PURESTORAGE®

TECHNICAL WHITE PAPER

Modernizing Midrange Data Protection with LaserVault and FlashBlade

Accelerate recovery outcomes and remove data silos.

Contents

Introduction

Midrange servers such as the IBM AS/400—though largely considered IT “dinosaurs”—are critical systems in many infrastructures. They operate in a majority of the Global 2000 and are here to stay for the foreseeable future. Protecting these systems and their critical data against disaster has been a challenge at a time when data centers are looking to eliminate single-use solutions and legacy big iron, and reduce CapEx and OpEx by shrinking their hardware footprints. Even with these other goals, disaster protection remains a priority. As a result, many data centers operate multiple and complicated backup solutions to protect both modern and legacy hardware and applications.

Midrange systems have relied on replication to other systems or on large and complex tape operations that can require offsiting and increase storage costs. Tape solutions are prone to mechanical failures, and restoration from tape after an outage can be time consuming and prone to errors and failures. These solutions also tend to leave data siloed, with no straightforward way of migrating data or applications to modern types of servers, storage, or cloud.

For organizations looking to protect these valuable systems and data, reduce their data center footprints, and ensure fast recoveries, LaserVault and Pure Storage® have created a modern data protection solution specifically for midrange servers.

Introducing LaserVault ViTL

ViTL is a Fibre Channel or SAS-connected all-flash appliance that appears as a tape library to your IBM Power Systems host (AS/400, iSeries, System i, AIX, Linux), automating the backup and recovery process. It emulates an IBM TS3200 tape library (3573-L4U) with extra cartridges. You can replace your existing tape library with ViTL and continue to back up and restore without programming changes. ViTL works with BRMS, Robot Save, and other backup software. Virtual tapes appear as tape media to the IBM i or other host and to BRMS or Robot Save. You get more options for more drives and more slots than tape without adding additional hardware. ViTL can also be used as a gateway device to store IBM i backups on Pure Storage FlashBlade®.

Introducing Pure Storage FlashBlade

Pure Storage FlashBlade is a Unified Fast File and Object (UFFO) storage platform that serves as the storage foundation for modern applications and data services.

FlashBlade consolidates workloads, eliminate data silos, and simplify infrastructure by consolidating diverse file and object workloads onto a single storage platform, all while delivering better performance for legacy and modern data and application demands.

FlashBlade capabilities include:

- **Multi-dimensional performance**: High throughput and IOPS with low latency support multiple workloads simultaneously, including those with small or large files, sequential or random I/O access, batched or real-time jobs, and large numbers of files.
- **Intelligent architecture**: The storage system is built for flash from the ground up to truly leverage the performance and efficiencies of flash. It is simple to deploy, manage, and upgrade without requiring constant tuning. A modern storage solution must be simple enough so that the operations don't overwhelm storage admins with mundane tasks like managing networking complexities when deploying the system, volumes, cluster pairs, aggregates, and flash caches or configuring replication.
- **Cloud-ready**: With on-premises control, you gain cloud-like agility, flexibility, and consumption choices.
- **Always available**: FlashBlade goes beyond traditional platform resiliency. Maintenance operations, software upgrades, and capacity expansions are completed without disruption. The foundational software design makes it possible for the solution to deliver high availability over multiple years and upgrade scenarios.
- **Dynamic scalability**: With FlashBlade, you can seamlessly scale not only capacity but also performance, metadata, number of files and objects, and more.
- **Multi-protocol support**: You gain a single platform provides native file and native object protocol support without compromising performance or any functionality.

