



Buyer Case Study

AMD Drives Performance and Manageability by Moving Virtual Infrastructure to Tintri Arrays

Eric Burgener

IDC OPINION

3rd Platform computing environments, driven as they are by the needs of both older applications (Oracle, SQL, Exchange, etc.) and newer applications (mobile computing, social media, big data and analytics, and cloud), demand significant performance, scalability, and agility. As enterprises consolidate applications onto virtual infrastructure, they are also dealing with very different I/O workloads at the storage array level. These requirements are all converging to drive the need for a new storage architecture based around flash technologies and designed for virtual environments. The rapid business growth of Advanced Micro Devices (AMD), a leading supplier in the microprocessor industry, had begun to stress the company's legacy storage infrastructure. Performance, complexity, and cost issues associated with AMD's legacy midrange storage array were causing problems for the company's most mission-critical environment – a manufacturing application that directly impacted factory throughput. The very visible need to meet service-level agreements (SLAs) with the manufacturing group drove IT to explore newer, flash-optimized storage architectures. An ability to cost effectively meet service-level performance agreements, ease of use, and data path visibility (for troubleshooting purposes) were the key factors in AMD's decision to deploy Tintri's VMstore hybrid flash arrays (HFAs). Further:

- Although the initial deployment was just for the mission-critical manufacturing application, AMD's experience with Tintri's solution was such that AMD moved more workloads onto that platform, purchased well over 100TB of storage capacity on VMstore HFAs, and is actively planning to move all of its virtual infrastructure onto these flash-based arrays.
- With Tintri's HFAs, AMD has implemented an architecture that can continue to meet performance SLAs as its business grows, provided end-to-end visibility across the data path to simplify troubleshooting, and removed the need to manually tune storage for performance reasons.

IN THIS BUYER CASE STUDY

This IDC Buyer Case Study summarizes how AMD brought flash storage technology in the form of HFAs into its datacenter for improved performance, easier manageability, and better troubleshooting visibility. This Buyer Case Study explores what drove AMD's initial interest in flash, how AMD has deployed the technology in its environment, and what AMD's experience with flash has been.

SITUATION OVERVIEW

Organization Overview

AMD is a \$5 billion computational technology company that was originally incorporated in Sunnyvale, California, in 1969. In addition to its PC and server microprocessor products, AMD's portfolio also includes embedded processors and other chips for graphics, video and multimedia applications, memory, and solid state drives (SSDs). AMD's internal IT infrastructure is built around VMware's vSphere for the company's most mission-critical applications and split across two independent datacenters for high availability. AMD manages four different datacenter sites, with application "pairs" load balanced across the sites.

AMD's most important manufacturing application leverages Oracle databases and is used to drive the company's manufacturing process. Rather than replicating between two datacenter locations, production data was batch loaded into the Oracle instances at each site several times a day. Socket server front ends can attach to either location to perform queries, providing a measure of fault tolerance for end users. This application requires 24 hours-a-day uptime, and the IT group manages it according to an SLA that demands a very low average query time on a daily basis. The storage for this environment had been based on a scale-up, monolithic storage architecture design, but as AMD's business grew, the company was outrunning the capabilities of this legacy platform.

Challenges and Solution

Several issues about the legacy platform were becoming problematic. Batch loading was taking too long, and Ross Alaspa, the Member of Technical Staff product engineer responsible for performance, was having difficulty consistently meeting his SLAs for the key manufacturing application. Query performance directly impacted the throughput of the factories, so storage performance issues had top-line visibility with the CIO. Storage performance troubleshooting was also a very time-consuming and manual task. Alaspa didn't have enough visibility into what was going on in the storage array to diagnose and fix problems directly in the production environment: the problem had to be recreated in the development environment to perform troubleshooting, and performance issues were increasing as AMD's business grew.

While attending a VMware User Group meeting in Austin, Texas, in 2012, Alaspa discovered Tintri. Tintri sells HFAs that ensure flash performance for 99% of reads and writes, deliver submillisecond latencies, and include a graphical user interface (GUI) that provides real-time, on-screen visibility of IOPS, latencies, capacity consumption, and flash hit ratios at a virtual machine (VM) – not a LUN – level. Within the Tintri VMstore arrays, the Tintri operating system (OS) allows quality of service (QoS) parameters to be established at the VM level, and then autonomously manages array resources to meet performance requirements without manual intervention. Based on these two capabilities, Alaspa brought a Tintri VMstore T540 in for a proof-of-concept (POC) test.

