP stands for Internet Protocol. IP addresses look like 4 numbers ranging between 0 to 255 each separated by a decimal (for e.g., one of Google's IP addresses is 142.250.192.142). This stands true for the 32 bits IPv4 address. The newest IPv6 address has 128 bits and hence solves the issue of providing unique IP addresses for newer devices in the future. It has the capability of addressing 2¹²⁸ devices uniquely which is somewhere around 340 trillion IP addresses.

The key difference is that while IP works on getting the data over the internet to the correct computer (or any other device), TCP's job is to get the data to the correct application at the destination. TCP operates at the Transport layer in the TCP/IP model while IP operates at the internet layer. Let us look at the TCP/IP model to understand it better.




Figure 1: The TCP/IP Model

The application layer acts as the interface between the application and the network. It consists of various protocols relating to a variety of application services and is responsible for requesting connection and also formatting/forwarding data to the transport layer. Whenever data from the application layer approaches the transport layer, at first the data is broken into smaller chunks so that it is easily manageable over the Internet. These small chunks of data are called packets. Once these packets are ready, the transport layer adds the header containing the port number of the application that it needs to be sent. Imagine how we multitask on our computers, for example, using Gmail, surfing the web on chrome, and using MS teams. All this happens seamlessly because the transport layer at the destination computer exactly knows which application it must send the packets to, from the port number. A few common port numbers and their functions are listed below for reference –

Port Number	Used for
20 & 21	File Transfer Protocol (FTP)
22	Secure Shell (SSH)
25	Simple Mail Transfer Protocol (SMTP)
53	Domain Name System (DNS)
80	Hypertext Transfer Protocol (HTTP)
123	Network Time Protocol (NTP)
179	Border Gateway Protocol (BGP)
443	HTTP Secure (HTTPS)
587	Modern SMTP with encryption
3389	Remote Desktop Protocol (RDP)
8080	Web Server

Table 1: Ports and their uses

TCP is not aware nor does it care if the packet reaches the right computer on the network. All it cares about is making it reach the right application on the destination device. The task of routing the packet to the right destination, that is the work of IP. IP exists in the internet layer. Remember how the TCP used to add the header at the transport layer to carry the port numbers? IP works in an exactly comparable way. It adds the IP header which contains the IP address of the source and destination to the incoming frames from the transport layer. Hence, TCP/IP work together for our data to be transported reliably over the internet.

Source Port (16 bits)		Destination Port (16 bits)	
Sequence Number (32 bits)			
Acknowledgment Number (32 bits)			
Data Offset (4 bits)	Reserved Bits (6 bits)	Control Bits (6 bits)	Window (16 bits)
Checksum (16 bits)		Urgent Pointer (16 bits)	
Options (Optional)			
Data (Optional)			

Table 2: TCP Header Format

Post all this, the packets are sent to the network access layer where it converts signals to be transferred via physical media. The packets are then received at the other end before they proceed further through the stack. The internet layer removes the IP headers. After obtaining the port number from the TCP headers the transport layer then redirects the packets to the right application. The packets are then sequenced back together to form up the entire data.

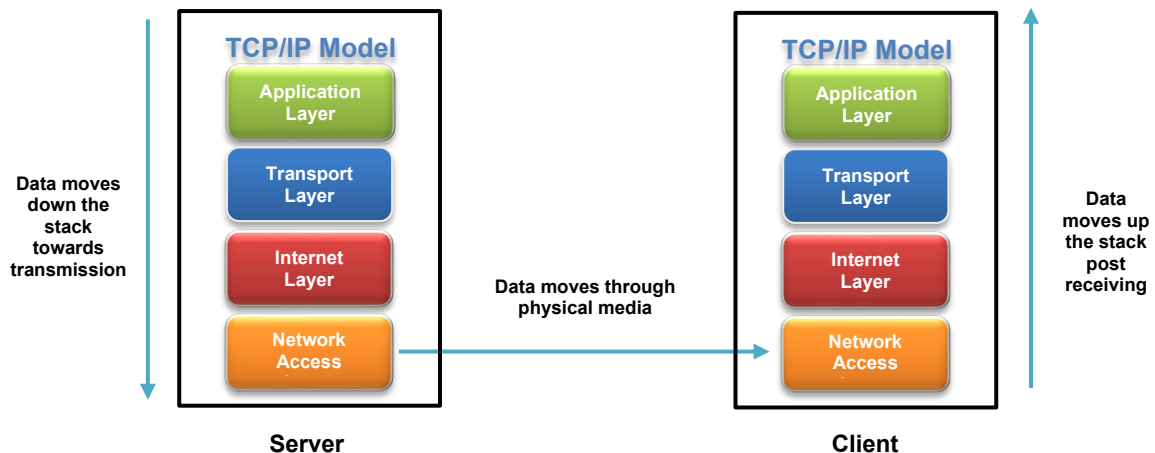


Figure 2: TCP/IP Stack flow process

3.2 TCP over UDP

Another protocol other than TCP, existing in the Transport layer of the TCP/IP stack is UDP. UDP or User Datagram Protocol is built with the same intention as TCP, i.e., to transport the data to the correct application at the destination, yet in a vastly different manner. While TCP is a reliable, connection-oriented protocol, UDP is a connection-less protocol. TCP focuses on a 3-way handshake (SYN, SYN-ACK, ACK), thus establishing a grounded connection between the sender and the receiver before starting the data transfer. This is backed with the TCP error checksum and various other useful functionalities like 16-bit urgent pointer or window size.

The TCP urgent pointer allows the data that has to be sent first on priority basis to move up the waiting queue. It also contains information on where the urgent data ends, hence signifying the position of the priority

data. The Window size, on the other hand, lets the host (sender) who is sending the data to know the maximum size of data the receiver is willing to accept.

UDP avoids all this to become a quite simple, faster transport protocol that is at least 3 to 5 times faster than TCP. UDP removes all the functionalities and the handshake creating the reliable connection, thus removing all the overhead from the header that leads to a decrease in overall latency. With this trade-off of reliability, it attains speed. UDP sends the data to the receiver continuously, least bothered about the fact that whether or not it is being received at the receiver end.

Hence, in most cases despite the speed that UDP provides, it is not the most reliable or widely used protocol. It is used in cases where the loss of data is not too critical like audio-video streaming, VoIP (Voice over IP, where we all have faced call quality due to dropping packets), gaming, or any other time-sensitive communication. Despite all this, huge video streaming platforms like Amazon Prime and Netflix still use TCP for reliable transmission. Hence, TCP which is the most common transport protocol over the internet is used in more than 90-95% of the use cases.

In summary, TCP is designed to provide reliable communication, but this reliability comes at a cost in terms of increased overhead and latency, which can affect the overall throughput. The impact of this on throughput will vary depending upon the specific network conditions.

3.3 How is TCP affected by Throughput and Latency?

3.3.1 Effect of Throughput

TCP is a connection-oriented protocol designed to ensure the delivery of all packets and minimize packet loss. However, its reliability can lead to re-transmission delays, particularly in networks with high latency, resulting in a poor quality of service.