Solution Architecture

The solution consists of four major components:

- IBM Power systems host
- ViTL appliance
- BEdupe appliance (can be standalone or run on the LaserVault appliance)
- Pure Storage FlashBlade

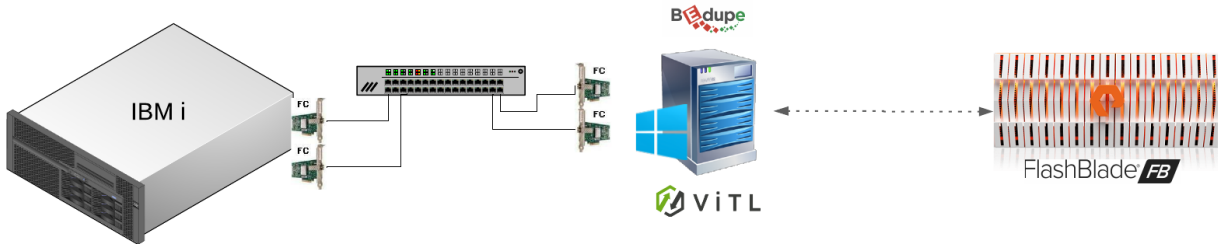

Figure 1: Solution architecture

ViTL Appliance

An all-flash appliance which appears to your IBM Power Systems host (AS/400, iSeries, System i, AIX, Linux) as a tape library. The ViTL is connected via Fibre Channel switch to your IBM Power System. The LaserVault appliance can replicate backups at the file level (i.e., tapes) to other storage systems, including cloud-based targets for archival purposes. In the joint FlashBlade-

LaserVault solution the ViTL appliance will write directly over an NFS or SMB connection to FlashBlade or the BEdupe appliance (see below).

Deduplication

Backups are notoriously storage inefficient with redundant data causing unnecessary storage growth. BEdupe can either be a standalone Windows all-flash appliance, or it can be run the LaserVault Appliance. It works by comparing blocks of data to detect duplicates. Once BEdupe finds a block of data that already exists in the deduplication repository, the block is replaced with a virtual index pointer linking the new pseudo-block to the existing block of data. When you restore a file, BEdupe recreates the file from its reparse point file, using its own special index and unique chunk store.

Testing Background

Best practices in data protection call for having multiple copies of your data stored in at least two different places to ensure recoverability. In this scenario, we backed up an IBM Power system host (AS/400, iSeries, System i, AIX, Linux) to a LaserVault ViTL virtual tape appliance, and then created an additional backup of the ViTL appliance with a Pure Storage FlashBlade. This ensures that there are multiple backups available for recovery and adds the capability to replicate offsite, faster recoverability, and data reuse.

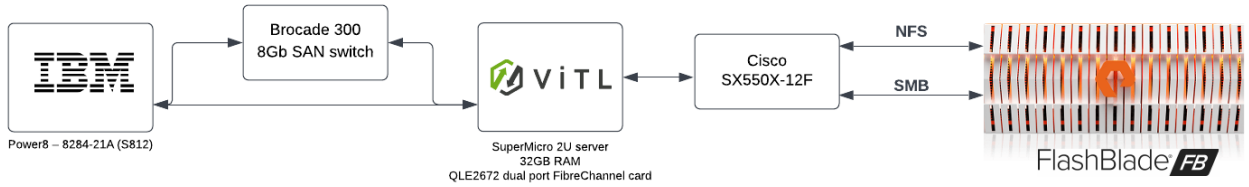

Figure 2: Test setup

Test Systems

Multiple test systems were used. See below for details.

- LaserVault Server:
  - SuperMicro 2U server with dual 4215R Silver processors 3.2 to 4.0 GHz
  - 32GB RAM
  - Intel X-520 dual port 10Gbe Ethernet card
  - Both 10Gbe ports teamed and using LACP; QLE2672 dual port Fibre Channel card with (2) 16Gb FC directly connected to IBM i system (see 16Gb tests)
  - LaserVault ViTL version: 2.13.3
  - O/S Windows server 2019

- 1 FlashBlade 7x17TB
  - Purity//3.2.0
  - 2 x 10GbE, port-channel
  - No vPC due to single switch
- 1 Cisco SX550X-12F
  - 12-port 10GbE switch
- 1 8Gb Brocade 300 8Gb SAN switch
  - 1 8Gb between IBM i and LaserVault server (see 8Gb tests)
- 1 IBM server
  - Power8 – 8284-21A (S812)
  - Spin disks
  - EN0A/577F 16Gb Fibre Channel HBA
  - O/S: IBM i V7R3M0

Test Details

Below you will find more information about the tests and how we conducted them.

