



FreeNAS

Latest: 9.2.1.9 on November 18, 2014
[Details / Download](#)

- Home
- Forums**
- Resources
- Members
- FreeNAS Manuals
- BugReports
- Forum Rules

Mark Forums Read Search Forums Watched Forums

Home > Forums > Pre-RELEASE testing > 9.3 Testing >

SOLVED ESX 5.0 - ISCSI appears as SSD in ESX

Discussion in '9.3 Testing' started by depasseg, Nov 17, 2014.

Thread Status: Removed from public view.

Page 1 of 2 [1](#) [2](#) [Next >](#)

Thread Tools [Watch Thread](#) [Selected Posts: 0](#)

depasseg
FreeNAS Aware

Member Since: Sep 16, 2014

Messages: Message 104
Count: 104

Likes Received: 9

Trophy Points: 18

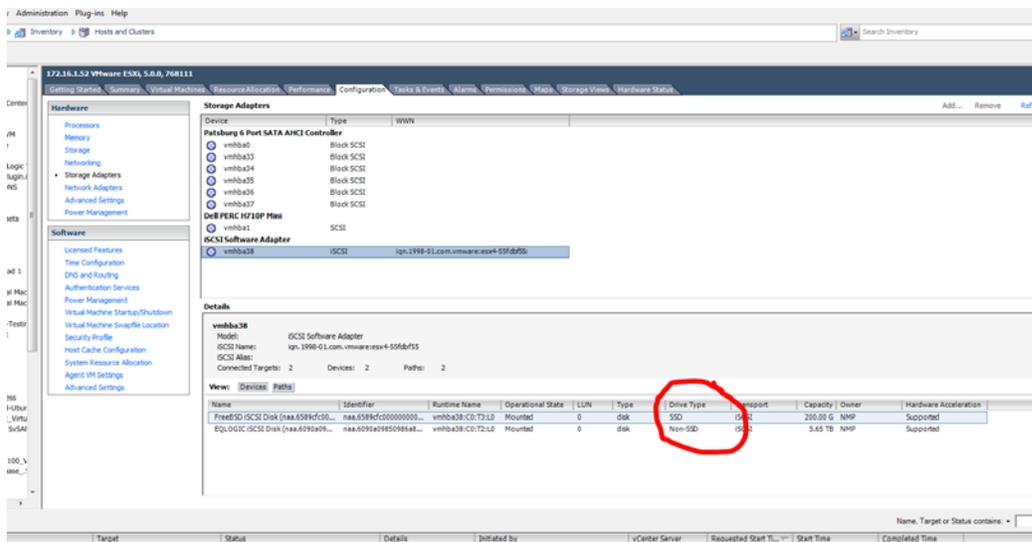
Occupation:
Solutions Architect

Online

depasseg, Nov 17, 2014

Tested out adding an iSCSI datastore to ESX. Surprisingly ESX thinks it's SSD backed. Has anyone else experienced this?

I have 6 WDC 3TB 7200rpm SATA in RAID 10 (used the wizard) as the backing.



9.3-beta | Xeon E5-2637V3 | 128GB DDR4 ECC | X10SRH-CLN4F | On-board LSI 3008 (Flashed v5-IT) | 12x 4TB WD-Re SAS (striped mirrors) | 2x Intel S3700 (SLOG and L2ARC) | SuperMicro SuperStorage Server 5028R-E1CR12L

depasseg, Nov 17, 2014 [Edit](#) [Delete](#) [IP](#) [Warn](#) [Report](#) #1 Like

jkh, Nov 17, 2014



Offline

jkh

Administrator

Member Since: Jul 22, 2013

Messages: Message 474
Count: 474

Likes Received: 161

Trophy Points: 43

That's how it's supposed to work, since an iSCSI device looks more like an SSD than a HD. You can even TRIM it.

jkh, Nov 17, 2014 [Edit](#) [Delete](#) [IP](#) [Report](#)

#2 [Like](#)



Online

depasseg

FreeNAS Aware

Member Since: Sep 16, 2014

Messages: Message 104
Count: 104

Likes Received: 9

Trophy Points: 18

Occupation:
Solutions Architect

depasseg, Nov 17, 2014

Tagging it as an SSD opens up a world of possible pain if it's not able to provide the IOPS of an SSD. Having this provided automatically is not desired. My EqualLogic array with 12x 7200 RPM drives doesn't appear as SSD (and for good reason).

Any chance of making this an option? I presume there is some sort of variable that gets based to vcenter to provide the identification?

9.3-beta | Xeon E5-2637V3 | 128GB DDR4 ECC | X10SRH-CLN4F | On-board LSI 3008 (Flashed v5-IT) | 12x 4TB WD-Re SAS (striped mirrors) | 2x Intel S3700 (SLOG and L2ARC) | SuperMicro SuperStorage Server 5028R-E1CR12L

depasseg, Nov 17, 2014 [Edit](#) [Delete](#) [IP](#) [Warn](#) [Report](#)

#3 [Like](#)



Offline

jpaetzel

FreeNAS Core Team

jpaetzel, Nov 17, 2014

What sort of pain could you experience? In reality under most workloads a FreeNAS iSCSI target does provide SSD performance. Sure with microbenchmarks you can blow out the cache and get disk performance but under typical usage scenarios you'll get SSD (or better) levels of performance.

jpaetzel, Nov 17, 2014 [Edit](#) [Delete](#) [IP](#) [Report](#)

#4 [Like](#)

Member Since: May 27, 2011
 Messages: Message 194
 Count: 194
 Likes Received: 33
 Trophy Points: 28
 Occupation: IT
 Location: Gilroy, CA
 Home page: <http://www.freenas.org>

jgreco, Nov 17, 2014



Offline

jgreco
 Resident Grinch
 Member Since: May 29, 2011
 Messages: Message 4,231
 Count: 4,231
 Likes Received: 441
 Trophy Points: 83
 Location: WHO(1)ville, Unixland

It might seem like it provides SSD-like performance on a sufficiently large array, but I think most FreeNAS users won't have that. Certainly I've never seen anything resembling "SSD performance" or even SSD-like performance out of FreeNAS, even on a 12 drive array.