There are several factors that can affect the throughput of the Transmission Control Protocol (TCP), including:

- **Network congestion:** When the network is congested, there may be delays and packet losses, which can lead to retransmissions and decreased throughput.
- **Bandwidth:** The available bandwidth on the network can affect the rate at which data can be transferred, and therefore the throughput.
- **Latency:** The time it takes for a packet to travel from the sender to the receiver can affect the throughput, as the sender may have to wait for an acknowledgement (ACK) before sending more data.
- **Window size:** The size of the sliding window used in TCP flow control can affect the throughput, as a larger window size allows for more data to be sent before receiving an ACK, but it can also lead to increased delays if the buffer is too large. The purpose of windowing is to prevent the sender from overwhelming the receiver with too much data. If the receiver's buffer is full, it will send an acknowledgment packet with a reduced window size, indicating that it can only accept a smaller amount of data. The sender will then slow down its transmission rate accordingly.
- **Retransmission timeout:** The time it takes for retransmission to occur can affect the throughput. A shorter timeout can lead to faster retransmissions, but it can also increase the number of retransmissions and decrease the throughput.

- **Quality of Service (QoS) settings:** A QoS setting that prioritizes certain types of traffic over others can impact the throughput, as prioritized traffic may receive higher bandwidth allocation, while non-prioritized traffic is slowed down.
- **Buffering:** Buffering can improve throughput by allowing more data to be sent before receiving an ACK, but it can also lead to increased delays and decreased performance if the buffer is too large or too small. If the buffer is too large, it can lead to increased delays, as the sending device may wait for the buffer to fill up before sending the data. On the other hand, if the receiving device has a small buffer, it may not be able to store all the incoming data, and may have to drop packets, which can lead to retransmissions and decreased performance.
- **Application layer protocol:** The application layer protocol used on top of TCP can also affect the throughput, as some protocols are more efficient than others in terms of data transfer.

A rule of thumb to obtain the TCP rate is,

$$TCP\ Rate = MSS/RTT,$$

where *MSS* = Maximum Segment Size and *RTD* = Round Trip Delay

Hence, as TCP rate is inversely proportional to RTD, where RTD is affected by a range of factors like distance, number of routers/switches, queuing latency, etc.

3.3.2 Effect of Latency

Latency is the time taken for a packet to travel from the sender to the receiver, can have a significant impact on the performance of the Transmission Control Protocol (TCP).

One way that latency can affect TCP is through the use of sliding window mechanism in TCP flow control. The sliding window is used to control the amount of data that can be sent before receiving an acknowledgement (ACK) from the receiver. As the latency increases, the time it takes for an ACK to be received by the sender also increases, which can lead to the sender having to wait longer before sending more data.

Another way latency can affect TCP is through the retransmission timeout. TCP uses a retransmission timeout to detect lost packets and retransmit them. As the latency increases, the time it takes for a retransmission timeout to occur also increases, which can lead to longer retransmission delays and decreased throughput.

Additionally, high Latency can lead to the sender's buffer being filled with packets that are waiting for ACK, this can cause the sender to stop sending new packets and wait till the buffer is empty before sending new packets, this can lead to further delays.

In summary, latency can have a significant impact on TCP performance by affecting the flow control mechanism and the retransmission timeout, which in turn leads to decreased throughput and increased delays.

Jorge Crichigno's paper (Ref.1) effectively illustrates the relationship between RTT and throughput with near-zero packet loss.




Figure 3: Throughput vs RTT (Ref. 1)

TCP Reno is a specific variant of the Transmission Control Protocol (TCP) congestion control algorithm. It is named after the city of Reno, Nevada, where it was first implemented. TCP Reno is an evolution of the original TCP algorithm, known as TCP Tahoe. It improves upon TCP Tahoe by quickly reducing the sending rate in the event of packet loss, rather than waiting for a timeout. This helps to prevent congestion in the network.

TCP HTCP (High Speed TCP) is another variant of TCP congestion control algorithm. It is designed to work well with high-bandwidth, high-latency networks. It uses a larger congestion window than TCP Reno, which allows it to transfer more data before experiencing packet loss. Additionally, it uses a more aggressive retransmission scheme to quickly recover from packet loss.

Both TCP Reno and TCP HTCP are widely used in today's networks. They are both designed to improve the efficiency and performance of TCP over high-speed networks, by allowing more data to be transferred before experiencing packet loss and quickly recovering from packet loss when it does occur.

The illustration labeled "Figure 3" illustrates how TCP Reno and TCP HTCP are implemented in real-world scenarios, where the round-trip time (RTT) still has an impact on even the most advanced and commonly used TCP protocols.

Vcinity, Inc. revolutionized the industry by introducing a revolutionary novel approach for handling and utilizing data, breaking away from traditional methods.

4. VCINITY

Vcinity is a software company that focuses on creating solutions for data storage and management. These solutions address issues with network latency by utilizing techniques such as UDP, WAN optimization, edge caching, and deduplication. However, for companies with media and entertainment workloads, these methods may not be effective. As a result, there is a need for enterprise companies to duplicate data in multiple locations to improve processing and user access.

With a goal to "connect any compute anywhere, to any data everywhere it exists, from any physical or virtual device to any other physical or virtual device," they crafted a solution to bridge the gap between data and application/users.

Vcinity has developed a novel approach to data movement and access over wide area networks (WANs) by rewriting the Transmission Control Protocol (TCP) from the ground up using Remote Direct Memory Access (RDMA). This technology allows for extremely fast and reliable data transfer, enabling unprecedented speeds and distances for on-premises and cloud-based systems. With such minimal latency, it feels as if the data is right next to the computer, eliminating the need for caching. The technology, based on RDMA, allows for LAN-like speeds over a WAN connection, effectively turning the entire network into a global LAN-like configuration. We will discuss more on RDMA in section 4.4.

4.1 Understanding the traditional issues and roadblocks

The traditional method for data processing relied on copying data to be close to the compute because network standards were not advanced enough to allow fast data movement. This resulted in the need for multiple copies of data at various locations for data processing. Vcinity's technology breaks this traditional paradigm by allowing the user to access data over a WAN at speeds that are comparable to a LAN as if the data is located next to the compute, eliminating the need for multiple copies of data.

The technology was developed with few ground rules such as not using –

(i) WAN optimization - It focuses on various methods such as dedupe or compression to reduce the data that must be transferred over WAN. This helps in decreasing the overall time and bandwidth required while improving the speed and responsiveness of the application.

WAN optimization is an issue in the case of M&E workloads as they do not allow or show any significant downsizing post-compression methods. Additionally, for critical enterprise or Federal government data, which is heavily encrypted, alternative solutions are needed.

(ii) UDP – Despite the fact that UDP can hold stable throughput even at high RTTs and packet loss, it does not have any flow control or mechanism to keep a track of the packets sent. We would not want to transmit data to a client who is not ready to receive it, which will end up dropping packets. Right?

(iii) Edge caching – Edge caching is a technique that involves predicting which data will be accessed next based on the history of previous access and preloading it into memory or providing a local copy to the edge computer to improve access speeds. However, this approach may not be effective in scenarios with a large number of clients accessing a Distributed File System (DFS) or in situations where data is rapidly changing. Additionally, edge caching may not be a suitable option for sensitive information as it involves copying or manipulating the data. As a result, it may not be the best solution for achieving quick access and fast data transfer.

