



**NetApp™**  
Go further, faster

Technical Report

# Using Performance Acceleration Module in Exchange 2007 Solutions

Wei Liu, Lee Dorrier, NetApp  
June 2009 | TR-3766

## **EXECUTIVE SUMMARY**

This report describes the benefits of the Performance Acceleration Module (PAM) for Microsoft® Exchange 2007 workloads based on lab studies. It also describes the processes that enable one to answer some frequently asked questions regarding PAM and Exchange, such as how many disks a PAM card would be able to save and how many more users a configuration could support.

TABLE OF CONTENTS

<b>1</b>	<b>INTRODUCTION .....</b>	<b>3</b>
<b>2</b>	<b>FAS3040 AND PAM WITH SATA DRIVES.....</b>	<b>3</b>
2.1	FAS3040 TEST CONFIGURATION.....	3
2.2	8,000-USER TEST RESULTS .....	3
<b>3</b>	<b>FAS3170 AND PAM WITH FC DRIVES.....</b>	<b>4</b>
3.1	FAS3170 TEST CONFIGURATION.....	4
3.2	60,000-USER TEST RESULTS .....	5
<b>4</b>	<b>DISCUSSIONS .....</b>	<b>6</b>
4.1	HOW MANY DRIVES CAN I SAVE? .....	6
4.2	HOW MANY MORE USERS CAN I SUPPORT?.....	6
4.3	DOES PAM ENABLE THE USE OF SATA DRIVES FOR EXCHANGE 2007?.....	6
4.4	WHEN SHOULD I USE PAM?.....	6
4.5	HOW MUCH CAN PAM HELP? .....	6
<b>5</b>	<b>CONCLUSIONS .....</b>	<b>7</b>
<b>6</b>	<b>ACKNOWLEDGEMENTS .....</b>	<b>7</b>
<b>7</b>	<b>REFERENCES .....</b>	<b>8</b>

# 1 INTRODUCTION

The Performance Acceleration Module is a second-layer cache [1] that can accelerate an application's read operations, especially small random reads such as those seen in Microsoft Exchange. Using PAM can enhance Exchange's responsiveness, support more users on a given controller with a fixed number of hard disk drives (HDDs), or reduce the number of hard drives needed for a set Exchange configuration.

In this technical report, a specific implementation of PAM, namely PAM I, a PCI Express X4 card with 16GB of SDRAM (together with its software component), was tested using two vastly different Exchange 2007 configurations. Both configurations benefited from the use of PAM.

Microsoft Exchange Server is an I/O-intensive application that requires a reliable and high-performance storage system to function correctly [2]. Jetstress is a Microsoft tool that can accurately simulate Exchange I/O characteristics and validate the performance and reliability of the Exchange storage solutions [2]. Jetstress is also a testing harness of ESRP, a Microsoft Exchange Solution Reviewed Program, for facilitating third-party Exchange storage testing and solution publishing [3].

The purpose of this technical report is to demonstrate the benefits of PAM for Microsoft Exchange 2007, using ESRP-like test procedures and environment. As such, the test results/data presented in this report should not be quoted directly for any pre-deployment validation. Also, this report is not a substitute for best practices nor deployment guides.

Sections 2 and 3 below discuss the Jetstress tests of Exchange 2007 on a FAS3040 with SATA drives<sup>†</sup>, and on a FAS3170 with high-performance Fibre Channel (FC) drives, respectively. Section 4 outlines a process that can be used to determine under what circumstances PAM will be helpful for an Exchange workload in general.

## 2 FAS3040 AND PAM WITH SATA DRIVES

This section describes briefly the Jetstress tests of an 8,000-user configuration for Exchange 2007 on a pair of FAS3040 storage controllers using SATA hard disk drives, with and without PAM.

### 2.1 FAS3040 TEST CONFIGURATION

Data ONTAP 7.3 was installed on the pair of FAS3040 storage controllers. Each controller was configured with one PAM module. Table 1 below shows the most important parameters of the test configuration.

Table 1) Summary of FAS3040 and PAM Jetstress test configuration.

Test Parameters	PAM Enabled	PAM Disabled
8,000 users	✓	✓
0.30 IOPS per user	✓	✓
1000MB per mailbox	✓	✓
2 servers	✓	✓
2 FAS3040 controllers	✓	✓
2 PAM modules (1 per controller)	✓	✓
82 HDDs (SATA 500GB 7.2K RPM)	✓	✓
FlexScale settings	<code>flexscale.enable on</code> <code>flexscale.normal_data_blocks off</code>	<code>flexscale.enable off</code>

Each Jetstress test was run twice: once with PAM enabled to use its *metadata mode* [1] and once with PAM disabled. The FlexScale settings row in Table 1 shows how this was done. The reason to use PAM metadata mode (the command, `options flexscale.normal_data_blocks off`, enables metadata mode) is that typically the Exchange working set size is too large for a 16GB PAM to cache effectively in *default mode* [1].