I'm convinced enough that I can get "better-than-EqualLogic" performance (or the much lower "better-than-NetApp" bar) if I build a system right, so I actually do plan to try again to build a VM storage appliance on FreeNAS this year, but I've seen nothing that would cause me to expect magic SSD-like performance out of it even with a 24 drive array. So I'm truly curious what sort of magic pixie dust you're sprinkling in the code, hehehe.

jgreco, Nov 17, 2014 [Edit](#) [Delete](#) [IP](#) [Warn](#) [Report](#)

#5 Like

depasseg, Nov 17, 2014



Online

depasseg
 FreeNAS Aware
 Member Since: Sep 16, 2014
 Messages: Message 104
 Count: 104
 Likes Received: 9
 Trophy Points: 18
 Occupation: Solutions Architect

jpaetzel said: ↑

What sort of pain could you experience? In reality under most workloads a FreeNAS iSCSI target does provide SSD performance. Sure with microbenchmarks you can blow out the cache and get disk performance but under typical usage scenarios you'll get SSD (or better) levels of performance.

VMware will/can do things differently if it believes the underlying datastore is SSD (vs non-SSD). From VMware Docs (<http://pubs.vmware.com/vsphere-50/i..UID-E461B938-0C74-4F9A-8862-05CFB43407A2.html>):

Benefits of SSD Enablement
SSDs are very resilient and provide faster access to data.
SSD enablement results in several benefits:

-

It enables usage of SSD as swap space for improved system performance. For information about using SSD datastores to allocate space for host cache, see the vSphere Resource Management documentation.

- *It increases virtual machine consolidation ratio as SSDs can provide very high I/O throughput.*
- *It supports identification of virtual SSD device by the guest*

What if I configure storage policies to allow the use of SSD for Swap Space and the array isn't actually backed by SSDs?

What happens when I ask VMware to consolidate and it decides to increase the consolidation ratio because it thinks it's on a fast datastore?

9.3-beta | Xeon E5-2637V3 | 128GB DDR4 ECC | X10SRH-CLN4F | On-board LSI 3008 (Flashed v5-IT) | 12x 4TB WD-Re SAS (striped mirrors) | 2x Intel S3700 (SLOG and L2ARC) | SuperMicro SuperStorage Server 5028R-E1CR12L

depasseg, Nov 17, 2014 [Edit](#) [Delete](#) [IP](#) [Warn](#) [Report](#)

#6 [Like](#)

cyberjock, Nov 17, 2014



Offline

cyberjock
Forum Guard
Dog/Admin

Member Since: Mar 25, 2012
Messages: 15,112
Message Count: 15,111
Likes Received: 806
Trophy Points: 113

The first bullet only enabled your ability to use it. You must actually choose to use it.

The second bullet seems like a good idea because a properly running FreeNAS/TrueNAS machine running with good L2ARC and ARC sizes should perform better than spinning rust (but perhaps not as perfectly as all-SSD).

The third bullet either doesn't matter or might be useful for the unmap feature in iSCSI. That's something a Josh Paetzel would know the answer too.

So I'm not sure there's any critical downsides to it listing itself as SSD.

[Hardware recommendations](#) • [RAID5/RAIDZ1 is dead](#)

[Read my guide!](#) It answers common questions newbies to FreeNAS have.

Only you can prevent flame wars! Read the [FreeNAS manual](#).

cyberjock, Nov 17, 2014 [Edit](#) [Delete](#) [IP](#) [Report](#)

#7 [Like](#)

depasseg, Nov 17, 2014



There is a reason that iSCSI targets can be given the tag of non-SSD or SSD. There are policies that use this information. I'm not aware of something that ESX will change to increase performance because this is identified as an SSD. So why would we want to cause possible performance problems by calling it

depasseg
FreeNAS Aware

Member Since: Sep 16, 2014

Messages: Message Count: 104

Likes Received: 9

Trophy Points: 18

Occupation: Solutions Architect

Online

something it isn't. I get it, it could be faster than spinning disk, but you know what, it could be a lot slower than an actual SSD.

9.3-beta | Xeon E5-2637V3 | 128GB DDR4 ECC | X10SRH-CLN4F | On-board LSI 3008 (Flashed v5-IT) | 12x 4TB WD-Re SAS (striped mirrors) | 2x Intel S3700 (SLOG and L2ARC) | SuperMicro SuperStorage Server 5028R-E1CR12L

depasseg, Nov 17, 2014 [Edit](#) [Delete](#) [IP](#) [Warn](#) [Report](#) #8 [Like](#)

cyberjock
Forum Guard
Dog/Admin

Member Since: Mar 25, 2012

Messages: Message Count: 15,112

Likes Received: 806

Trophy Points: 113

Offline

cyberjock, Nov 17, 2014

depasseg said: ↑

So why would we want to cause possible performance problems by calling it something it isn't. I get it, it could be faster than spinning disk, but you know what, it could be a lot slower than an actual SSD.

You do know there's SSDs out there that perform MUCH worse than platter disks? Look at the first few generations of SSDs. They were beyond horrible!

[Hardware recommendations](#) • [RAID5/RAIDZ1 is dead](#)

[Read my guide!](#) It answers common questions newbies to FreeNAS have.

Only you can prevent flame wars! Read the [FreeNAS manual](#).

cyberjock, Nov 17, 2014 [Edit](#) [Delete](#) [IP](#) [Report](#) #9 [Like](#)

depasseg
FreeNAS Aware

Member Since: Sep 16, 2014

Messages: Message Count: 104

Likes Received: 9

Trophy Points: 18

Online

depasseg, Nov 17, 2014

There is no upside to marking a non-SSD iSCSI share as an SSD in the eyes of ESX. There is downside however. Now, perhaps if there was logic in FreeNAS to test the performance prior to stating it was as fast as/faster than an SSD, the down side would be removed. But I doubt that is the case. Is it?