Two separate file systems were created for testing. The file system named “Data0” had only NFSv3 enabled while the file system named “Data1” had only Native SMB 2.1 enabled. UNC paths were used for both NFS and SMB (e.g., no drive mapping). Native SMB was enabled by Pure Support and subsequently added to the Active Directory installation.

Multiple tests were run with various ViTL configurations. Differing numbers of data files, buffer sizes, data sizes and paths were used to find the best combination. Given the large number of testing permutations possible, it was not practical to test all of them. Only a subset was run (as noted in the tables below) and narrowed in on based on initial test run results.

We ran only a single backup or restore from the IBM i machine at a time. A single file was presented as a single tape cartridge to the IBM i machine. Backups were run from the IBM i using a single tape/file (standard/serial) as well as multiple tape drives and cartridges/files (parallel).

The 16Gb tests were with the IBM i directly connected to the LaserVault server and the 8Gb tests were with the IBM i connected to the LaserVault server via a Brocade SAN switch. ViTL software used two buffers of equal size. One buffer was filled with incoming data, while the other was being flushed to the file on the FlashBlade for a backup.

Failure tests were performed by starting a backup job and then physically pulling various blades. Blades were pulled both randomly and specifically by identifying the actual blade that terminated the TCP connection. The blades were pulled and reinserted after 30-45 seconds, causing session disconnects on all protocols.

In the testing we did not include BEDupe.

LaserVault Replication

LaserVault software can replicate backups at the file level (i.e., tapes) to other storage systems, including cloud-based targets for archival purposes.

For backups with LaserVault compression enabled, data files are transferred in their compressed (smaller) size, reducing time and resources. For backups with LaserVault compression disabled, data files will be transferred in their full (larger) size as the data would be uncompressed as it is delivered to the application (in this case LaserVault) by the FlashBlade. This may result in greater time and resources to complete a transfer.

We strongly recommended disabling LaserVault compression, as FlashBlade compression is always-on and cannot be disabled. Testing showed overall data reduction with FlashBlade compression yielded data saving of 3.6 to 1.

FlashBlade Replication

Data files can also be replicated at the FlashBlade array-level and referred to by another instance of the LaserVault ViTL software. However, FlashBlade native replication sends the data uncompressed, eliminating the data reduction benefits over the wire.

Performance Summary

For both NFS and SMB, the best test results using real customer data was achieved with the combination of:

- 1 mount path
- 8 data files
- 8MB buffer sizes

FlashBlade allows for fast backups and restores. With the recommended settings and the documented lab hardware, backup, and restore speeds of greater than 500MB/sec (1.8TB/hr.) speeds were easily achievable. Results suggest that even faster backup and restore speeds are possible in environments with greater FC connectivity when running parallel backups from IBM i.

The minimally configured 7x17TB FlashBlade supports up to ~2.3GB/sec writes, and 7GB/sec reads in aggregate when all blades are utilized. There is additional headroom for more backup streams from additional clients with this configuration.

