

A new paradigm of no move, no copy

The file stays in one place!

Executive Summary

Enterprises will waste billions of dollars by 2025 moving files around as applications and data have to be collocated. Explosive data growth and hybrid, multi-cloud infrastructures exacerbate the problem by copying files everywhere and moving racks and trucks full of data around. The pervasive move-copy paradigm also affects productivity due to the delays in the arrival of data.

Vcinity's revolutionary technology creates a new paradigm of no move, no copy – where the enterprise applications no longer have to be collocated with their data. It transforms your WAN into a Global LAN by enabling instant, secure access to data anywhere across the globe by keeping it in-place without downloading, uploading, copying, pre-caching or pre-staging.

As a result, fresh and complete datasets are accessed where (on-prem, cloud or colo) and when they are created to achieve productivity and efficiency never possible before. Reducing the number of copies of data results in better control and security of valuable data assets and results in substantial reduction in storage, data management and data transfer costs. In cases where the enterprise does not want to commit data to the cloud, it can now leverage cloud compute while keeping data on-prem.

Vcinity allows for data- and application-agnostic performance without using compression, de-duplication or any other processing of data. The integration of our solutions with existing hybrid, multi-cloud infrastructures requires no rip and replace as they support open standard protocols for client and WAN interfaces.

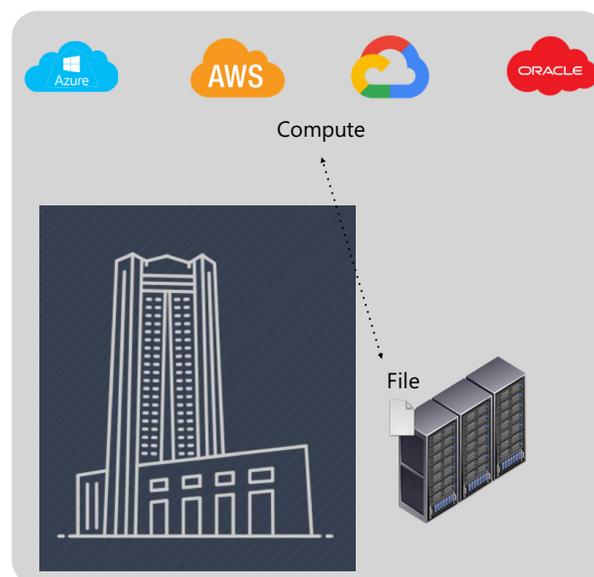


Figure 1. The Vcinity Paradigm—No Copy

Challenges

The current dilemma facing today's IT architects is the increasing gap between the requirements of business applications and the ability to access data needed for them to run successfully. Existing approaches require the application to move to the data site, the data to be transferred to the application site, or data to be accessed from private/public or hybrid cloud. Even with data transfer, data arrives at the destination in an unpredictable time and performance varies with data type/size and application. Utilizing network data transfer optimization appliances and applications are so cumbersome and costly that even physical transportation of storage media or data is considered an acceptable compromise.

Common legacy hardware and software-based methods are deployed to address data movement or pre-positioning of data at the appropriate location before it can be used:

- WAN acceleration or optimizers that provide compression, data de-duplication, and/or protocol overhead reduction
- Proprietary protocols and tunneling to address the inadequacies of the underlying TCP/IP networks
- UDP-based protocols for fast transfers (Extreme File Transfer - EFT solutions) while managing data loss recovery by requiring special software be placed on both ends of the data transfer
- Edge caching solutions that achieve performance by pre-staging or caching a portion of data by predicting usage patterns

These approaches have helped to address data access issues by moving data, and in most cases only focus on the transport part of the overall problem. To fully solve the need for dynamic, agile, and timely data access for modern usage patterns, a solution must address the entire I/O from the application's perspective. This includes how the solution integrates into the architecture, as well as the way data moves through the enterprise and across long distance connectivity.

The current technologies also do not address one of the major issues with copying and moving full data sets: the perishability of data. In some cases, by the time a copy of the data is transferred the information obtained from the data is stale or obsolete. This is especially true with modern AI and ML workflows where instant access to information from the edge is required. Even with the proliferation of edge-based compute having to wait for data to be sent to larger compute farms for true data fusion leaves the edge in a stranded or isolated state.

The Vcinity Solution

With Vcinity's no-copy, no-move paradigm, enterprise applications no longer have to be collocated with their data. Vcinity transforms the WAN into a Global LAN enabling instant, secure access to data anywhere without downloading, uploading, copying, pre-caching or pre-staging. This allows for application and data-agnostic operation to truly leverage enterprise data when and where it is created.