If I configure a perfect pool and it is super high performance, and it isn't listed as SSD, I'm not going to miss out on performance.

If I configure a mediocre or poor performing pool (perhaps for ISO image storage or backups), and a storage policy thinks it's an SSD, I'm going to see performance problems.

All this does is create more work on the end user to go into the vsphere CLI and untag the device as an SSD. For what?

9.3-beta | Xeon E5-2637V3 | 128GB DDR4 ECC | X10SRH-CLN4F | On-board LSI 3008 (Flashed v5-IT) | 12x 4TB WD-Re SAS (striped mirrors) | 2x Intel S3700 (SLOG and L2ARC) | SuperMicro SuperStorage Server 5028R-E1CR12L

depasseg, Nov 17, 2014 [Edit](#) [Delete](#) [IP](#) [Warn](#) [Report](#)

#10 [Like](#)

cyberjock, Nov 17, 2014

Ok, you've been asked to cite specific examples and have none so far... do you not have an example and are just hypothesizing???

You say things like 'not going to miss out on performance' and 'I'm going to see performance problems' but have cited nothing to suggest that mis-labeling something is going to matter.

Is this just spreading FUD?

[Hardware recommendations](#) • [RAID5/RAIDZ1 is dead](#)

[Read my guide!](#) It answers common questions newbies to FreeNAS have.

Only you can prevent flame wars! Read the [FreeNAS manual](#).

cyberjock, Nov 17, 2014 [Edit](#) [Delete](#) [IP](#) [Report](#)

#11 [Like](#)

Ericloewe, Nov 17, 2014

depasseg said: ↑

There is no upside to marking a non-SSD iSCSI share as an SSD in the eyes of ESX. There is downside however. Now, perhaps if there was logic in FreeNAS to test the performance prior to stating it was as fast as/faster than an SSD, the down side would be removed. But I doubt that is the case. Is it?

If I configure a perfect pool and it is super high performance, and it isn't listed as SSD, I'm not going to miss out on performance.

If I configure a mediocre or poor performing pool (perhaps for ISO

TRIM allows ZFS to properly handle empty blocks, which should improve performance.

Trying to defrag a ZFS-backed iSCSI is silly and pointless, so exposing it as an SSD lets software know not to try that (Windows, in particular, automatically defrags HDDs by default).

FreeNAS 9.2.1.9-Release

Supermicro X10 SLM+-F with Intel Core i3 4330 and 2*8GB Crucial ECC 1.35V DDR3 1600MHz
6 * WD30EFRX WD Red 3TB in RAIDZ2 and 8GB Toshiba USB drive
Sharkoon T9 Value with 2 * Icy Dock FatCage MB153SP-B 3-in-2 drive cages
Seasonic G-550Gee, you're bored. APC Back-UPS Pro 900



Offline

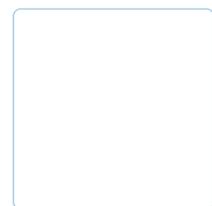
cyberjock
Forum Guard
Dog/Admin

Member Since: Mar 25, 2012

Messages: 15,112
Message Count: 15,112
Likes Received: 2

Likes Received: 806

Trophy Points: 113



Offline

Ericloewe
FreeNAS Guru

Member Since: Feb 15, 2014

Messages: 2,158
Message Count: 2,158

Likes Received: 151

Trophy Points: 63

Location: Portugal

Hardware Recommendations - ZFS basics (a.k.a Cyberjock's Guide)
 Supermicro X10 FAQ and Memory Recommendations

Ericloewe, Nov 17, 2014 [Edit](#) [Delete](#) [IP](#) [Warn](#) [Report](#)

#12 [Like](#)

mav@ likes this.

depasseg, Nov 17, 2014



Online

depasseg

FreeNAS Aware

Member Since: Sep 16, 2014

Messages: Message 104
Count:

104

Likes Received: 9

Trophy Points: 18

Occupation:

Solutions Architect

cyberjock said: ↑

Ok, you've been asked to cite specific examples and have none so far... do you not have an example and are just hypothesizing???
You say things like 'not going to miss out on performance' and 'I'm going to see performance problems' but have cited nothing to suggest that mis-labeling something is going to matter.
Is this just spreading FUD?

Maybe I missed it, but I can't find where I have been asked to cite specific examples. I was asked to explain how it could impact performance, and I did. If there is something more you need please let me know. As for FUD, I certainly have the F: Fear that an inattentive admin will configure a storage policy to take advantage of SSD storage, only to find that the non-ssd storage performance can't keep up.

So far the only supporting reason I've heard for keeping the incorrectly assigned "SSD" tag is for TRIM. TRIM is not supported by ESXi as CyberJock's own extremely detailed analysis has shown:

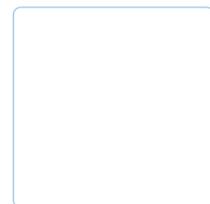
<https://forums.freenas.org/index.php...pport-trim-on-ssds-and-does-it-function.18167>

9.3-beta | Xeon E5-2637V3 | 128GB DDR4 ECC | X10SRH-CLN4F | On-board LSI 3008 (Flashed v5-IT) | 12x 4TB WD-Re SAS (striped mirrors) | 2x Intel S3700 (SLOG and L2ARC) | SuperMicro SuperStorage Server 5028R-E1CR12L

depasseg, Nov 17, 2014 [Edit](#) [Delete](#) [IP](#) [Warn](#) [Report](#)

#13 [Like](#)

depasseg, Nov 17, 2014



Online

depasseg

FreeNAS Aware

When I mention storage policies, I am referring to vmware things like Storage I/O Control and associated resource management. Admins can configure profiles and policies for VM's which can be given the appropriate resources. For instance, I might have a critical VM that requires higher than average IOPS. As I assign it to a resource pool, the resources will get divided up. If storage resource manger thinks a datastore is an SSD, it presumes a specific performance profile and allocates those resources. If the pool doesn't actually have those resources, performance will naturally degrade.