### 2.2 8,000-USER TEST RESULTS

<sup>†</sup> The Jetstress test results on the FAS3040 with SATA drives in this technical report should apply in general to a similarly configured FAS3140 with SATA drives.

Key metrics from the 8,000-user test for the two-hour Jetstress performance test, both with PAM and without PAM, are summarized in Table 2 below.

Table 2) Summary of FAS3040 Jetstress test results with and without PAM.

	PAM Enabled	PAM Disabled	Improvement Due to PAM
Achieved database IOPS	2,508	2,086	20%
Average database reads/sec	1,254	1,036	21%
Average database writes/sec	1,254	1,050	19%
Avg. database read latency (ms)	16	19	16%
Avg. database write latency (ms)	2	2	N/A
Avg. log writes/sec	860	724	19%
Avg. log write latency (ms)	1	1	N/A

Two main benefits of PAM were observed: More database IOPS (20% in this case), and lower read latency (16% in this case). These benefits also apply to the FAS3170 using FC drives and will be described in more details in Section 3.

Another benefit of using PAM was an increased number of I/Os achieved during the test, resulting in a successful execution of the e-mail profile. This is an example of how PAM can enable the use of low-cost SATA drives in Exchange storage solutions.

### 3 FAS3170 AND PAM WITH FC DRIVES

This section describes the Jetstress tests of a 60,000-user configuration for Exchange 2007 on a pair of FAS3170 storage controllers using FC hard disk drives, with and without PAM.

#### 3.1 FAS3170 TEST CONFIGURATION

Data ONTAP® 7.3.1<sup>‡</sup> was installed on the pair of FAS3170 storage controllers. Each controller was configured with one PAM module. Table 3 below shows the most important parameters of the test configuration. More details can be found in the 60,000-user ESRP report [4].

Table 3) Summary of FAS3170 and PAM Jetstress test configuration.

Test Parameters	PAM Enabled	PAM Disabled
60,000 users	✓	✓
0.48 IOPS per user	✓	✓
400MB per mailbox	✓	✓
6 servers	✓	✓
2 FAS3170 controllers	✓	✓
2 PAM modules (1 per controller)	✓	✓
184 HDDs (FC 300GB 15K RPM)	✓	✓
FlexScale settings	flexscale.enable on flexscale.normal_data_blocks off	flexscale.enable off

Each Jetstress test was run twice: once with PAM enabled to use its *metadata mode* [1] and once with PAM disabled.

<sup>‡</sup> The results in Section 3 should apply in general to 7.3.1.x version of Data ONTAP.

### 3.2 60,000-USER TEST RESULTS

Key metrics from the two-hour Jetstress performance test, both with PAM and without PAM, are summarized in Table 4 below.

Table 4) Summary of FAS3170 Jetstress test results with and without PAM.

	PAM Enabled	PAM Disabled	Improvement Due to PAM
Achieved database IOPS	32,635	28,871	13%
Average database reads/sec	16,094	14,487	11%
Average database writes/sec	16,541	14,383	15%
Avg. database read latency (ms)	14	16	11%
Avg. database write latency (ms)	3	2	N/A <sup>††</sup>
Avg. log writes/sec	12,850	11,832	9%
Avg. log write latency (ms)	2	1	N/A <sup>††</sup>

The 60,000-user test again demonstrates several benefits of using PAM. First, when PAM was enabled, 3,764 more database IOPS were achieved. This increase in IOPS is mainly due to the reduction in overall I/O latency when PAM is enabled. These extra IOPS due to PAM can potentially lead to support of more users or make it possible to configure fewer hard drives for the given user count. How many more users can be supported or how many hard drives can be saved? Figure 1 shows a simple approach, suggesting that 7,841 more users could be supported or that 18 hard drives could be eliminated from the 184-drive configuration by using PAM.

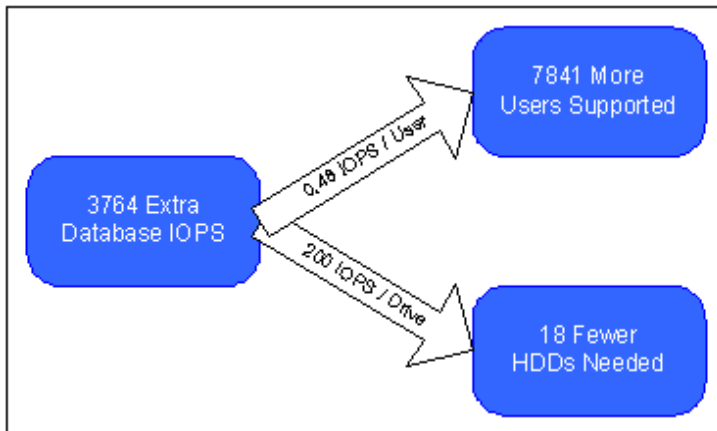


Figure 1) A coarse estimate of users or fewer HDDs based on extra IOPS due to PAM.

