

Why Windows Is a Good Platform for Hosting Virtual Storage Software

Only One Right Choice

The resiliency, performance and cost-effectiveness of applications are governed in large part by the platform on which they are deployed. That holds particularly true for large-scale storage control software – arguably one of the most mission critical applications under IT’s purview.

DataCore’s design team, composed chiefly of Unix and proprietary real-time operating system designers, deliberated numerous tradeoffs when selecting the hardware and software platform for the SANsymphony and SANmelody products. They ultimately broke with past allegiances to Unix and Linux, choosing instead Windows on Intel/AMD PC servers as their foundation.

This white paper examines the rationale behind that calculated decision and validates the choice with undeniable evidence of commercial success.

Improving on Conventional Storage Arrays

Most conventional storage arrays, whether monolithic or distributed, share a widespread problem – they are prone to a common type of undetectable programming errors. Essentially, their designs have no way to prevent one firmware program from unintentionally overwriting memory addresses that affect other modules. In software engineering lingo this condition is referred to as stray pointers. Stray pointers account for many data corruption incidents encountered in high-end storage controllers. Some vendors go to great lengths to hide this dirty little secret. Their purchasing contracts prevent customers from publicly discussing occurrences of data corruption.

The likelihood of many such errors occurring somewhere in millions of lines of firmware code is shockingly high. What’s worse, no amount of testing can ensure they’ve been discovered. Usually, the culprit is never found even after much irreparable damage has been done to the customer’s data.

Windows’ Memory Management Catches Hidden Programming Errors

DataCore, on the other hand, takes advantage of Windows’ built-in memory segmentation to prevent programming errors from addressing space outside their range. Any such errors are immediately detected and stopped by the operating system. The origin can be traced and corrected to prevent it from causing any damage. This is a key benefit of having a robust commercial OS acting as an untiring traffic cop.

In addition, DataCore engineers can take advantage of Windows’ modern development tools and high-level languages to substantially reduce the lines of code manually written. These tools dramatically reduce the incidence and consequences of programming errors.

As a result of the above, SANsymphony and SANmelody software can be statistically proven to be more reliable than the firmware so many trust implicitly.

Unix & Linux Lack Critical Features for Storage Control

Some are curious why DataCore picked Microsoft Windows over one of the many Unix or Linux variants.

Given that all of these commercial operating systems offer similar memory management features and development tools, using a Linux offshoot, for example, would potentially be cheaper.

On the face of it, Unix and Linux were the natural first choices for SANsymphony. DataCore's principal architect, one of the company's founders, is referenced in the Unix literature as an authority on multi-threading. DataCore's design team is composed of real-time Unix kernel designers who earned their stripes writing proprietary operating systems for high-fidelity military flight simulators and nuclear power plant controllers.

However, the team's first-hand experiences trying to adapt multi-processing Unix and its derivatives to closely coupled storage subsystems had earlier revealed a key flaw.

-
- ▶ Commercial Unix and Linux operating systems lack deterministic interrupt processing.
-

In other words, standard distributions of Unix and Linux cannot ensure that a given task will complete in a predictable and consistent window – a critical failing for real-time processes necessary in enterprise-class storage control.

While this is not an issue for time-sharing applications measured in minutes and seconds, it is absolutely imperative when trying to process more than 400,000 IOPS on a pair of redundant PC servers.

The lack of determinism explains why storage appliances running on off-the-shelf Unix or Linux kernels suffer from poor and erratic performance. Despite their reliance on considerably more expensive hardware, these Unix/Linux products are unable to achieve SANsymphony's and SANmelody's extraordinary response and throughput.

Real-time Adaptations Violate 3rd Party Qualifications

In an effort to reduce the performance handicap of their Unix or Linux foundation, some companies make proprietary modifications to the standard operating system. These so-called “kernel hacks” bring significant but seldom-publicized consequences:

First, the vendor must accept responsibility for maintaining and supporting its proprietary Unix or Linux variant – which eliminates any benefits from standardization. Though the suppliers might claim to use a particular base OS, their custom modifications invalidate all software and hardware qualifications against the original OS version.

These suppliers find themselves weighed down with the monumental task of re-qualifying hardware and software add-ons. This forces them to devote a large number of people and numerous person-years to retesting. Invariably, the added burden delays their time to market, increases their R&D costs and prevents them from incorporating bug fixes and security patches available from the latest commercial updates.

Their proprietary Unix or Linux version rapidly becomes obsolete, as does their choice of hardware.

Windows Offers DataCore Customers Determinism & Compatibility

Windows on the other hand, provides deterministic interrupt processing in its micro-kernel. This unique property allows DataCore to plug in its low-latency, feature-rich, I/O stack without modifying the base operating system.