NFS							
Test Method	Data	Data Size	Average Write Speed (Bytes/s)	Average Read Speed (Bytes/s)	Data Files	Buffer Size (MB)	Paths
Save From IBM i (16Gb)	Highly Compressible Test Data	160,332,963,840	210,165,073	257,108,246	1	8	1
Save From IBM i (16Gb)	Highly Compressible Test Data	160,332,963,840	563,591,459	473,312,908	4	8	1
Save From IBM i (16Gb)	Highly Compressible Test Data	160,332,963,840	575,540,394	467,994,785	8	8	1
Save From IBM i (16Gb)	Highly Compressible Test Data	160,332,963,840	567,518,925	467,994,785	4	8	4
Save From IBM i (16Gb)	Real Customer Data	51,680,116,736	151,111,452	213,554,201	1	8	1
Save From IBM i (16Gb)	Real Customer Data	51,680,116,736	549,788,476	506,667,811	4	8	1
Save From IBM i (16Gb)	Real Customer Data	51,680,116,736	587,274,054	492,191,588	8	8	1
Save From IBM i (16Gb)	Real Customer Data	51,680,116,736	434,286,695	353,973,402	4	8	4
Save From IBM i (16Gb)	Real Customer Data	51,680,116,736	427,108,403	441,710,399	4	16	4
Save From IBM i (16Gb)	Real Customer Data	51,680,116,736	544,001,229	474,129,511	4	16	1
Save From IBM i (16Gb)	Real Customer Data	51,680,116,736	527,348,130	478,519,599	8	16	1
Save From IBM I (8Gb)	Real Customer Data	51,680,116,736	461,429,614	353,973,402	7	8	1
Save From IBM I (8Gb)	Real Customer Data	51,680,116,736	496,924,199	346,846,421	7	8	7
Save From IBM i Parallel (8Gb)	Real Customer Data	103,360,233,472	760,001,716	571,051,014	8	8	1
Save From IBM i Parallel (8Gb)	Real Customer Data	103,360,233,472	658,345,436	538,334,549	4	8	4
Save From IBM i Parallel (8Gb)	Real Customer Data	103,360,233,472	760,001,717	560,551,531	7	8	7
TapeUtil Disk Read Test	Real Customer Data	203,522,004,896		908,439,558	7	8	1
TapeUtil Disk Read/Write Test	Random Data	214,748,364,800	560,700,691	1,106,950,334	7	8	1
TapeUtil Disk Read/Write Test	Random Data	214,748,364,800	307,222,267	499,414,801	7	8	7
TapeUtil Disk Read/Write Test	Repeating Data (0xFF)	214,748,364,800	989,623,800	1,136,234,734	7	8	1
TapeUtil Disk Read/Write Test	Repeating Data (0xFF)	214,748,364,800	466,844,271	516,222,030	7	8	7

Figure 3: NFSv3 using Windows NFS Client

SMB							
Test Method	Data	Data Size	Average Write Speed (Bytes/s)	Average Read Speed (Bytes/s)	Data Files	Buffer Size (MB)	Paths
Save From IBM i (16Gb)	Real Customer Data	51,680,116,736	587,274,054	474,129,511	7	8	1
Save From IBM i (16Gb)	Highly Compressible Test Data	160,332,963,840	596,622,460	473,312,908	7	8	1
Save From IBM I (8Gb)	Real Customer Data	51,680,116,736	506,667,811	400,621,060	7	8	1
Save From IBM i Parallel (8Gb)	Real Customer Data	103,360,233,472	771,345,026	574,223,519	7	8	1
TapeUtil Disk Read Test	Real Customer Data	203,522,006,392		1,169,485,408	7	8	1
TapeUtil Disk Read/Write Test	Random Data	214,748,364,800	778,073,785	1,167,110,678	7	8	1
TapeUtil Disk Read/Write Test	Repeating Data (0xFF)	214,748,364,800	1,106,950,334	1,173,488,332	7	8	1

Figure 4: SMB 2.1 using Windows SMB client

Conclusion

PureStorage FlashBlade and LaserVault offer a compelling solution for organizations that want to remove tape backup and recovery and associated costs and complexity, transform recovery speed, and consolidate other backup and recovery solutions, as well as any other applications which require a high performance, scale out storage architecture.