(iv) Platform and infrastructure - Another important consideration is the development of a solution that can easily integrate with the current standards of wide area networks (WANs) and existing infrastructure. Many organizations may be unwilling to overhaul their data centers to achieve high speeds, and the internet may not easily adapt to changes in protocols. A desirable solution would be one, which is application-agnostic, file system-agnostic, storage-agnostic, and network-agnostic, thus embracing current technology without modifying it, and enhancing it to reach modern standards for accessing data at LAN-like speeds from anywhere in the world.

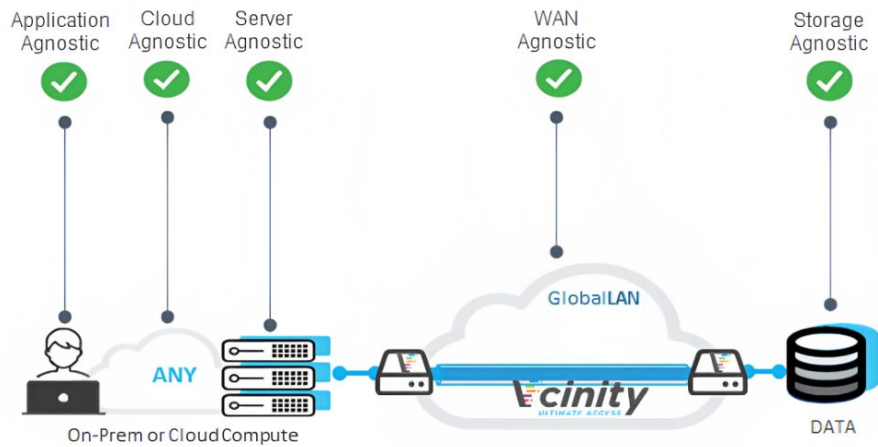


Figure 4: Completely agnostic of the entire environment (Ref. 2)

4.2 HOW DOES IT WORK?

The VDAP (discussed in section 5) utilizes Vcinity Transmission Protocol (VTP) which came out to be an innovation marvel as it provides 90% space in a packet for the data leaving only 10% as the overhead. This is a significant improvement compared to traditional protocols such as TCP, which use 90% of a packet for overhead and only 10% for data. In contrast to TCP, UDP boasts a substantial enhancement in throughput, with a **threefold increase**. However, this advantage is accompanied by the drawback of an unreliable connection, an absence of flow control, and a lack of acknowledgment. Nonetheless, these characteristics can be incorporated at the application level, thereby resolving the issue by obtaining a solution at a higher point in the stack.

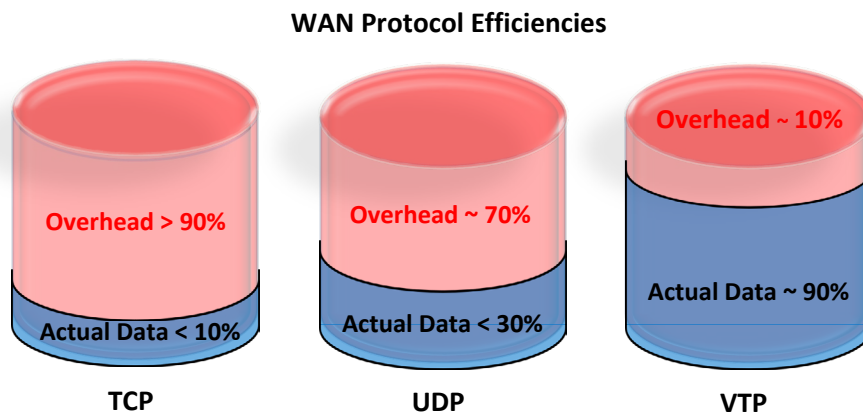


Figure 5: WAN Protocol Efficiencies

VTP leverages various methods to provide real-time speeds while creating a virtually lossless network by sustaining 90% of the frame containing data, as opposed to the 90% overhead of TCP. This results in remarkable throughput while enabling the transfer of near-maximum data over a WAN during a single packet transfer.

One may question the impact of the round-trip time (RTT) and packet drops on this technology. It is at this juncture that the underlying technology comes into play. VTP is built upon the foundation of Remote Direct Memory Access (RDMA) technology, which centers on the remote access of memory-to-memory. Before delving further into VTP, it would be beneficial to gain a deeper understanding of RDMA.

4.3. REMOTE DIRECT MEMORY ACCESS (RDMA)

It is a technology that allows a computer's network interface to directly access the memory of a remote computer, without involving the remote computer's CPU. This can greatly improve the performance of network-based operations, such as data transfer, by reducing the amount of CPU overhead required. RDMA is used in several high-performance computing and storage applications, including InfiniBand, iWARP (Internet Wide-area RDMA Protocol), and RoCE (RDMA over Converged Ethernet).

In the case of RDMA, the data does not need to be copied within memory. Instead, RDMA allows a computer to directly access the memory of another computer over a network, without the need to copy the data to its own memory first. This is known as "zero-copy" networking, because the data does not need to be copied or moved between memories, which makes the data transfer extremely fast and efficient.

With RDMA, the data transfer takes place directly between the memory of the two computers, with the data being moved in large chunks, rather than being copied into a temporary buffer first. This reduces the amount of memory bandwidth needed, which can increase the performance of the data transfer.

This is the main advantage of RDMA over traditional networking methods, where data is typically copied multiple times between different memory buffers, adding overhead to the data transfer process. It is also worth noting that RDMA can also be used to transfer data from one memory location on the remote computer to another, without ever bringing the data into the local computer's memory. This is sometimes referred to as "in-place" data transfer.

RDMA requires the use of specialized network adapters, also known as Host Channel Adapters (HCA) or Converged Network Adapters (CNA), that are designed to support RDMA protocols. These adapters have specialized hardware engines that can perform the direct memory accesses required for RDMA, and they typically connect to the host computer via a high-speed interface such as PCI Express (PCIe).

4.4. Implementing RDMA in VTP

VTP implements RDMA in its architecture to obtain direct access to the memory of the system in another location for rapid access. These RDMA headers are added to the chunks of data before sending it via the WAN. It is due to this, RTT has barely any effect over VTP. The process of keeping a track of data loss is discussed later in this section. Using VTP, the VDAP now establishes a virtual lossless network tunnel that operates in place of traditional TCP/IP protocols. It is called virtually lossless because of the exceedingly high utilization ratio of total available bandwidth. Almost 90% of the bandwidth is used to transfer only the data and the rest is the overhead. This is achieved without the use of compression, deduplication, or caching techniques, but rather through memory-to-memory transfer via byte-range access. This approach allows for the provision of specific byte-level data requested by the user, rather than the entire piece of information, thus increasing the efficiency of the architecture at a granular level.




Figure 6: VTP Header

4.5. Packet Tracking and Retransmission

You may be wondering about the measures taken to ensure data integrity in the process of accessing data via VTP.

Similar to TCP, VTP uses sequence numbers to track outgoing packets. In the event of data loss, the Vcinity initiator sends the data to the Vcinity peer at the receiving end. The peer then informs the initiator of any lost sequence numbers and only that specific data is resent. For instance, if 100 packets were sent to the peer and the sequence number 30 was lost, rather than resending the entire window of data, only the data relative to that sequence number is resent. This helps in reducing data loss and improves performance.