Unlike any current WAN acceleration, file transfer or caching technique (which leverages and enhances existing protocols), Vcinity's solutions address the WAN problem head-on and not just optimize the data being sent. They do not manipulate data preserving its integrity and the ability to

handle any type of file-based data including video or encrypted data. Unlike UDP-based EFT solutions, Vcinity uses a novel approach of leveraging HPC technologies achieving predictably scalable performance regardless of distance and bandwidth and without requiring any special purpose software on client machines.

In addition, Vcinity provides better savings compared to Edge Caching solutions as data does not need to be pre-staged or cached at all. Data is cached or replicated only if the workflow absolutely demands it. This results in a maximum of two copies of data including backup copy required for disaster recovery or continuity of operations. Edge caching solutions' claim of high performance with rapid file access for easy collaboration is true only for the pre-cached data and performance is compromised if data is not in the cache. On the contrary, Vcinity provides predictable performance for accessing any and all of the data regardless of where it may be located within the enterprise infrastructure.

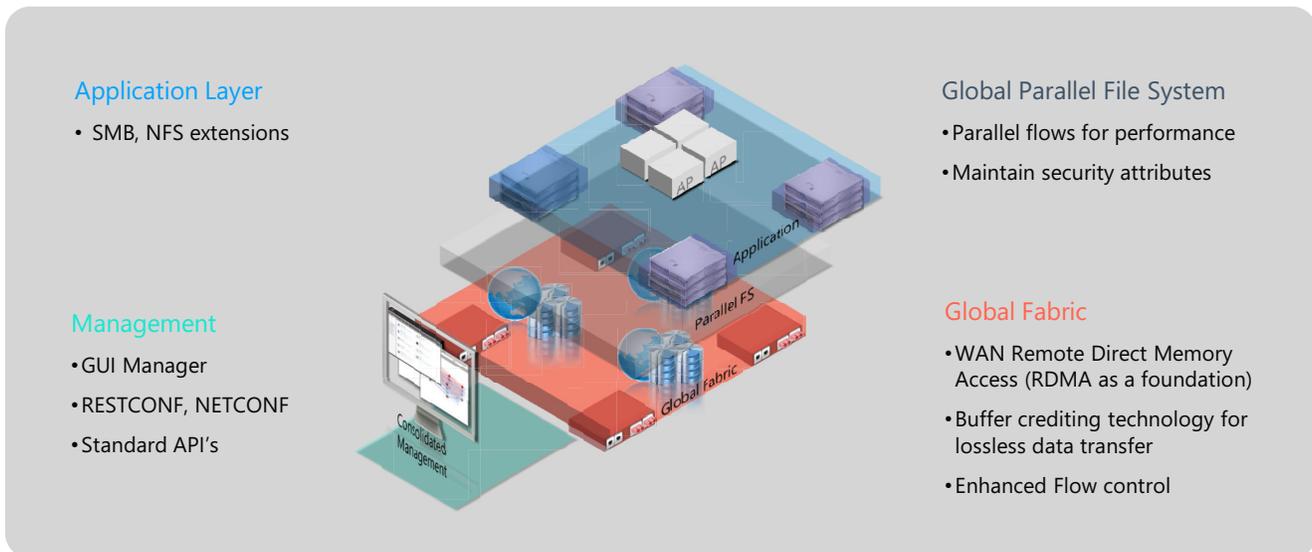


Figure 2. Vcinity's Methodology: Leveraging HPC Technologies to Remove Distance Barriers

Vcinity's unique, holistic approach (outlined in Figure 2) addresses various layers of the overall stack, through efficient transport, integration, and application transparency. This methodology combines proven, patented network processing technologies at both hardware and software levels, with open standards-based approaches to interoperability and implementation. An open API-based interface coupled with powerful data management and synchronization tools offers a solution for enterprise customers to close the IT gap between applications and data and reduces operational costs commonly attributed to closed proprietary systems. The important elements of the Vcinity approach are:

- **Global Fabric**—A highly efficient and secure global fabric by leveraging High Performance Based protocols such as RDMA, advanced traffic engineering and flow control methods across any distance
- **Integrated Parallel File System**—Industry-proven parallel file system designed for high performance I/O of the global fabric, global names space views and distributed seamless access controls

- **Simplified Integration**—Open and well understood interfaces and APIs, easy-to-use tools and multi-site management system for easy integration into existing IT architectures