Appendix

Failure Testing Results

NFSv3

Here we can see the original TCP source port 985 connected terminated at blade 1. NFS uses port 2049.

```
ir@FB01-ch1-fm1:~$ exec.py -na "sudo netstat -antp | grep -v LISTEN | grep 10.0.0.20"
[ir1.out] tcp        0      0 10.0.0.14:2049          10.0.0.20:985           ESTABLISHED 5249/nfs
[ir2.rc]  exit 1
[ir3.out] tcp        0      0 10.0.0.14:2049          10.0.0.20:997           ESTABLISHED 5190/nfs
[ir3.out] tcp        0      0 10.0.0.14:445           10.0.0.20:63085         ESTABLISHED 5190/nfs
[ir4.rc]  exit 1
[ir5.rc]  exit 1
[ir6.rc]  exit 1
[ir7.rc]  exit 1
```

After the blade was pulled, we can see the NFS client reconnected with new source port 960 to blade 7:

```
ir@FB01-ch1-fm1:~$ exec.py -na "sudo netstat -antp | grep -v LISTEN | grep 10.0.0.20"
[ir1.err] ssh: connect to host ir1 port 22: No route to host
[ir1.rc]  exit 255
[ir2.rc]  exit 1
[ir3.out] tcp        0      0 10.0.0.14:2049          10.0.0.20:997           ESTABLISHED 5190/nfs
[ir3.out] tcp        0      0 10.0.0.14:445           10.0.0.20:63085         ESTABLISHED 5190/nfs
[ir4.rc]  exit 1
[ir5.rc]  exit 1
[ir6.rc]  exit 1
[ir7.out] tcp        0    328 10.0.0.14:2049          10.0.0.20:960           ESTABLISHED 5220/nfs
ir@FB01-ch1-fm1:~$
```

During NFS failure testing, the LaserVault application did not report any errors and completed without issue. There were no errors in the application logs; LaserVault was completely unaware of the blade failure.

The application continued without issue when the actual blade that terminated the TCP connection was pulled as well as when a random blade was pulled.

SMB 2.1

Here we see the original TCP source port 58639 (verified on the Windows Server) connected to blade 6. SMB uses port 445.

```
ir@FB01-ch1-fm1:~$ exec.py -na "sudo netstat -antp | grep -v LISTEN | grep 10.0.0.20"
[ir1.out] tcp        0      0 10.0.0.14:2049          10.0.0.20:985           ESTABLISHED 5249/nfs
[ir2.rc]  exit 1
[ir3.out] tcp        0      0 10.0.0.14:2049          10.0.0.20:997           ESTABLISHED 5190/nfs
[ir4.rc]  exit 1
[ir5.rc]  exit 1
[ir6.out] tcp        0      0 10.0.0.14:445           10.0.0.20:58639         ESTABLISHED 5299/nfs
[ir7.out] tcp        0      0 10.0.0.14:2049          10.0.0.20:950           ESTABLISHED 5220/nfs
ir@FB01-ch1-fm1:~$
```

After blade 6 was pulled, we can see the SMB client reconnect with new source port 63085 to blade 3:

```
ir@FB01-ch1-fm1:~$ exec.py -na "sudo netstat -antp | grep -v LISTEN | grep 10.0.0.20"
[ir1.out] tcp        0      0 10.0.0.14:2049          10.0.0.20:985           ESTABLISHED 5249/nfs
[ir2.rc]  exit 1
[ir3.out] tcp        0      0 10.0.0.14:2049          10.0.0.20:997           ESTABLISHED 5190/nfs
[ir3.out] tcp        0      0 10.0.0.14:445           10.0.0.20:63085         ESTABLISHED 5190/nfs
[ir4.rc]  exit 1
[ir5.rc]  exit 1
[ir7.out] tcp        0      0 10.0.0.14:2049          10.0.0.20:950           ESTABLISHED 5220/nfs
ir@FB01-ch1-fm1:~$
```

When the blade was pulled during the backup using SMB, the LaserVault application threw an IO “media error” and the backup stopped completely. The application hung waiting for a reconnection to the data file and needed to be manually stopped and restarted. These same results occurred when the actual blade that terminated the TCP connection was pulled, as well as when a random blade was pulled.

Due to the backup being stopped, the data already written to the data file on FlashBlade is incomplete. In this scenario the user would be notified (via email) that the backup job failed and would need to restart the backup. The new backup would overwrite the data on the existing data file or select a new tape cartridge (data file). There is no risk of data corruption as the failed backup would be useless and abandoned.