4.6. Role of FPGA

Processing substantial amounts of data and sending it to multiple users and applications can be challenging. This task can often be daunting to a few servers and hence there are FPGA-based hardware acceleration devices like RAD X Family of Vcinity portfolio. FPGAs (Field-Programmable Gate Arrays) is a solution to this problem as they are widely used in networking. They provide high performance and low power consumption. They can be used to speed up networking functions such as packet processing, encryption, and compression, and also to offload these functions from general-purpose processors, resulting in reduced power consumption and better performance.

The following graphs illustrate the performance of a 1 Gbps link with and without VTP. The graph without VTP shows that the TCP speed starts off high but drops significantly as latency increases. On the other hand, the graph with VTP demonstrates that the speed is sustained even at high latency levels, reaching as high as 1800 ms, which is equivalent to several round trips around the Earth.

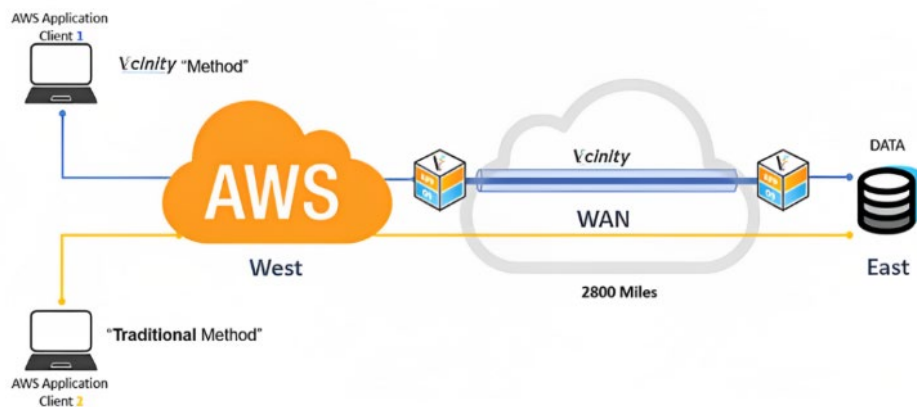


Figure 7: Demonstration setup to compare traditional and Vcinity methods (Ref. 3)

In an experiment, Vcinity used VTP (Virtual Tunneling Protocol) to transfer data from the East Coast to the West Coast of the United States, which typically has a latency of 64 ms for 2800 miles. The results (Figure 8) showed that with VTP, it was possible to send 11TB of data per day over a 1Gbps network, 100TB over a 10 Gbps network and 1PB per day over a 100 Gbps network. This demonstrates that VTP can effectively overcome the latency caused by long distances.

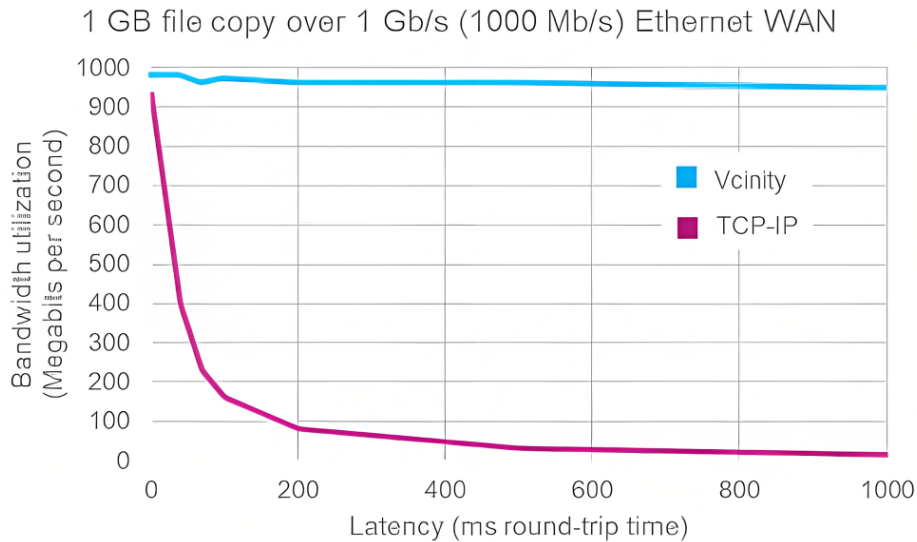


Figure 8: Figure depicting the effect of Latency over TCP and VTP over a 1 GbE link (Ref. 4)

5. Vcinity Data Access Platform

The Vcinity Data Access Platform (VDAP) enables fast data transfer over global wide area networks (WANs) by utilizing various products from the Vcinity portfolio. These products, which can include virtual machines and hardware devices, are tailored to the specific needs of each location, and provide a seamless experience for the user, who can access remote data as if it were local.

5.1. VDAP Application and Working

Depending on the requirement each location gets some part of the Vcinity portfolio. It could be a Virtual Machine (VM) over the existing server, or it can also be other hardware configurations coupled with the server. The underlying architecture of the Vcinity portfolio is designed to be simple and user-friendly, abstracting itself from the mainframe architecture and presenting itself as a network-attached storage (NAS) device. This allows users to access and perform tasks on files and folders located thousands of miles away as if they were locally available.

To access these remote files, the user simply needs to create a folder that points to the local Vcinity platform in their environment. As a result, the user is able to access and perform tasks on remote files without being aware of the underlying architecture and can work with the files as if they were locally available.

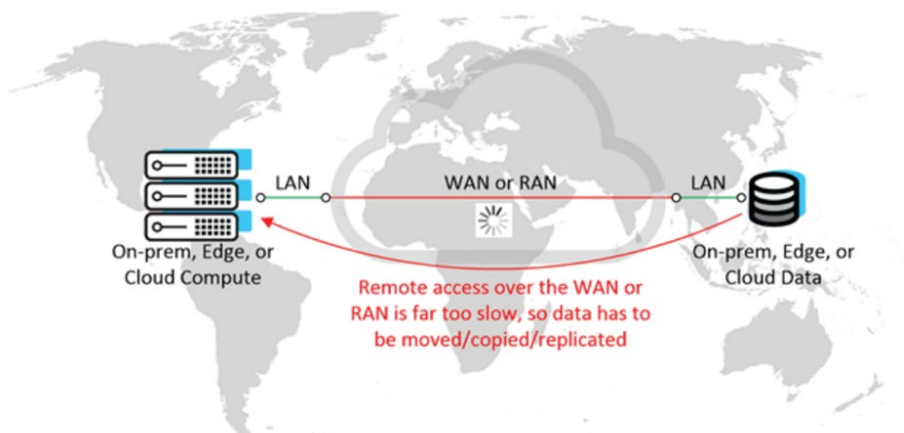


Figure 9: The traditional computational paradigm (Ref. 3)

The underlying architecture creates the virtual lossless tunnel and then the metadata is fetched. The client requests for the data from the server over RDMA hidden under IP headers via WAN. This happens so fast that the application fetching the data is unaware that the data is physically present across the globe while it is logically sitting right beside the compute. This allows the compute to operate on the data wherever it is created and the instant it is created, thereby significantly helping in AI/ML workloads or in cloud bursting. The user interface presents familiar file and directory structures, such as folders and directories, in a simple and easy-to-use format, which users are accustomed to. This allows users to navigate and interact with the system in a way that feels natural to them. At the same time, the complexities of the underlying architecture are hidden from the user, allowing for a seamless and effortless user experience. This abstraction of technical details ensures that users can focus on the task at hand rather than worrying about the inner workings of the system.