Global Fabric

Extending TCP/IP and UDP over 10G Ethernet across the MAN/WAN is susceptible to inefficient utilization of the network resources due to congestion, retransmission, and other factors even at distances beyond a typical campus environment. Thus, the performance of interconnects between the locations is diminished by the least efficient element (i.e., the WAN connection). Vcinity's technology transforms the WAN into a Global LAN by extending High Performance Computing technology—Remote Direct Memory Access (RDMA) and storage-based fabrics—over longer distances. These techniques create a lossless and deterministic Global Fabric that stretches across global distances. This provides consistent and repeatable performance—in terms of end-to-end “goodput” with different data sets and under varying traffic conditions. The resulting efficiencies on the MAN/WAN transport reach in excess of 90% “goodput” from WAN entry to WAN exit and are sustained regardless of distance. By creating virtual point-to-point or point-to-multi-point tunnels, traffic engineering and flow control are applied end to end minimizing the risk of errors and dropped packets.

WAN-facing interfaces on Vcinity products leverage standard Ethernet ports that connect to the customer's terrestrial or satellite WAN infrastructure and leverage L2TPv3 tunneling protocol to provide secure transport for customer data. There are no changes required to the existing WAN infrastructure or services. This means that once the WAN ports are connected to the initial WAN Router at each site (for example, a Gateway or Border Router), the WAN path can operate and data can flow across any WAN technology including Ethernet, SONET, OTN, MPLS, ATM and IP networks and any infrastructure services such as Firewalls, WAN Encryption, etc.

Vcinity supports lossless, multi-fabric extension over a WAN connection at distances exceeding 250,000km for 10G or 25,000km for 100G systems. This is done by accelerating the data processing and managing all data in-flight. The approach leverages Vcinity's patented traffic engineering, multiple concurrent flows, multi-protocol support, and in-flight line-rate AES-256 or SIMON-256 encryption.

In addition, DataPrizm™ feature provides significant assurance against compromise of data in-flight over WAN. DataPrizm is similar to Link Aggregation however, each network flow is broken down to the packet level with each sequential packet going over up to eight physical paths and reconstructed on the receiving end. This obfuscates data in-flight across multiple paths such that access to any one path never results in access to more than a portion of the transmitted data.

Integration with Parallel File System

From a performance perspective, a scalable parallel RDMA-enabled filesystem is integrated within the solution to guarantee the highest possible data transfer rates in and out of the Vcinity solution. This technology allows data to be striped across multiple drives within a cluster node and I/O to be performed in parallel, thus increasing the overall data read/writes beyond the serial process of a single drive. By utilizing multiple streams of data, the solution can continuously keep pace with high-speed networks at speeds enabling the maximum possible data throughput.

This file system also enables local error checking and correction in the event of lost data in flight. Unlike standard TCP/IP, in the event of dropped data in flight, the two nodes (one at each end of the WAN circuit) retransmit the missing information without local sender or remote target intervention. In

traditional TCP/IP-based systems, a missing or dropped packet requires the originating sender to retransmit the missing data plus the data in flight once again. Doing this with an inefficient protocol such as TCP/IP or a connectionless protocol like UDP across the WAN can cause a significant impact on performance. As an example, having to retransmit lost information on a 10G network with a latency of 40ms can result in retransmitting over 50MB of data. With Vcinity, terminating TCP/IP locally rather than extending it over the WAN allows for enhanced error recovery and higher throughput particularly at geo-scale distances. By adding advanced technologies such as Packet Loss Recovery, retransmission is completely avoided by creating a construct to rebuild lost information on the receiving end rather than retransmission of lost information.

Simplified Integration

It is imperative for an enterprise solution to integrate into an existing architecture easily, and to use industry standard protocols for interoperability and longevity. Ensuring a seamless integration for presenting shared data, Vcinity leverages standard folders, shares, and network path features. By incorporating SMB and NFS mount points, users can access data over distance by simply mounting a folder or NFS export. This is shown in Figure 3 with sites able to access data at other sites by accessing corresponding folders as if there were local to the site. Placing or retrieving content into well understood file shares using these methods is a non-disruptive mechanism for accessing data. To transfer data, the user simply moves the data into the folder associated with the local node and the solution presents the data in a similar folder at the remote end. No additional licenses are required for the clients to access the information. In fact, the users and applications need no knowledge of the underlying system; it appears to local users like a NAS access by NFS and SMB.