Tintri had the T540 set up on AMD's development network, with VMs mounted, within 30 minutes. Alaspa has a performance suite that he uses to test new software releases, and AMD used VMware's Storage vMotion solution to move that suite over to the T540 and began testing. The impact of the flash performance was immediately evident. Features like the Tintri HFA's ability to lock selected virtual disks into flash allowed AMD to ensure flash performance for critical elements like redo logs 100% of the time. Compared with AMD's older array, initial installation and configuration was much easier, and the system's ability to self-manage to meet SLAs without operator intervention promised to

improve not only the reliability of operations but also the administrative span of control. Alaspa deployed two T540s into production in late 2012 for the manufacturing application.

Results

The Tintri HFA deployments had several major impacts. Batch loading now happens in near real time, giving analysts up-to-date information to work with. Query completion times are now well within specified SLAs, and there is significant room for growth while continuing to meet these requirements with the Tintri architecture. Administrative productivity has improved. Each site has a single administrator responsible for managing the entire virtual environment where the manufacturing application is running, and these personnel are not classic storage administrators by background or training. Outside of increasing capacity when needed and establishing SLAs, almost no time is now spent on managing storage performance.

"The visibility on storage performance, allowing us an end-to-end view from the guest OS layer to the storage layer, lets us pinpoint exactly the source of a problem – host, network, storage, or something else – and take appropriate remedial action," said Alaspa. "This makes it easier than ever to meet our SLAs, and we basically don't have to manage storage performance issues anymore."

Tintri arrays also provide the level of availability Alaspa needs for his most mission-critical application. Tintri's VMstore solutions feature dual-controller architectures, redundant power and cooling, hot pluggable components, and flash-optimized RAID 6 implementation and can support nondisruptive controllers and drive firmware upgrades. Although AMD is not using replication today, Tintri supports a very bandwidth-efficient and space-efficient snapshot-based replication capability (ReplicateVM) that is also well integrated with VMware Site Recovery Manager. Since Alaspa went live with the Tintri arrays in 2012, he has not had any downtime despite experiencing a controller failure in 2014.

Although many enterprises may initially experiment with flash technology on less critical applications, Alaspa deployed it in AMD's most mission-critical environment first. As AMD has grown since the initial deployments in 2012, Alaspa has purchased additional Tintri VMstore solutions and is moving other workloads to these as well. He brought in six additional VMstore T620s and has a total of 108TB of raw flash capacity under management for his manufacturing applications. Alaspa recently brought in a VMstore T820 to provide the storage for an engineering cluster as well and is taking advantage of some of the other features of Tintri's solutions here, using both the space-efficient, almost instantaneously provisioning clones and the inline data reduction. Using both compression and deduplication, Alaspa is seeing, on average, a 6:1 data reduction ratio in the engineering environment. Since Alaspa runs a lot of Oracle databases, he is particularly happy with the inline compression capability that he is using instead of Oracle's columnar compression. "Oracle's compression is extremely CPU intensive, driving up the need to license Oracle across more cores," noted Alaspa. "With Tintri's compression, we can offload that to the array, drop some cores, and save on Oracle licensing costs."

Alaspa also has a fledgling virtual desktop infrastructure (VDI) environment that AMD is still running on its legacy array. Although AMD has less than 50 virtual desktops under management today, that has already become cost prohibitive from a storage perspective, and plans are in the works to move that environment to Tintri products. "Given what we've seen with Tintri, we want to move all our virtual environments onto that platform," explained Alaspa. "Tintri products are purpose built for virtual environments, and the VM-level storage management is not only more efficient and intuitive than the LUN-level management we're forced to use on our older array, but it frees us from having to deal with storage performance issues on a regular basis like we were before."

In closing, Alaspa noted that Tintri has worked closely with him on road map items that allow him to manage his environment more effectively. Although Tintri supports AES-256 encryption, AMD's engineering group was also looking for additional security capabilities to help them make new builds available more selectively and securely. A Tintri OS release in late 2014 included the export controls that Alaspa was looking for. "That's a level of responsiveness that I just don't see from the larger enterprise storage vendors," said Alaspa. "Tintri got with us on this request and was able to roll that feature out very quickly."

