

Transferring 1 Petabyte Across the U.S in Just Over 23 Hours

1PB/23Hours@100Gbps / ~2,800 Miles

EXECUTIVE SUMMARY

Challenge:

- Transfer 1 Petabyte (PB) of file sets across a single 100Gbps WAN connection with 70ms latency
- Complete data movement in under 24 hours
- Perform the transfer using commonly available tools (i.e., Linux "cp")

Solution:

Vcinity's ULT X software suite and RAD X-1040 provide a compelling package for high performance servers to transfer or remotely access any data, anytime, anywhere. The integrated solution seamlessly fits into existing enterprise LAN and WAN environments and unleashes the true value of distributed data across global distances.


Results:

The test continued for just over 23 hours to transfer a total of 1.0556PB of data. The integrated solution is not optimized specifically for file transfer applications, yet it achieves optimal, predictable and repeatable performance leveraging existing resources.

Vcinity™ in collaboration with Penguin Computing® recently showcased the unprecedented ability to migrate 1 Petabyte (PB) of data across the U.S. in just over 23 hours. The data was transferred using a 100Gbps WAN link across approximately 2,800 miles. The demonstration integrated Vcinity's technologies with those of Penguin Computing using generally available hardware/software and industry-standard protocols.

Test Configuration

The following figure illustrates the test configuration.



Configuration	Datasets
<p>Six Penguin Computing Relion XE1112 source servers:</p> <ul style="list-style-type: none"> • Dual Intel Xeon Gold 6150 with 192GB DRAM, 8x1.6TB NVMe drives • Running Vcinity Ultimate X™ (ULT X™) software <p>Six Penguin Computing IceBreaker 4936 target servers:</p> <ul style="list-style-type: none"> • Dual Intel Xeon with 128GB DRAM, 24x1.4TB SSD drives • Running Vcinity ULT X software <p>Six Vcinity Radical X™ (RAD X-1040) fabric extenders</p>	<p>Total data transferred: 1.00556 PBs</p> <p>Mixed filesets: 100x10GB, 20x50GB, 20x100GB, 10x500GB, and 5x1TB = -12.6TB per run per server pair</p> <p>2 server pairs — 13 runs transferred 164.01TB per server pair: 13 runs * -12.6TB per run per server pair</p> <p>4 server pairs — 14 runs transferred 169.38TB per server pair: 13 runs * -12.6TB per run per server pair + last run (14th run) transferred 5.368TB per server pair</p> <p>Starting after the second run, data from the previous run was deleted from the target directory after each server pair finished a run. This process was done due to storage limitations on the target directories and continued until all transfers per server pair were completed.</p> <p>The test was performed 3 times over 3 days.</p>

The test configuration consisted of six pairs of Penguin Computing’s high-performance servers running Vcinity’s ULT X software suite and RAD X-1040 WAN Fabric Extenders connected to a 100G Ethernet switch. Each RAD X-1040 WAN connection was rate-shaped at a maximum bandwidth of 16.666Gbps feeding the aggregate 100Gbps WAN connection with 70 milliseconds of latency, typical of a dedicated circuit between the U.S. East and West Coasts. Each pair of two servers was connected to a pair of RAD X-1040 fabric extenders on either side of the WAN connection via InfiniBand connections. Using file synchronization capabilities of Access X™, the basic test procedure repeatedly sent datasets of mixed files and sizes totaling ~12.6TB per run from the source server bank to the target server bank for a total of 13 runs. Then the last run on four servers used a reduced fileset (so as not to greatly exceed an aggregate of 1 PB).

Results

Results (refer to the table below for details) include:

- 1.00556PB of data transferred in three different tests on three separate days in just over 23 hours with times of 23:19:19, 23:15:48 and 23:13:28 hours for each run averaging to 23:16:12 hours.
- Approximately 12.6TB data sets averaged 1:44:17 hours across all three runs and all servers.
- Four servers with fastest transfer times averaged at 1:43:10 hours with a standard deviation of 9.065 seconds.

- Average bandwidth utilization across all three runs and all servers was 96.9%. Four servers with the fastest transfer times averaged 97.8% utilization with a standard deviation of just 0.1%.

Summary

The PB test results illustrate clear advantages of using Vcinity’s technologies:

- **Unprecedented performance:** While not specifically optimized for file transfer applications, the combined solution still achieves unprecedented performance for transferring large volumes of data across the WAN.
- **Performance predictability:** The file transfer time variation is minimal across multiple runs and server pairs confirming the solution’s deterministic performance. Running the test three times over three days resulted in very repeatable performance of moving a PB in just over 23 hours.
- **Linear scalability:** From 100TB in 23:05:46 hours over 10G to 1PB in an average of 23:16:12 hours over 100G, Vcinity’s technology scales its performance linearly with larger data sets and network links.
- **Ease of integration:** Vcinity’s technologies work with Commercial Off-The-Shelf (COTS), generally available hardware/software and industry-standard protocols instead of highly customized solutions. They also interoperate seamlessly with existing LAN/WAN infrastructures.

Server Transfer Times

	Across All Servers		Across Four Servers with Fastest Times		Across Two Servers with Longest Times	
	Transfer Time	Utilization	Transfer Time	Utilization	Transfer Time	Utilization
Average Time	1:44:17	96.86%	1:43:09	97.84%	1:46:33	94.73%
Fastest Time	1:43:02	97.98%	1:43:02	97.98%	1:43:29	97.53%
Longest Time	1:50:43	91.16%	1:43:58	97.08%	1:50:43	91.16%
Difference	0:07:41	6.82%	0:00:56	0.90%	0:07:14	6.37%
Standard Deviation	0:01:49	1.62%	0:00:07	0.12%	0:01:28	1.32%



Some features listed in the specifications may be under development. ©Vcinity, Inc. 2019. All Rights Reserved. Vcinity, Inc., the Vcinity logo, Radical X, Ultimate X, Command X, Access X, and Ultimate Access are trademarks and/or registered trademarks of Vcinity, Inc. Any other trademarks are the property of their respective owners. Doc ID: 20-0200-007 Rev. C 08/12/19