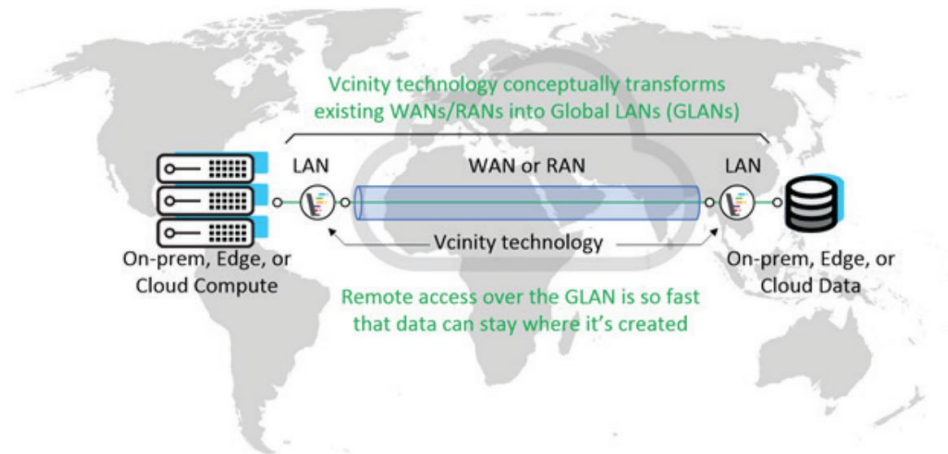


Figure 10: The Vcinity computational paradigm: Making Global data Local (Ref. 3)

One might wonder in case of huge workloads such as 8K video editing at M&E enterprises, how is a 500 GB file sent over RDMA to the application so fast that it does not even notice. That is because whenever the application requests for the data that is being worked upon, it never works on entire dataset by loading it from the hard drive into the memory. Rather it loads a chunk of that data that is being edited and works upon it. Whenever the requirement changes, the new data is loaded into the memory. RDMA directly accesses that particular data without the knowledge of the CPU and directly provides it at the peer side for processing.

6. Security Aspects: Maintaining Data Integrity over Transmission

How do we decide who has access to the existing data?

What if RDMA is misused to manipulate the data without the admin's permission?

Let us understand the security aspects of VDAP -

- The platform is made to allow current permissions to be reapplied on the peer side as well as to take those existing rights to those files and folders from anywhere in the world. The goal is to improve upon and supplement current security measures without adding to their complexity. As a result, the original permissions remain the same and there is no need to change them in any way.
- Data can be encrypted in flight via Advanced Encryption Standard 256-bit (AES 256) symmetric encryption. The other option which is provided by the OEM is Simon Encryption with Fernetix Key exchange.
- Data Prism™ – This technology that utilizes multiple WAN paths to obfuscate the data in multiple physical paths. The number of maximum paths allowed for data prism is 8. In this method, the data is striped across the multiple available WAN paths. This striped data is again encrypted using a unique

encryption key in each link which can be dynamically modified. Hence, even if one of the network paths is compromised and also the unique key is obtained, only a part of the data is lost. As an example, if we have 3 WAN paths available, and 1 of them is compromised, rest 2/3 of the data is still safe and away from harm. This acts as added layer of security.⁵

7. PRODUCT OFFERINGS BY VCINITY – ULTIMATE X® (ULT X) AND RADICAL X (RAD X) FAMILY

Vcinity, Inc. offers a diverse range of products that cater to a variety of hardware and software needs, enabling users to select the solutions that best meet their requirements. Vcinity's Ultimate X® (ULT X) is specifically designed for today's enterprise environment. It allows applications to access remote storage without having to move data, and if necessary, it can move data with unparalleled speed and efficiency. Utilizing its proprietary Remote Direct Memory Access (RDMA) over WAN technology. ULT X is application and server agnostic, as it does not require any changes to existing applications and there is no need to install any software or agents.²

The Radical X® (RAD X) is a Vcinity's product which is available in the form of a PCIe card, which can be easily installed in an X-86 based server. RAD X uses FPGA to perform various real-time tasks, which helps to speed up a large number of single-threaded functions, resulting in high-speed throughput and minimal latency.⁶

In addition to performance, one way in which Vcinity's product offerings can be differentiated is by their ability to fan out to multiple locations. For example, when one of the client locations utilizes VDAP to connect to a main storage facility at a different location for data access, it is known as a 1:1 fan out. In contrast, if a client location connects to three other locations for data access, it is referred to as a 1:3 fan out.

All the offerings support AES – 256 with IPsec as well as protocols namely, NFSv3, NFSv4 and SMBv3. With this concept in mind, let us take a closer look at Vcinity's product offerings.

	ULT X HW (ULT X-1000)	ULT X FPGA (ULT X-1000e)	ULT X SW (ULT X-1000s)	ULT X VM (ULT X-1000v)
Deployment	Hub/Edge	Hub/Edge/Cloud	Edge	Edge/Cloud
Deployment Options	Hardware Appliance	FPGA Edge/Hub: PCIe card Cloud: AWS F1	Converged Instance Bare Metal	Virtual Instance Edge: KVM™, ESXi™ Cloud: AWS™, Azure™
Bandwidth*	>10Gbps	<10Gbps	<5Gbps	<2Gbps
*Performance varies with factors including network configuration				

Table 3: ULT X family offering (Ref. 2)

	RAD X-1010e*	RAD X-1040
Description	PCIe card	Hardware appliance
Fabric Reach	Round-Trip Time (RTT) of 2500ms@10Gbps	Round-Trip Time (RTT) of 2500ms@10Gbps, 625ms@40Gbps
Switching	Not applicable (the RAD X-1010e is an endpoint, not a switch)	120Gbps non-blocking
LAN/WAN Interfaces	Two 1/10GE SFP+ ports	Two 4X FDR InfiniBand QSFP+ ports Two 40GE QSFP+ ports Six 1/10GE SFP+ port
Management Interfaces	Not applicable (the RAD X-1010e is managed through the host server)	One 10/100/1000Base-T Ethernet/ One serial RS-232 console port
Form Factor	Half-height, half-length, single width PCIe card for a server*	1U (1/2 width), 19 in. front rackmount, or two units per 1RU with rear rackmount kit for use with four post racks, or desktop use
*Installed in a server		

Table 4: RADX family offering (Ref. 6)

8. NETAPP'S FLEXCACHE VS VCINITY'S VDAP

Vcinity's FlexCache is a remote caching feature that makes file distribution more efficient, reduces latency on wide area networks (WANs), and decreases WAN bandwidth costs. It allows for distributed product development across different sites, as well as branch office access to corporate datasets. FlexCache is supported on FAS (Fabric Attached Storage), AFF (All Flash FAS), and ONTAP Select running ONTAP 9.5 or later versions. ⁷

This feature helps to increase productivity and collaboration when working with data, while also providing faster read-intensive workloads. By serving frequently accessed data, known as "Hot Data," from multiple controllers in a cluster, FlexCache improves the performance of key applications. Additionally, by caching Hot Data locally for users at various locations, FlexCache enhances collaboration by enabling simultaneous access to centralized datasets, while also reducing the time it takes for users to access Hot Data.