<http://pubs.vmware.com/vsphere-50/i...UID-37CC0E44-7BC7-479C-81DC-FFFC21C1C4E3.html>

Member Since: Sep 16, 2014
 Messages: Message Count: 104
 Likes Received: 9
 Trophy Points: 18
 Occupation: Solutions Architect

For comparison, imagine a VM required 4 cores and I told ESX that my dual core host CPU was actually 4 cores because in windows hyper threading is amazing and almost doubles the performance of my windows machine. Except imagine that ESX doesn't speak hyperthreading and only has 2 cores to divide the work against. Performance likely isn't going to be good for that VM.

9.3-beta | Xeon E5-2637V3 | 128GB DDR4 ECC | X10SRH-CLN4F | On-board LSI 3008 (Flashed v5-IT) | 12x 4TB WD-Re SAS (striped mirrors) | 2x Intel S3700 (SLOG and L2ARC) | SuperMicro SuperStorage Server 5028R-E1CR12L

depasseg, Nov 17, 2014 [Edit](#) [History](#) [Delete](#) [IP](#) [Warn](#) [Report](#)

#14 [Like](#)



Offline

cyberjock
 Forum Guard
 Dog/Admin

Member Since: Mar 25, 2012
 Messages: Message Count: 15,112
 Likes Received: 806
 Trophy Points: 113

cyberjock, Nov 17, 2014

ZFS's resource management is nothing like SSD or HDD. That's the reality of it. If everything is in RAM you are limited by your network ability and protocol limitations. So arguing SSD and/or HDD isn't really a very good argument for multiple reasons.

As for being asked to cite examples, I thought this was plenty good enough:

jpaetzel said: ↑

What sort of pain could you experience? In reality under most workloads a FreeNAS iSCSI target does provide SSD performance. Sure with microbenchmarks you can blow out the cache and get disk performance but under typical usage scenarios you'll get SSD (or better) levels of performance.

No example has been cited that could make ESXi explode because it can't get "SSD performance" (whatever threshold that is since SSDs themselves aren't necessarily fast, as dozens of threads have shown).

I'm just saying that calling something as SSD performance (versus HDD performance) is subjected and you've cited no specific reason why setting it to the (in your opinion) incorrect setting of "SSD" is particularly bad.

If your concern is your admin being stupid, you don't even want to try FreeNAS. It's easy as hell to check a box or two and totally fubar your server in extremely awesome ways.

[Hardware recommendations](#) • [RAID5/RAIDZ1 is dead](#)

[Read my guide!](#) It answers common questions newbies to FreeNAS have.

Only you can prevent flame wars! Read the [FreeNAS manual](#).

cyberjock, Nov 17, 2014 [Edit](#) [Delete](#) [IP](#) [Report](#)

#15 [Like](#)



depasseg
FreeNAS Aware

Member Since: Sep 16, 2014

Messages: Message 104
Count: 104

Likes Received: 9

Trophy Points: 18

Occupation:
Solutions Architect

Online

depasseg, Nov 17, 2014

Guys - I'm not the one who wrote the ESXi code to determine if a storage array was using SSD drives. I'm not the one who wrote the ESXi code that uses the T10 standard to provide an inquiry response mechanism by which a storage array will inform ESX of it's capabilities. <http://pubs.vmware.com/vsphere-50/i...UID-01FDF27C-77D8-46C4-931D-CCAAAACAE97B.html>

I am saying that since there is a concept in ESX of identifying both SSD and Non-SSD backed storage, FreeNAS should honor and respect that. The two labels exist for a reason.

9.3-beta | Xeon E5-2637V3 | 128GB DDR4 ECC | X10SRH-CLN4F | On-board LSI 3008 (Flashed v5-IT) | 12x 4TB WD-Re SAS (striped mirrors) | 2x Intel S3700 (SLOG and L2ARC) | SuperMicro SuperStorage Server 5028R-E1CR12L

depasseg, Nov 17, 2014 [Edit](#) [Delete](#) [IP](#) [Warn](#) [Report](#) #16 [Like](#)

Mlovelace likes this.



Mlovelace
FreeNAS Aware

Member Since: Aug 19, 2014

Messages: Message 139
Count: 139

Likes Received: 2

Trophy Points: 18

Location: Arizona

Offline

Mlovelace, Nov 17, 2014

jgreco said: ↑

It might seem like it provides SSD-like performance on a sufficiently large array, but I think most FreeNAS users won't have that. Certainly I've never seen anything resembling "SSD performance" or even SSD-like performance out of FreeNAS, even on a 12 drive array.

I'm convinced enough that I can get "better-than-EqualLogic" performance (or the much lower "better-than-NetApp" bar) if I build a system right, so I actually do plan to try again to build a

How does a TrueNAS stack up against a Netapp FAS2240-2 (48 x 600gig 10k SAS drives)? Given the same drives active/active controllers would TrueNAS be a better ESXi backend?

Supermicro x9slr-f, xeon E5-1620v2, 64gb ddr3-1866 ecc memory, lsi 9201-16i hba, chelsio t520-cr 10gbe, 12 x 3tb wd-se

RTFM

Mlovelace, Nov 17, 2014 [Edit](#) [History](#) [Delete](#) [IP](#) [Warn](#) [Report](#) #17 [Like](#)



cyberjock, Nov 17, 2014

depasseg said: ↑

I am saying that since there is a concept in ESX of identifying both SSD and Non-SSD backed storage, FreeNAS should honor and

respect that. The two labels exist for a reason.

cyberjock

Forum Guard
Dog/Admin

Member Since: Mar 25,
2012

Messages: Message
Count:
15,112
15,11
2

Likes Received: 806
Trophy Points: 113

Offline

We aren't disagreeing with you on that. We're asking you to provide an example where this differentiation really makes a difference. It's like saying a gasoline car uses gasoline, which makes it different from an E85 which can do E85 or gasoline, and you are upset because we chose to use gasoline. Ok, you are upset. But is there a justification to be upset that is more than "because E85 is a special misnomer and should be treated as such."