By prioritizing SANsymphony and SANmelody's execution ahead of background services, polling I/O channels for inputs, and using on-board RAM as inexpensive high-speed cache, we eliminate unpredictable I/O processing delays found on competitor's implementations.

-
- ▶ These traits enable DataCore's Virtual Storage Software to achieve ten times more throughput than other products running on equivalently equipped PC-servers.
-

As importantly, the clean Windows application programming interfaces (APIs) make it practical to develop SANsymphony and SANmelody modules that maintain hardware and software compatibility with the complete spectrum of 3rd party qualified devices and applications. Strictly adhering to the Windows conventions ensures DataCore's software stays current with new technology without any extra effort, fully leveraging the development and qualifications of other best-of-breed suppliers.

Precautions Ensure a Tightly Controlled Robust Environment

While no one can guarantee that operating systems won't fail, DataCore and its customers employ several measures to minimize the occurrence of such incidents.

Some actions are under DataCore's software control, while others fall under the realm of IT best practices. All take the perspective that SANsymphony and SANmelody nodes are special function mission-critical engines, not general-purpose Windows hosts.

For example, SANsymphony and SANmelody both create a controlled and compartmented environment by executing as a well-behaved I/O stack layered on top of the Windows micro-kernel. The core software reserves processor, memory and I/O resources for the exclusive purpose of providing dependable target disk emulation and value-added block I/O services. Essentially, the same loop of code runs repetitively as a virtual machine inside the OS shell. Windows merely serves as a convenient host for the administrative GUI, as well as the conduit for standard IP networking services between collaborating SANsymphony and SANmelody nodes.

Standard operating procedures dictate that no user applications or unqualified third-party device drivers are permitted on the SANsymphony and SANmelody servers. Avoiding unknown interactions between variables eliminates the root cause behind most OS panics. This precaution alone sidesteps numerous instability issues seen in the general-purpose use of Windows and other popular operating systems.

Security and Virus Protection

DataCore customers also employ prudent, security-conscious best practices to physically isolate their Windows-based virtual storage control nodes from virus assaults and other security threats.

Just like any mission critical high-end disk controller, customers take practical steps to keep their SANsymphony and SANmelody nodes away from untrusted Internet or e-mail services. ¹

¹ Contact DataCore's service and support organization for more explicit security recommendations.

Of course, DataCore storage servers may be serving disk blocks over Fibre Channel or iSCSI to web servers and e-mail servers subject to virus attacks. In these cases, the pooled disk drives could potentially house a virus or worm. To shield against contamination and other potential security concerns, SANsymphony and SANmelody prevent the underlying OS from accessing the disk drives or any executable programs stored there.

Network Design: More Resilient Than Its Parts

While we've focused on the rationale for selecting Windows throughout this paper, DataCore's end-to-end design criteria goes much further:

-
- ▶ No single component of the storage infrastructure can possibly be trusted to be bulletproof.
-

Failure to acknowledge this simple fact would be foolish.

Storage networks are made up of many parts constantly in flux: some failing, new ones being added, still others being decommissioned. Turmoil is the norm, not the exception. Therefore, comprehensive networking solutions, like SANsymphony and SANmelody, account for equipment that goes down unexpectedly or intentionally for periodic software updates, hardware upgrades and maintenance.

The SANsymphony and SANmelody storage architecture is designed to insulate applications from this chaos, removing all single points of failure and continually providing access to data despite planned and unplanned outages.

Network of Loosely Coupled Redundant Nodes

DataCore's solutions make use of path, platform and block level replication to eliminate all single points of failure. SANsymphony's and SANmelody's asynchronous network design is borrowed from the telecommunications industry, where it has been successfully applied to large-scale, highly distributed voice and data networks for many years. Redundant paths from applications through the infrastructure (FC or iSCSI) ensure that critical data is accessible over alternate routes should the primary connection be disrupted for any reason.

Redundant Windows-based storage control servers take over for each other in the event of an outage. Yet the nodes do not incur the synchronization overhead or the scalability constraints necessary to keep server clusters in lock step.² Our loosely coupled approach is by definition more highly available, less complex, and much more economical than trying to bolster the reliability of any one box or attempting to closely tie a pair of them.

Nodal independence and loose inter-node cooperation form the foundation for SANsymphony and SANmelody to achieve high-availability, unmatched performance and optimum scalability at the lowest cost.

A long list of satisfied customers using DataCore's solutions in enterprise and SMB environments can attest to these benefits.

² Although SANsymphony and SANmelody do not rely on clustering to achieve fail-over, they do provide enhanced storage services to clustered applications running in Windows, Unix, Linux and NetWare.