Below is a figure depicting the flex cache implementation on any of their storage offerings on Amazon EC2 running ONTAP 9. ⁷

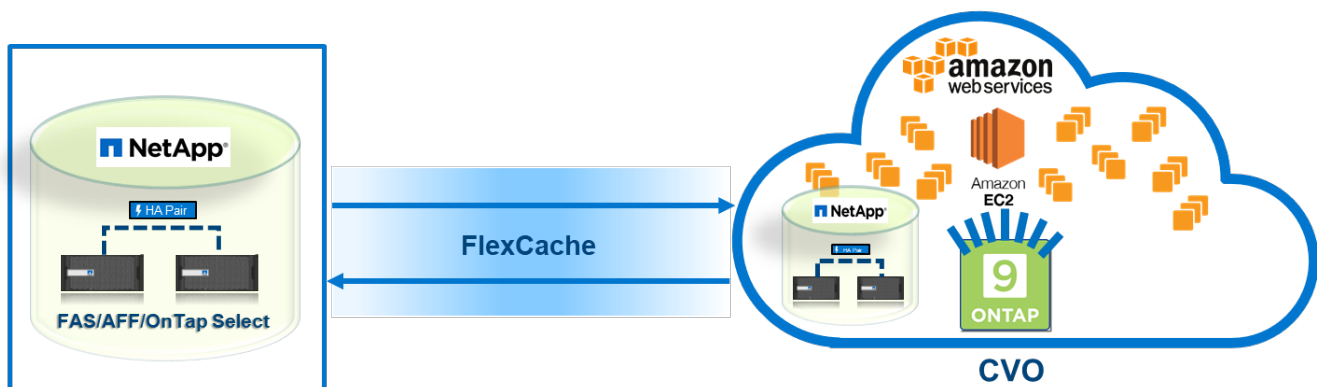


Figure 11: FlexCache implementation

As previously mentioned, Vcinity's VDAP does not employ any caching or network optimization mechanisms. Instead, it is based on the principles of high-performance computing and utilizes various enhancements of Remote Direct Memory Access (RDMA) over wide area networks (WANs) for long-distance communication. Field-programmable gate arrays (FPGAs) are also utilized at certain instances where increased throughput and further processing power is needed. The VDAP concept is built around the idea that data is not always located near where it is computed, and therefore it aims to make data available in real-time, while abstracting

the underlying infrastructure and file system elements to provide a seamless and infrastructure-agnostic experience. This makes all data, whether located locally or in the cloud, easily accessible by converting the WAN into a global LAN and hence the location on the opposite side acts as an extension of local shares.

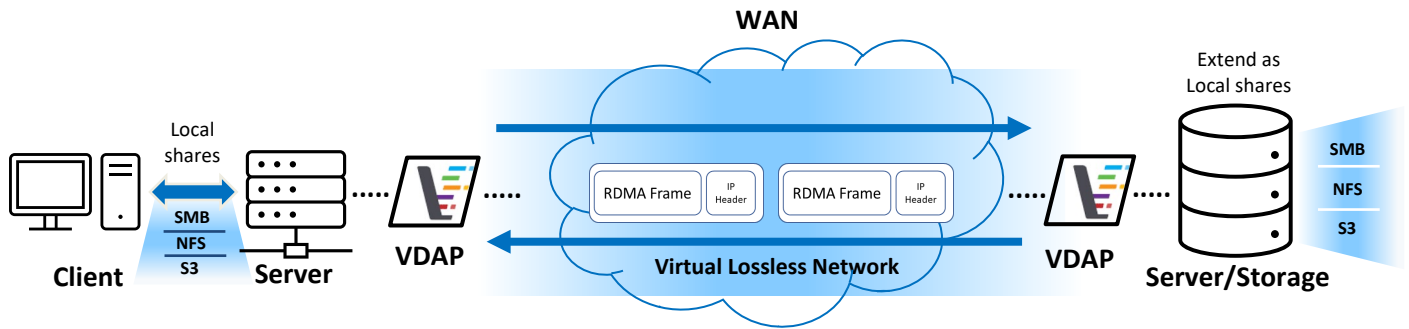


Figure 12: Simple VDAP environment

9. EXPLORING THE EDGE, HYBRID, AND MULTI-CLOUD SOLUTIONS OF VCINITY AND DELL: USE CASES AND BENEFITS

There is a flexible deployment approach that allows for custom configurations to meet the specific needs of our customers. The technology can be utilized to establish connections between servers, clouds, and even between multiple clouds, enabling high-speed data transfer and processing at LAN like speeds over WAN.

9.1 EDGE DEPLOYMENT

Vcinity's Edge deployment solutions address the common challenges of remote data access and slow transfer speeds caused by traditional TCP/IP protocols. These deployments can effectively tackle issues related to data being located at the edge, far away from the main network and difficult to transfer in a timely manner.

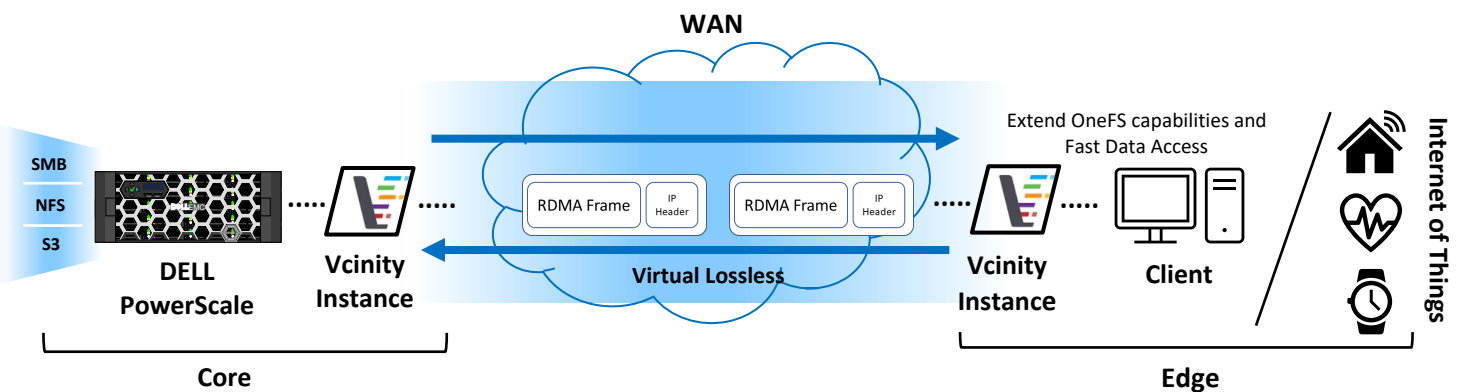


Figure 13: Edge Deployment Model with Vcinity & DELL

9.2 HYBRID CLOUD DEPLOYMENT

With the increasing trend of digital transformation, organizations are under pressure to optimize the use of cloud resources to improve on-prem resource utilization, break down IT silos, and reduce excessive expenditure on maintaining resources for peak usage times. To address these challenges, organizations often utilize pay-per-use models or cloud-bursting techniques.