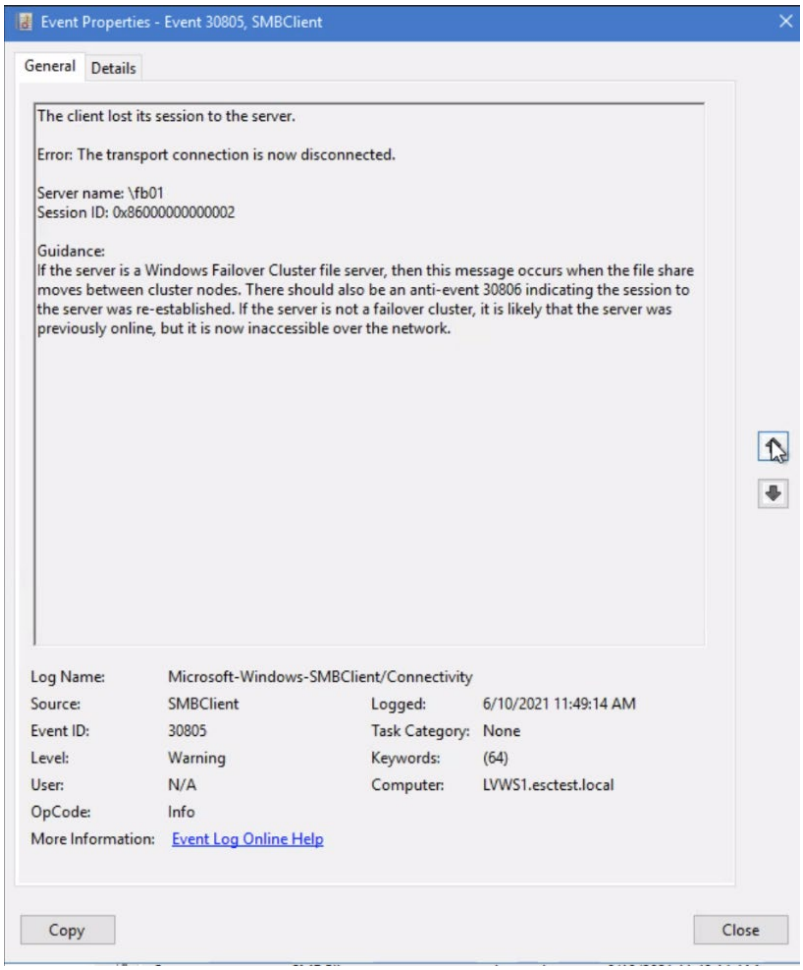

Figure 6: SMB Client event log error. There were no 30806 anti-events in the logs indicating a recovery.

NFS Data Validation Test

Another test was performed to validate the backup data after a blade failure when connected via NFS. The steps were:

1. Run backup from IBM i to internal disk in LaserVault server.
2. Data file is copied with Windows Explorer from local/internal disk to FlashBlade via NFS.
3. A blade is pulled during the copy process and a short drop in throughput is observed (as expected) and the copy resumes.
4. Perform binary comparison of the two backup files using `fc.exe /b`.

This test showed that the local backup file and the FlashBlade backup file were an exact binary match, proving the blade failures have no impact on job completion or potential for data corruption.

Data Reduction

LaserVault compression was disabled for all tests. FlashBlade compression is always-on and cannot be disabled.

The FlashBlade compression ratio on the customer daily/weekly saves is 3.6 to 1. For comparison, LaserVault's compression ratio on the same data is 4.29 to 1. (Taken from different backup with same dataset). Actual real-world compression will vary based on the data type in other environments.

Additional Resources

Next Steps

- Try FlashBlade for free and take a test drive.
- Want to see LaserVault in action? Sign up for a demo here.

Supporting Information

- Learn more about FlashBlade Unified Fast File and Object Storage.
- Pure Storage is Positioned as a Leader in the 2021 Gartner® Magic Quadrant™ for Distributed File Systems and Object Storage

Pure Storage, Inc.
650 Castro Street, #400
Mountain View, CA 94041