[Hardware recommendations](#) • [RAID5/RAIDZ1 is dead](#)

[Read my guide!](#) It answers common questions newbies to FreeNAS have.

Only you can prevent flame wars! Read the [FreeNAS manual](#).

cyberjock, Nov 17, 2014 [Edit](#) [Delete](#) [IP](#) [Report](#)

#18 [Like](#)

depasseg, Nov 17, 2014

Cyberjock you are disagreeing with exactly that statement. Vmware says call SSD - "SSD" and all other storage "Non-SSD". And FreeNAS says, "we don't disagree, but we are going to call Non-SSD - "SSD"". That is the definition of disagree.

depasseg

FreeNAS Aware

Member Since: Sep 16,
2014

Messages: Message
Count:
104
104

Likes Received: 9
Trophy Points: 18

Occupation:
Solutions Architect

Online

In your gas example. I'm upset because FreeNAS called their product gasoline (but was really, really good E85, in fact some would call it better than most gasoline), and I put the FreeNAS "Gasoline" in my antique boat tank and the ethanol disintegrated the epoxy in my fuel tank and gummed up the engine. Or the Ethanol absorbed so much moisture from the air that water collected in the tank over the winter storage. Does the reason really matter? There is a reason Gas is called Gas and E85 is called E85.

9.3-beta | Xeon E5-2637V3 | 128GB DDR4 ECC | X10SRH-CLN4F | On-board LSI 3008 (Flashed v5-IT) | 12x 4TB WD-Re SAS (striped mirrors) | 2x Intel S3700 (SLOG and L2ARC) | SuperMicro SuperStorage Server 5028R-E1CR12L

depasseg, Nov 17, 2014 [Edit](#) [Delete](#) [IP](#) [Warn](#) [Report](#)

#19 [Like](#)

jkh, Nov 18, 2014

Guys, this conversation really isn't convincing anyone to change the current behavior. Adding additional UI to change the identification type of the iSCSI device is a non-zero investment in time and coding, it doesn't just magically appear when folks ask for it, so leaving the status quo as it is (and investing that time and energy on more fruitful pursuits - we have plenty of bugs to fix in the bug tracker already) will be the default decision until/unless this conversation can get away from "Fnarr! Fnarr!" and into more convincing territory. Just saying... Thanks.

jkh

Administrator

Offline

Member Since: Jul 22, 2013
Messages: Message 474
Count: 474
Likes Received: 161
Trophy Points: 43

jkh, Nov 18, 2014 [Edit](#) [Delete](#) [IP](#) [Report](#)

#20 [Like](#)

Page 1 of 2 [1](#) [2](#) [Next >](#)

(You have insufficient privileges to reply here.)

Thread Status: Removed from public view.

Share This Page

[Tweet](#) [g+](#) **Like** Be the first of your friends to like this.

[Home](#) [Forums](#) [Pre-RELEASE testing](#) [9.3 Testing](#)

[Help](#) [Home](#) [Top](#)

Forum software by XenForo™ ©2010-2014 XenForo Ltd.

Tac Anti Spam from Surrey Forum



FreeNAS

Latest: 9.2.1.9 on November 18, 2014

[Details](#) / [Download](#)[Home](#)**Forums**[Resources](#)[Members](#)[FreeNAS Manuals](#)[BugReports](#)[Forum Rules](#)[Mark Forums Read](#)[Search Forums](#)[Watched Forums](#)[Home](#)[Forums](#)[Pre-RELEASE testing](#)[9.3 Testing](#)

SOLVED ESX 5.0 - iSCSI appears as SSD in ESX

Discussion in '9.3 Testing' started by depasseg, Nov 17, 2014.

Thread Status: Removed from public view.

Page 2 of 2 [< Prev](#) [1](#) [2](#)[Thread Tools](#)[Watch Thread](#)[Selected Posts: 0](#)

jgreco, Nov 18, 2014

Jordan, you've had a new user, [@depasseg](#), come in and give you the view that an experienced storage admin would give you, and you've been dismissive. It is basically similar to what I would have said had I been awake. What you're doing is requiring vSphere admins to fix your brokenness on their side, which leaves a poor impression. The SSD tag is there to denote the underlying storage technology, because different rules apply to platters and flash. Even when the platters or flash are buried behind a layer of ZFS.

jgreco

Resident Grinch

Member Since: May 29, 2011

Messages: Messag 4,231
eCount:
4,231

Likes Received: 441

Trophy Points: 83

Location:

WHO(1)ville, Unixland

Offline

cyberjock said: ↑

We aren't disagreeing with you on that. We're asking you to provide an example where this differentiation really makes a difference.

I love seeing cyberjock argue the side of "it's OK to use non-ECC in your server" though; there's a fair bit of cognitive dissonance going on here. [@cyberjock](#) - this isn't about providing an example. He already did, with links to VMware's documents, in #6 and #14. You're playing the part of the guy who doesn't understand why ECC makes a difference because you haven't experienced something firsthand.

But those of us out here in the trenches do rely on things like SIOC, because manually configuring every aspect of every bit of hardware FRELLING SUCKS.

jkh said: ↑

Adding additional UI to change the identification type of the iSCSI device is a non-zero investment in time and coding, it doesn't just

magically appear when folks ask for it,

This particularly disingenuous and dismissive response to add what is most likely a passthru flag to a conf file is fine, for me, it marks the end of the conversation because it is clear that there's no interest in technical correctness even when explained.

I leave you with the following amusing image.



■ jgreco, Nov 18, 2014 [Edit](#) [Delete](#) [IP](#) [Warn](#) [Report](#)

#21 [Like](#)

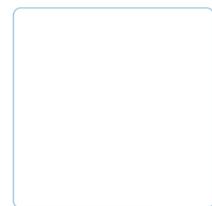
[Mlovelace](#) and [enemy85](#) like this.

cyberjock, Nov 18, 2014

I totally disagree with your assessment of my opinion. I've provided a very good example in my ECC thread of precisely why ECC RAM is very important for ZFS. Right now I've only seen hand-waiving and "this should be corrected to be proper". Sure, ideally you'd want everything to be correct. But he hasn't really cited a real-world scenario where this would be particularly bad in an automatic sort of way (like the ECC discussion).