Hybrid cloud deployments provide a useful solution for cloud-native applications that require real-time access to data stored in on-premises environments, whether at the core or at the edge. These deployments offer a flexible and cost-effective way to manage peak usage times and ensure seamless data access for applications.

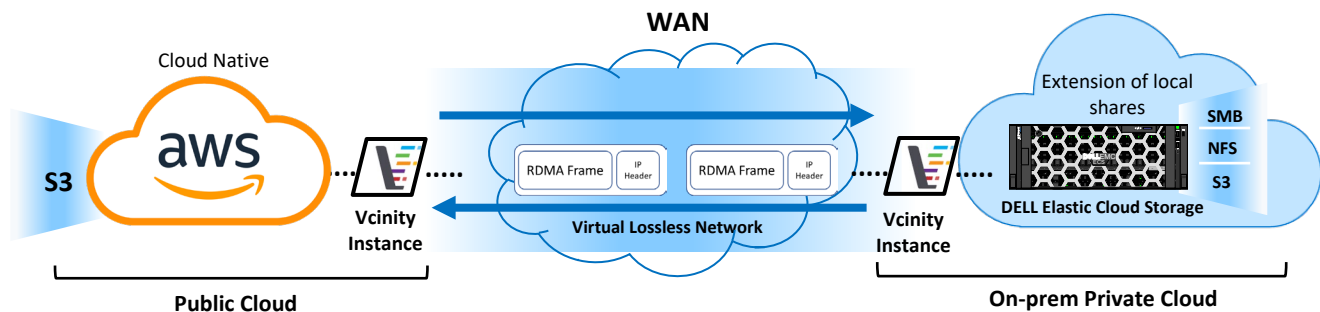


Figure 14: Hybrid Cloud Deployment Model with Vcinity & DELL

9.3 MULTI-CLOUD DEPLOYMENT

Multi-cloud refers to the use of multiple cloud computing services from different providers in a single, unified architecture. This approach allows organizations to take advantage of the different strengths and capabilities of each cloud provider, such as better security, scalability, or specific services.

One of the main use cases for multi-cloud is to mitigate vendor lock-in, which occurs when an organization becomes dependent on a single cloud provider and is unable to easily switch to another provider. It also allows organizations to optimize their workloads by spreading them across different providers, ensuring that they are running on the most cost-effective and appropriate infrastructure for the task at hand.

For instance, organizations may have on-premises resources but also store data in Azure and AWS. Vcinity and DELL's latest technology enables seamless communication and data exchange between these diverse cloud platforms by abstracting the complexities of protocol-to-protocol conversion, making it easy for organizations to benefit from the advantages of multiple cloud providers without getting locked in to a single one.

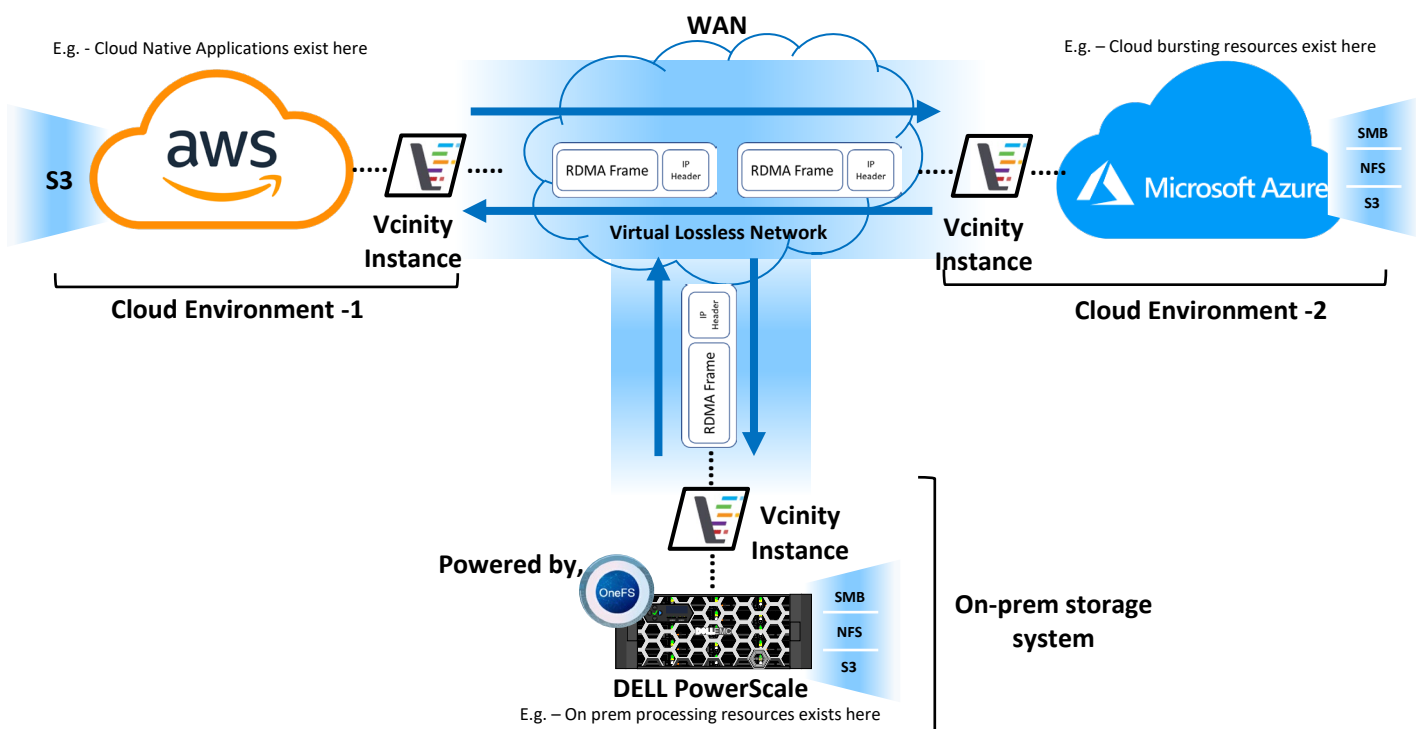


Figure 15: Multi-Cloud Deployment Model with Vcinity & DELL

10. VCINITY & DELL EDGE ACCESS SOLUTION (EAS)

With the aim to create a global namespace (GNS) for all PowerScale and ECS environments of an organization for seamless communication between the resources, and reduce silos, DELL partnership with Vcinity produced the idea to link PowerScale and ECS globally to provide unified and universal access of data. There are several benefits to this – ⁸

1) Extending OneFS to the Edge

Extend OneFS to remote and branch offices to improve collaboration, speed, and productivity for tasks such as software development, file sharing, and animation. Provide a local-like experience when accessing remote data and minimize additional storage costs by caching only accessed data.

2) Extending OneFS to the Cloud

EAS enables you to use the cost benefits of public cloud compute capabilities and persistent storage by extending OneFS to the cloud (e.g., AWS or Azure), while avoiding data lock-in, achieving low total cost of ownership, and maintaining content security through multi-cloud tools.re

3) Rapid Core to Edge access for quickly deriving insights

It allows you to quickly and reliably bring data created at the edge, such as in customer offices, remote sites, or IoT devices, back to the core for immediate use, eliminating the need to wait for edge data.