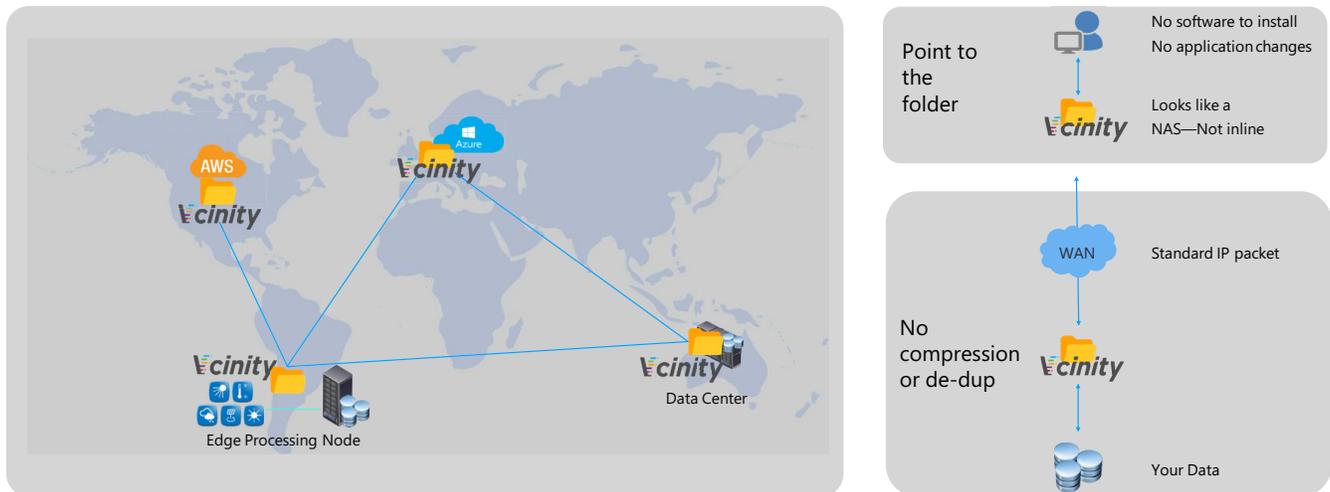


Figure 3. Easy to Use: Seamless and Transparent

This well-understood method allows the solution to be dropped into an existing architecture with minimal up-front effort and does not require the users to have any training or specific proprietary knowledge of the solution or its components. The system includes a management console that may be reached via any internet browser for obtaining status, statistics, and events.

Enabling Global Federated Data

With the techniques outlined above, Vcinity enables a federated data platform that is crucial for today's data-driven workflows with compelling benefits:

- Unparalleled Accessibility—Access data remotely by keeping it in-place without moving it
- No copy sprawl—No keeping track of redundant copies, keeping them current or securing them consistently across different locations. This results in substantial infrastructure cost savings and better control/security of the data overall.
- Transparent to applications—Seamlessly integrates and presents as standard storage types. No changes to the applications.
- Transparent to data type—No compression, deduplication and maintains data integrity
- Flexible, hybrid IT Environment—On-premise, multi-cloud or true hybrid environments. Scalable cloud-based compute instances against data on-premise or compute on-premise against data stored in the cloud, or any mix.

When the use case (such as backup, disaster recovery) demands data replication, Vcinity enables movement of terabytes to petabytes of data at unmatched speeds over any distance. We have proven to transfer 100 Terabytes in less than a day across the U.S. over 10 Gbps link and 1 Petabyte in a similar timeframe over 100 Gbps link demonstrating the linear scalability and predictability of our architecture. Such performance enables the customers to operate at the highest levels of link efficiency and deploy various use cases such as data migration (to another data center, colo site, or the cloud), backup and disaster recovery achieve the most aggressive goals that scale with data growth.

Conclusion

The challenge for the enterprise is in capitalizing on the value of its data, which is dispersed across an ever-growing complex infrastructure spanning geographically- separated, hybrid—data center, colo, and cloud environments. This data continues to grow at an alarming rate, making the challenge of universal access a daunting task especially with the AI/ML models requiring instant access to information. Through proven technology, storage and networking expertise, Vcinity brings to market the ability to access data without downloading, copying, moving, uploading, pre-caching or pre-staging it. By providing a standards-based solution Vcinity truly unlocks the full potential of data anywhere for any enterprise across any geography.



Some features listed in the specifications may be under development. ©Vcinity, Inc. 2020. All Rights Reserved. Vcinity, Inc., the Vcinity logo, Radical X, Ultimate X, Command X, Access X, DataPrizm, and Ultimate Access are trademarks and/or registered trademarks of Vcinity, Inc. Any other trademarks are the property of their respective owners. Doc ID: 85-0200-001 Rev. B 11/23/20