It does little more than enable the admin to use some features. Features that either provide zero benefit because of conditions or won't hurt anything because of condition. The only argument I'm seeing made is the technical grounds for accurate representation of the storage architecture. I'm not seeing an example that says something like "as a result of the media being presented as SSD the storage architecture ESXi expects X to happen automatically. But we all know this won't happen and as a result the ESXi server will crash, blow up or sleep with your mom".

Yes, the argument could be made that it might allow a user to misconfigure a server. But, we've got plenty of stuff in the WebGUI that, if set, is often permanent and often uncorrectable. Should we disable dedup, RAIDZ1, and



Offline

cyberjock
Forum Guard
Dog/Admin

Member Since: Mar 25, 2012

Messages: 15,112
Message Count: 15,112
Trophy Points: 2

Likes Received: 806

Trophy Points: 113

RAID0 because we deem them to be stupid and dangerous? Should we program FreeNAS to refuse to boot with 8GB of RAM in 9.3? It seems to me you are trying to make the argument to do something "because we can" and not because there's actual technical reasons for doing it.

[Hardware recommendations](#) • [RAID5/RAIDZ1 is dead](#)

[Read my guide!](#) It answers common questions newbies to FreeNAS have.

Only you can prevent flame wars! Read the [FreeNAS manual](#).

cyberjock, Nov 18, 2014 [Edit](#) [Delete](#) [IP](#) [Report](#)

#22 [Like](#)

jkh, Nov 18, 2014

Umm. OK? I don't see any value being added by [@jgreco](#) here, he's just repeating the same things in a more nasal, whiny voice. Again, still waiting for an answer to [jpaetzel](#)'s question regarding why this will cause actual pain and an apocalyptic rain of frogs.

I will also note, just to add a little context, that this is the first and only complaint of this nature we have received, and we sell a lot of VMWare storage servers. We use ESXi ourselves for running tons of VMs in multiple locations with TrueNAS HA systems on the back end and never has this been a problem for us either. I'm still waiting for clear evidence that this is something other than someone's OCD twitching, and the nasal whining doesn't qualify. If the original poster would care to file a bug report with more justification, we will look at it. We will even run it by our VMWare rep for an informed 2nd opinion. I'm certainly not going to send VMWare into the forum for *this* conversation, however.

jkh, Nov 18, 2014 [Edit](#) [History](#) [Delete](#) [IP](#) [Report](#)

#23 [Like](#)

depasseg, Nov 18, 2014

Please talk to your VMware rep and engineer. Ask them to confirm that there isn't an issue with FreeNAS spinning disks getting reported as SSD. Ask if it will affect Storage DRS or other Policy and Profile Driven Storage capabilities.

I could write about VM performance degrading due to over-allocation on over-reported HW, but there are already a ton of examples out there.

9.3-beta | Xeon E5-2637V3 | 128GB DDR4 ECC | X10SRH-CLN4F | On-board LSI 3008 (Flashed v5-IT) | 12x 4TB WD-Re SAS (striped mirrors) | 2x Intel S3700 (SLOG and L2ARC) | SuperMicro SuperStorage Server 5028R-E1CR12L

depasseg, Nov 18, 2014 [Edit](#) [History](#) [Delete](#) [IP](#) [Warn](#) [Report](#)

#24 [Like](#)



Offline

jkh

Administrator

Member Since: Jul 22, 2013

Messages: Message 474
Count: 474

Likes Received: 161

Trophy Points: 43



Online

depasseg

FreeNAS Aware

Member Since: Sep 16, 2014

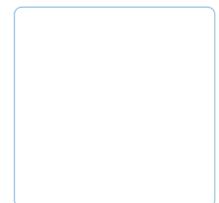
Messages: Message 104
Count:

diehard
Trophy Points: 104 18
Solutions Architect

diehard, Nov 18, 2014

I thought Storage DRS went off I/O latency, not reported datastore type.

Likening this case to other cases based on over-allocated VM's is kinda misguided IMO.



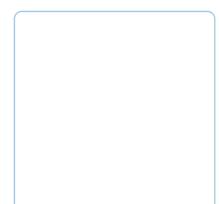
Offline

diehard
Newbie

Member Since: Mar 21, 2013
Messages: Message 75
Count: 75
Likes Received: 0
Trophy Points: 6

diehard, Nov 18, 2014 Edit History Delete IP Warn Report

#25 Like



cyberjock, Nov 18, 2014

diehard said: ↑

Likening this case to other cases based on over-allocated VM's is kinda misguided IMO.

Offline

cyberjock
Forum Guard
Dog/Admin

Member Since: Mar 25, 2012
Messages: Message 15,112
Count: 15,112
Likes Received: 806
Trophy Points: 113

That's why I didn't even respond. Arguing over what "overallocated hardware" is is so subjective that it's no argument at all. It's like claiming someone is tall. Well, compared to a 3 year old, anyone over 3 feet tall is "tall". To me, anyone that isn't over 6'4" is "short".

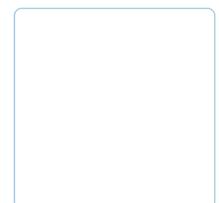
[Hardware recommendations](#) • [RAID5/RAIDZ1 is dead](#)

[Read my guide!](#) It answers common questions newbies to FreeNAS have.