4) Write caching and write back performance at the edge

EAS provides high-quality, high-performance write-caching and write-back at the edge over existing WAN connections, allowing you to protect your existing network infrastructure investments and reduce future network costs.

5) Data ecosystem with instant data access

You can access and operate on your data from anywhere, without requiring changes to your existing infrastructure. Now data is instantly and securely accessible, giving you the flexibility to deploy operations, workloads, and infrastructure from edge to core to cloud

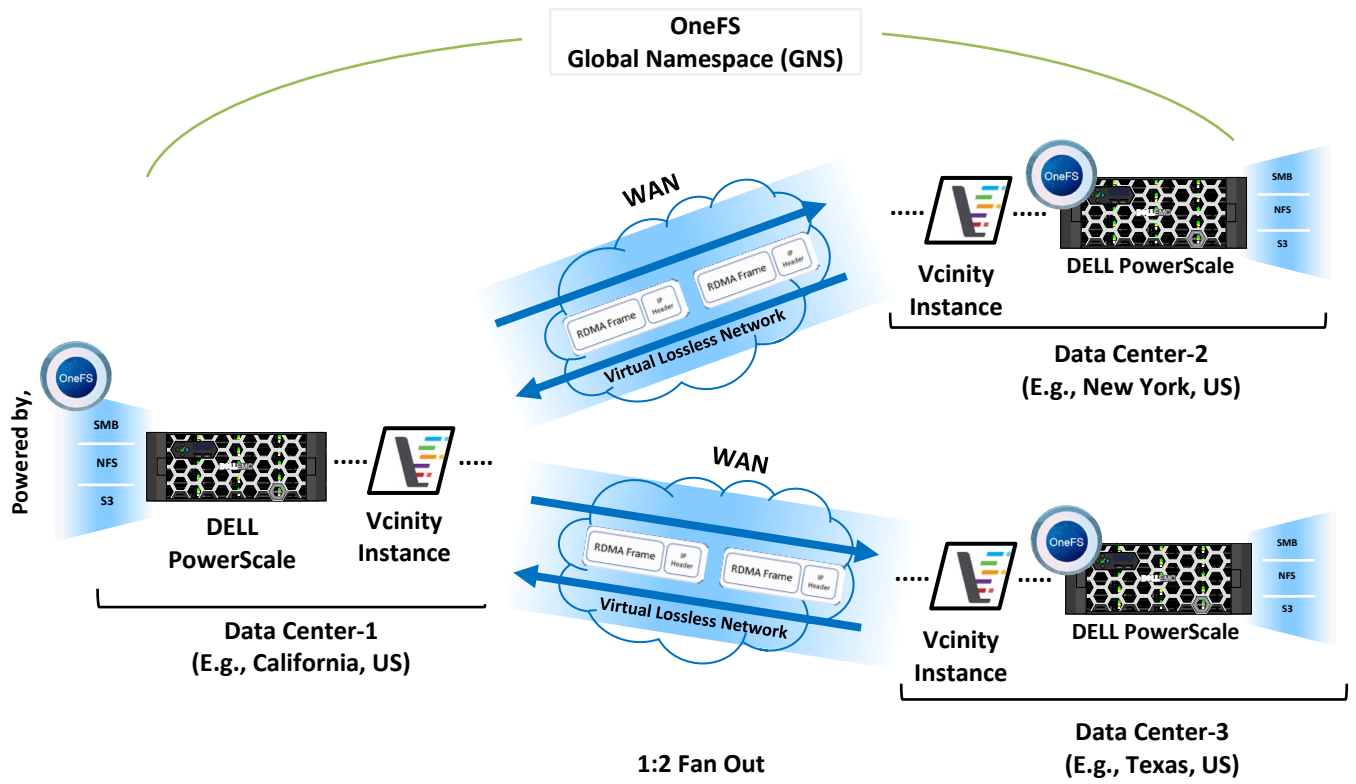


Figure 16: Deployment example for Edge Access Solution

11. Conclusion

With the rapid growth of data, it has become increasingly important for organizations to effectively track, manage, secure, and utilize data while minimizing the high costs associated with storage devices. As data continues to multiply at an unprecedented rate, it is often duplicated and moved multiple times before it is processed. Additionally, the limitations of current wide area network (WAN) standards in terms of data transfer speed and integrity have led organizations to continue copying data to multiple locations in order to meet the demands of users and compute resources. The advent of new technologies such as artificial intelligence and machine learning has further amplified the need to process data in real-time, without the need to move or copy data. Even with the emergence of 5G, traditional solutions such as edge caching, WAN optimization, and extreme file transfer methods, are unable to provide the necessary efficiency or throughput for data transfer over long distances.

Vcinity has addressed these challenges by introducing a new technology to the market, known as the Vcinity Data Access Platform. This technology is built on the foundation of over 30 patents in the field of networking and file systems and utilizes the Vcinity Transmission Protocol (VTP) to provide approximately 90% bandwidth for data transmission, with minimal overhead of only 10%. In comparison, traditional solutions such as TCP or UDP have 90% and 70% overhead respectively, making VTP a much faster and more efficient solution for transferring actual goodput of data in packets.

The Vcinity Data Access Platform offers a multitude of benefits, including the ability to process data in place, from thousands of miles away from the storage location, and the ability to copy data at local area network (LAN) speeds over a WAN. This effectively converts the entire internet into a global LAN, providing data that is so fast it appears to be available locally to the user or compute resource. Additionally, the solution is platform-agnostic, network-agnostic, application-agnostic, and WAN-agnostic, making it easy to implement over existing infrastructure. Data integrity is also maintained during transfer and is protected with AES 256-bit encryption and Vcinity's Data Prizm™ technology. Vcinity offers a variety of deployment mechanisms and, in partnership with DELL Technologies, provides some of the most robust solutions available for making data available on demand with real-time latencies to the user.

12. Bibliography –

1. Crichigno, Jorge & Bou-Harb, Elias & Ghani, Nasir. (2018). A Comprehensive Tutorial on Science DMZ. IEEE Communications Surveys & Tutorials. PP. 10.1109/COMST.2018.2876086
2. At a Glance: Ultimate X – <https://vcinity.io/documents/at-a-glance-ultimate-x/>
3. Intel Application Access Gateway Solution Brief – <https://vcinity.io/documents/application-access-gateway-for-intel-solution-brief/>
4. Solution Brief: Data Movement – <https://vcinity.io/documents/solution-brief-data-movement/>
5. Data Prizm™ Technology Brief – <https://vcinity.io/documents/dataprizm-technology-brief/>
6. At a Glance: Radical X - <https://vcinity.io/documents/at-a-glance-radical-x/>
7. FlexCache in ONTAP: NetApp - Chris Hurley, Product Manager Elliott Ecton, Technical Marketing Engineer - August 2022 | TR-4743
8. Edge Access Solution, powered by Vcinity and Dell PowerScale - <https://vcinity.io/documents/dell-powerscale-vcinity-edge-access-solution/>

Disclaimer: The views, processes or methodologies published in this article are those of the authors. They do not necessarily reflect Dell Technologies' views, processes, or methodologies.

Dell Technologies believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." DELL TECHNOLOGIES MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying and distribution of any Dell Technologies software described in this publication requires an applicable software license.

© 2023 Dell Inc. or its subsidiaries. All Rights Reserved. Dell and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.