Other Distinguishing Attributes

Like many successful Independent Software Vendors (ISVs), DataCore finds Windows attractive on several other fronts.

Many of the operating system's characteristics translate into a fabulous value proposition for both our end-users and our integration / OEM partners.

As a proven commercial environment that ships in huge volumes, the Windows platform provides rapid time-to-market and numerous price point advantages. Its familiar look-and-feel makes it easy for storage administrators – even those emotionally attached to Unix derivatives, to quickly be effective with little training.

Perhaps most significantly, the benefits of the Windows platform are enduring:

- Windows is first to market with support for all the latest, best of breed and relatively inexpensive hardware.
- It's not locked into a proprietary hardware platform. DataCore's customers get to choose from a wide range of server configurations available from numerous trusted suppliers.
- The large installed base persuades vendors to qualify their storage subsystems, host-bus adapters, switches and network interface cards on Windows first, freeing DataCore from such cost and responsibility while allowing us to concentrate on innovating within our core competencies.
- There is a wealth of 3rd party software to further enhance manageability.

- Windows servers consistently offer the best price-performance, while the massive consumer market ensures that they sustain that lead.
- Even if DataCore ceased development today, SANsymphony and SANmelody customers would enjoy faster and cheaper storage control just by non-disruptively migrating their software to next year's more powerful and more cost-effective Windows-powered PC platforms.

These are all lasting benefits that purpose-built appliances can never offer.

From an IT operations and business perspective, SANsymphony and SANmelody leverages Windows for competitive advantage by offering:

- Freedom to choose from a variety of hardware suppliers during acquisition and upgrades (no vendor lock-in)
- Successive generations of higher speed lower cost equipment from which to consolidate and automate their storage infrastructure

Reduced Total Cost of Ownership (TCO) arising from inexpensive off-the-shelf storage control nodes.

And Windows Outruns Everything Else

It may surprise a few of you to find that just a pair of economical 2-way Windows PC servers, "Powered by DataCore," combined with popular JBODs (just-a-bunch-of-disks), outperform top of the line EMC, HDS and IBM storage arrays - at just a fraction of the big guys' price!

Furthermore, in many storage configurations that use various high end arrays, customers have experienced substantial performance gains and simpler overall administration by front-ending them or their mid-range cousins with SANsymphony and SANmelody virtual storage solutions.

DataCore's virtual storage solutions achieve well over 400,000 I/Os per second (IOPS) and in excess of 2.1 GB/sec of throughput on just two inexpensive dual-processor storage servers. That's 3 to 5 times better than proprietary arrays. It's also considerably more powerful than Unix and Linux-based schemes.

With SANsymphony's N+1 scalable architecture throughput and IOPS scale linearly as SANsymphony nodes are added to expand the pool for growing workloads.

Best Overall on SPC-1™ Benchmark

The Storage Performance Council's SPC-1 benchmark provides another point of reference.

Just three similarly configured SANsymphony storage domain servers were first to drive 50,000 SPC-1 IOPS. DataCore established new records at twice the price/performance of distributed arrays and 5 times better than the fastest monolithic subsystems.

Evaluator Group has gone so far as to create a new category called "Open Platform Storage Systems" to account for DataCore's architectural and price performance advantages.

Peace of Mind

At the end of the day, personal experiences and those of trusted colleagues often provide the most comfortable basis for decision-making. DataCore can present an assortment of Global 2000 accounts that run DataCore software around the clock on redundant networks of Windows PC servers to consolidate and control very large heterogeneous storage pools.

These customers span such vertical sectors as healthcare, financial services, government, manufacturing and telecommunications with mission critical applications hosted on Unix, Linux, Netware, VMware and Windows.

Their successful implementations as well as the insight we've provided in this paper show that the selection of Windows for the SANsymphony and SANmelody platform was purposeful and carefully designed to yield the highest availability as well as the most benefits for the end-user. SANsymphony's and SANmelody's unique storage consolidation, management automation, data protection and performance capabilities fully leverage the technical, operational and business advantages that Windows offers.

*For more information, contact
DataCore Software Corporation
at www.datacore.com or email us at
info@datacore.com.*



Headquarters

6300 NW 5th Way

Fort Lauderdale, FL 33309

T 954.377.6000

Or toll free 877.780.5111

F 954.938.7953

E info@datacore.com

www.datacore.com

© Copyright 2007, DataCore Software Corporation. All rights reserved. DataCore, the DataCore logo and SANsymphony and SANmelody are trademarks of DataCore Software Corporation. All other products, services and company names mentioned herein may be trademarks of their respective owners. P052AC