Only you can prevent flame wars! Read the [FreeNAS manual](#).

cyberjock, Nov 18, 2014 Edit Delete IP Report

#26 Like



jkh, Nov 18, 2014

I don't know what kind of "over-allocation" @depasseg is specifically referring to, but if he's talking about the overcommit scenario that results from handing out a bunch of sparsely-allocated zvols to VMWare via iSCSI and then filling up the pool (and therefore causing those zvols to be unable to allocate more blocks on demand), that is also now dealt with properly in 9.3 given the enhanced VAAI support. For one thing, we now support STUN, which will phaser the relevant VMs into unconsciousness rather than let them freak out about not being able to write on what they perceive to be perfectly good SSDs (ha, you see

Offline

jkh
Administrator

Member Since: Jul 22, 2013
Messages: Message 474
Count: 474
Likes Received: 161
Trophy Points: 43

what I did there) and for the other, we also support the notion of a pool allocation threshold which will send low disk space warnings to VMWare when that threshold is exceeded by the pool, as the attached screenshot demonstrates.

Sharing

Apple (AFP) UNIX (NFS) WebDAV Windows (CIFS) **Block (iSCSI)**

Target Global Configuration Portals Initiators Authorized Access Targets Extents Associated Targets

Base Name: ⓘ

Discovery Auth Method:

Discovery Auth Group:

ISNS Servers:

Pool Available Size Threshold (%):

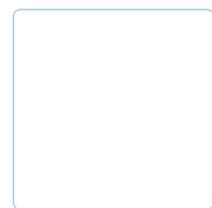
Save

jkh, Nov 18, 2014 Edit Delete IP Report

#27 Like

cyberjock, Nov 18, 2014

phaser!? You have to set the phaser to "kill" if you want any kind of impact.



Offline

cyberjock
Forum Guard
Dog/Admin

Member Since: Mar 25, 2012
Messages: Message 15,112
Count: 15,112
Likes Received: 806
Trophy Points: 113

Hardware recommendations • RAID5/RAIDZ1 is dead

[Read my guide!](#) It answers common questions newbies to FreeNAS have.

Only you can prevent flame wars! Read the [FreeNAS manual](#).

cyberjock, Nov 18, 2014 Edit Delete IP Report

#28 Like

depasseg, Nov 18, 2014



Online

cyberjock said: ↑

phaser!? You have to set the phaser to "kill" if you want any kind of impact.

I find your lack of faith disturbing. [darth voice]

depasseg

FreeNAS Aware

Member Since: Sep 16, 2014

Messages: Message 104
Count: 104

Likes Received: 9

Trophy Points: 18

Occupation:
Solutions Architect

9.3-beta | Xeon E5-2637V3 | 128GB DDR4 ECC | X10SRH-CLN4F | On-board LSI 3008 (Flashed v5-IT) | 12x 4TB WD-Re SAS (striped mirrors) | 2x Intel S3700 (SLOG and L2ARC) | SuperMicro SuperStorage Server 5028R-E1CR12L

depasseg, Nov 18, 2014 [Edit](#) [History](#) [Delete](#) [IP](#) [Warn](#) [Report](#)

#29 Like



Online

depasseg

FreeNAS Aware

Member Since: Sep 16, 2014

Messages: Message 104
Count: 104

Likes Received: 9

Trophy Points: 18

Occupation:
Solutions Architect

depasseg, Nov 18, 2014

The Over Allocation I was discussion isn't storage volume or capacity, it's performance/IOPS.

According to VMware, SSD backed data stores can provide additional capabilities like I've linked to in the past (host cache, increased VM Consolidation ratio, guest OS detection). Now, I am not going to argue the benefit of those capabilities. They exist. VMware tracks them. I'm not going to defend them. But they appear to exist for a reason. I've outlined from VMware documentation what that label means, what it should be applied to and what special things apply to SSD datastores.

diehard said: ↑

I thought Storage DRS went off I/O latency, not reported datastore type.

Likening this case to other cases based on over-allocated VM's is kinda misguided IMO.

The definition of Drive Type from VMware is:

Type of underlying storage device, a Solid State Drive (SSD) or a regular non-SSD hard drive. For details, see [Working with Solid State Disks](#).

<https://pubs.vmware.com/vsphere-55/...UID-057D6054-0A51-4023-B90A-D737DB0426F4.html>

Under "Working with Solid State Disks is the following:"

Working with Solid State Disks
In addition to regular storage hard disk drives, ESXi supports Solid State Disks (SSDs).
Unlike the regular hard disks that are electromechanical devices

containing moving parts, SSDs use semiconductors as their storage medium and have no moving parts. Typically, SSDs are very resilient and provide faster access to data.

Support of SSDs with ESXi provides the following benefits:

- *You can use VMFS datastores that are created on SSD storage devices to allocate space for ESXi host cache.*
- *You can use SSDs for such features as Flash Read Cache and*

<https://pubs.vmware.com/vsphere-55/...tml#GUID-A3457F81-0E12-4770-A2F0-F83C9CA1E2F3>

And under the Guest OS detection section is the following:

Identifying a Virtual SSD Device

ESXi allows operating systems to auto-detect VMDKs residing on SSD datastores as SSD devices.

To verify if this feature is enabled, guest operating systems can use standard inquirv commands such as SCSI VPD Page (B1h) for

Click to expand...

<https://pubs.vmware.com/vsphere-55/...UID-E9E146C9-E99C-4468-B70C-770B83788433.html>

So I guess regarding the highlighted line above from VMware, "Check with your vendor", Is it accurate to say the official stance of iXsystems, TrueNAS and FreeNAS is - "Our product meets or exceeds SSD performance, and we identify all of our storage as SSD". Is that accurate?

And as for why no one has mentioned this before, my guess would be they didn't understand it. Or maybe some didn't think it was worth the uphill battle. cyberjock - you're a smart guy but your tone can sometimes be very condescending and your avatar quote "Forum Guard Dog" isn't the least bit welcoming.

I have questions into 2 of my storage OEM partners that fully support VMware. I'm going to try to find out what their understanding is. Maybe I'm wrong and this affects nothing. In which case I will apologize for wasting everyone's time and effort. I would love to hear what anyone else's VMware or storage OEM has to say.

Greg

9.3-beta | Xeon E5-2637V3 | 128GB DDR4 ECC | X10SRH-CLN4F | On-board LSI 3008 (Flashed v5-IT) | 12x 4TB WD-Re SAS (striped mirrors) | 2x Intel S3700 (SLOG and L2ARC) | SuperMicro SuperStorage Server 5028R-E1CR12L

depasseg, Nov 18, 2014 [Edit](#) [History](#) [Delete](#) [IP](#) [Warn](#) [Report](#)

#30 [Like](#)

zambanini and Mlovelace like this.