ESSENTIAL GUIDANCE

The server consolidation enabled by virtual infrastructure has significantly changed I/O workloads at the host level. 3rd Platform computing data streams exhibit extremely random I/O patterns with a wide variety of read-to-write ratios, a wide distribution of different block sizes, a high percentage of data that is reducible, and I/O bands (hotspots) that drift over time. Flash-based arrays consistently deliver submillisecond latencies and IOPS in the hundreds of thousands and enable the use of inline data reduction for primary storage environments. In IDC's opinion, flash is required to cost effectively meet performance requirements in 3rd Platform computing environments and should be on the radar of every datacenter looking to refresh legacy storage technologies.

The LUN-based management of legacy storage architectures does not fit well with the VM-level management semantics of the virtual world. Administrators need to be able to perform storage operations – snapshots, clones, replication, and so forth – and see storage metrics (IOPS, latency, throughput, flash hit ratios, etc.) at a VM level. Tintri has offered this capability since its first products shipped in 2012, but this feature is permeating newer storage architectures in the hyperconverged space as well. And the most popular hypervisor platform (VMware vSphere) is offering an API that will allow storage to be managed at the VM level even for legacy arrays (with some work on the part of legacy array vendors). IDC sees VM-level management as the wave of the future, not only to improve the efficiency of storage operations but to make storage management more intuitive for the IT generalists who are increasingly managing storage in virtual environments.

AMD's flash deployment model follows the approach of much of the industry to date. Flash-based arrays are often brought in to meet requirements for a particular application, but the customer experience has been so positive that enterprises are eager to move more applications onto these platforms. As customers move from dedicated applications to more mixed, virtual workload-type deployments, other features besides performance start to become critical. In particular, enterprises want snapshots, clones, inline data reduction, storage tiering, encryption, QoS, and replication in a package that delivers five-nines or greater uptime. As flash capacity starts to be deployed at scale, enterprises will start to see a huge cost benefit from flash's secondary economic benefits: far fewer devices needed to meet performance requirements, significantly reduced energy and floor space consumption, far fewer servers needed to meet performance requirements, and lower software licensing costs (due to the need for fewer servers). The total cost of ownership (TCO) of flash-based systems that are roughly 100TB in size or larger already can be as much as 50% lower than that of legacy hard disk drive-based storage arrays.

Autonomous, self-managing storage is a significant differentiator for those vendors that offer it today, but this feature has such value that it will likely become another baseline feature requirement in enterprise storage solutions within the next 12-18 months as storage management tasks migrate more and more to IT generalists.

LEARN MORE

Related Research

- *Tintri Arrays Drive Business Improvements for Cloud Services Provider Cirrity* (IDC #254414, February 2015)
- *Worldwide All-Flash Array and Hybrid Flash Array 2014-2018 Forecast and 1H14 Vendor Shares* (IDC #252304, November 2014)
- *IDC's Worldwide Flash Storage Solutions in the Datacenter Taxonomy, 2014* (IDC #250560, September 2014)
- *Executive Interviews on Flash Array Deployments in the Enterprise* (IDC #249884, July 2014)
- *Flash-Optimized Storage Architectures* (IDC #249295, June 2014)

About IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications and consumer technology markets. IDC helps IT professionals, business executives, and the investment community make fact-based decisions on technology purchases and business strategy. More than 1,100 IDC analysts provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries worldwide. For 50 years, IDC has provided strategic insights to help our clients achieve their key business objectives. IDC is a subsidiary of IDG, the world's leading technology media, research, and events company.

Global Headquarters

5 Speen Street
Framingham, MA 01701
USA
508.872.8200
Twitter: @IDC
idc-insights-community.com
www.idc.com

Copyright Notice

This IDC research document was published as part of an IDC continuous intelligence service, providing written research, analyst interactions, telebriefings, and conferences. Visit www.idc.com to learn more about IDC subscription and consulting services. To view a list of IDC offices worldwide, visit www.idc.com/offices. Please contact the IDC Hotline at 800.343.4952, ext. 7988 (or +1.508.988.7988) or sales@idc.com for information on applying the price of this document toward the purchase of an IDC service or for information on additional copies or Web rights.

Copyright 2015 IDC. Reproduction is forbidden unless authorized. All rights reserved.