Mlovelace, Nov 19, 2014

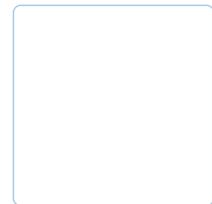
We have several Netapp SANs and a couple Nimble SANs (which use quite a bit of SSD cache) and they don't report their Datastores as SSD to ESXi. I'll put a call into our VMWare support engineer and see what he has to say.

Supermicro x9slr-f, xeon E5-1620v2, 64gb ddr3-1866 ecc memory, lsi 9201-16i hba, chelsio t520-cr 10gbe, 12 x 3tb wd-se

RTFM

Mlovelace, Nov 19, 2014 [Edit](#) [Delete](#) [IP](#) [Warn](#) [Report](#)

#31 [Like](#)



Offline

Mlovelace

FreeNAS Aware

Member Since: Aug 19, 2014

Messages: Message 139
Count: 139

Likes Received: 2

Trophy Points: 18

Location: Arizona

diehard, Nov 20, 2014

depasseg said: ↑

The definition of Drive Type from VMWare is:

Yes i know, but Storage DRS still operates off of I/O latency. It wont migrate storage to a datastore just because its reported as "SSD". I don't really see how this particular concern is an issue.

Most of the benefits under their "Working with Solid State Disks" documentation is the inherent benefit of having SSD's, not changing the way ESXi operates with those datastores. Yes, because it is reported as an SSD it will let you designate it as host swap cache, but if someone doesnt know not to do that then they really shouldn't be using either of these products.

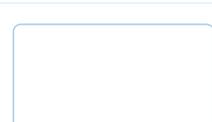
Having a guest OS pass TRIM commands is not a bad thing (you have to use Windows 8 for a guest OS to detect a SCSI device as an SSD btw).

diehard, Nov 20, 2014 [Edit](#) [History](#) [Delete](#) [IP](#) [Warn](#) [Report](#)

#32 [Like](#)

mav@, Nov 20, 2014

In my understanding, the principal difference between SSD and HDD is that for HDD linear reads are much faster then the random ones (or random are much





mav@
FreeNAS Aware

Offline

Member Since: Sep 29, 2011

Messages: Message 108
Count: 108

Likes Received: 12

Trophy Points: 18

Occupation:
FreeBSD developer

Location:
Dnepropetrovsk,
Ukraine

slower). Speaking about ZFS that is not obvious: if you have enough ARC/L2ARC for your workload -- any access pattern should give you the same speed; if ARC is not enough, then ZFS still does not guarantee that consecutive LBAs will be stored consequentially, so linear speed is not necessarily better than random.

As already mentioned above, reporting device type as SSD blocks Windows disk defragmenter. Defragmenting ZVOL is not necessary helpful. It reduces logical volume fragmentation by increasing physical pool fragmentation and, if snapshots enabled, space usage. Lower logical fragmentation should help read prefercher, but I am not sure that worth time spent and other consequences.

Kernel part of CTL in 9.3 already has config file option "rpm", that allows to control that SSD/HDD status (SSD has RPM=1). So if there is really good arguments to control that in UI, that should not be really difficult to add. From the other side every new option is an additional way to confuse users.

mav@, Nov 20, 2014 [Edit](#) [Delete](#) [IP](#) [Warn](#) [Report](#)

#33 [Like](#)



Dave Genton

Offline

Member Since: Feb 27, 2014

Messages: Message 20
Count: 20

Likes Received: 0

Trophy Points: 1

Occupation:
Cisco Consulting
Engineer

Location: Florida

Dave Genton, Nov 22, 2014

I use alot of iSCSI LUN targets and also do not like seeing it show as an SSD. One issue I have seen already is now that all SSD stores show up in the host cache listing to be selected and used as host cache. Obviously local flash/ssd is being utilized for many tools today in the areas of acceleration and caching but remote iSCSI storage by default getting put into the same category is not optimal in my mind.

Dave - CCIE #6746

Dave Genton, Nov 22, 2014 [Edit](#) [Delete](#) [IP](#) [Warn](#) [Report](#)

#34 [Like](#)



jgreco

Offline

jgreco, Nov 28, 2014

Nobody I've spoken to thinks that having NAS storage show up as "SSD" is brilliant. VMware allows several drive types, including "Unknown" and "Non-SSD", and I cannot ever recall having seen a non-SSD datastore just automatically show up as SSD. Heck, on a newly deployed ESXi 5.5 box just today, a pair of Intel S3500's on an HP SmartArray P400i shows up as "Non-SSD", apparently because it was behind an array controller, even though it's DAS.

jgreco

Resident Grinch

Member Since: May 29, 2011

Messages: 4,231
Message Count: 4,231

Likes Received: 441

Trophy Points: 83

Location: WHO(1)ville, Unixland

VMware's definition of "SSD" is extremely specific, and having a NAS vendor describe their spinny-rust-over-Ethernet with the words "a FreeNAS iSCSI target does provide SSD performance" - that basically comes out as a load of horse manure the moment you look at the link speed.

jkh said: ↑

Umm. OK? I don't see any value being added by @jgreco here, he's just repeating the same things in a more nasal, whiny voice.

Hey @jkh? If you can't argue based on technical merits and have to resort to ad-hominem attacks, go get bent.

jgreco, Nov 28, 2014 [Edit](#) [Delete](#) [IP](#) [Warn](#) [Report](#)

#35 [Like](#)

(You have insufficient privileges to reply here.)

Thread Status: Removed from public view.

Share This Page

Tweet < 0

g+1 < 0

Like Be the first of your friends to like this.

[Home](#)

[Forums](#)

[Pre-RELEASE testing](#)

[9.3 Testing](#)

[Help](#) [Home](#) [Top](#)