Be advised that this is a coarse estimate. A more precise estimate requires the analysis of constraints such as available capacity, cache utilization, working set size, and controller CPU utilization, in addition to the database IOPS. Section 4 describes a process that can be used to answer these questions semi quantitatively. Customers should work with their NetApp Exchange field specialists to conduct sizing with and without PAM to accurately assess the impact of PAM.

Another benefit of PAM is lower read latency: 14.4 ms compared 16.2 ms without PAM. Average database read and write latencies greater than 20 ms are considered by Microsoft to be poor, and can often result in slow performance and a poor user experience [5]. Therefore, a lower database read latency is very important to Exchange. PAM reduces this latency, ensuring a better Exchange solution. Note that the increase in database and log write latencies is due to more load on the controller overall and is not a function of PAM.

Yet another important benefit of PAM is better storage efficiency. Clearly, better storage efficiency can be a result of either supporting more users with a given number of hard drives or supporting the same number of users with fewer hard drives. In fact, the 60,000-user test with PAM yielded 53% storage efficiency or capacity utilization [6].

<sup>††</sup> The increase in write latency is due to the increased load and not a function of PAM.

## 4 DISCUSSIONS

The success of PAM for Exchange 2007 on a FAS3170 (with FC drives) and FAS3040 (with SATA drives) was based on a simple configuration model:

- Use enough hard drives to handle all user I/Os.
- Use PAM to cache metadata and reduce disk reads.
- Ensure controller CPU can handle the small overhead due to the PAM software component.

The tested FAS3170 configuration is an example of a very large working set, and the storage utilization in terms of both capacity and disk IOPS were optimized. On the other hand, the tested FAS3040 configuration is an example of a disk-limited configuration due to the low-cost SATA drives. It is clear that PAM is effective in reducing database read latencies in both scenarios.

The remainder of this section will answer some frequently asked questions about PAM and Exchange. Again, we strongly recommend that customers work with their NetApp Exchange specialists to conduct sizing with and without PAM to accurately assess the impact of PAM.

### 4.1 HOW MANY DRIVES CAN I SAVE?

It depends. You need to consider two factors: capacity and IOPS. PAM helps increase IOPS. Use the following three-step process to get a better answer.

Step 1. Capacity: How many disks can I reduce and still provide sufficient space to support all the mailboxes?

Step 2. IOPS: How many disks can I reduce because of PAM's boost of IOPS?

Step 3. Result: the minimum of steps 1 and 2.

### 4.2 HOW MANY MORE USERS CAN I SUPPORT?

It depends. You need to consider four factors: capacity, IOPS, cache utilization, and controller CPU. PAM can only help IOPS, not the other factors. Use the following five-step process to get a better answer.

Step 1. Capacity: How much spare capacity is available to support more users?

Step 2. IOPS: How many more IOPS can the existing disks support?

Step 3. Cache utilization: Can PAM cache all metadata with added users? NetApp's Predictive Cache Statistics (PCS) software [1] can be used to help answer this question.

Step 4. Storage controller CPU: Use the Exchange Sizer to verify the controller can support more users.

Step 5. Result: the minimum of steps 1–4.

### 4.3 DOES PAM ENABLE THE USE OF SATA DRIVES FOR EXCHANGE 2007?

The answer is yes, based on the findings in Section 2.

### 4.4 WHY SHOULD I USE PAM?

PAM allows you to scale more efficiently by increasing the number of users or deliver lower read latency without adding more drives.

### 4.5 HOW MUCH CAN PAM HELP?

It depends. The processes outlined in sections 4.1 and 4.2 can be used to do semi quantitative analysis.

## **5 CONCLUSIONS**

PAM can enhance Exchange 2007 performance and support more users on a given controller without the need to add more disk drives. Or it can save the number of hard drives needed by a given Exchange user profile. This paper provides hands-on test results and rough guidelines for evaluating the effectiveness of adding a PAM card to an Exchange 2007 configuration. For a more accurate analysis, please consult with a NetApp Exchange specialist to conduct a detailed sizing exercise tailored to your specific configuration requirements.

## **6 ACKNOWLEDGEMENTS**

The authors would like to thank Keith Griffin for setting up the test beds used in this report and Kesari Mishra and Dave Tanis for insightful discussions.

## 7 REFERENCES

[1] Boost Performance Without Adding Disk Drives

<http://www.netapp.com/us/communities/tech-ontap/pam.html>

[2] Microsoft Exchange Server Jetstress

<http://technet.microsoft.com/en-us/library/bb643093.aspx>

[3] Microsoft Exchange Solution Reviewed Program (ESRP)—Storage v1.2

<http://technet.microsoft.com/en-us/exchange/bb412165.aspx>

[4] NetApp FAS3170 60,000-Mailbox Fibre Channel Storage—Solution ESRP

<http://www.netapp.com/us/solutions/solution-partners/global-alliance/ms-esrp-fas3000.html>

[5] Planning Storage Configurations

<http://technet.microsoft.com/en-us/library/bb124518.aspx>

[6] *Storage efficiency or capacity utilization* is defined as: (total DB size) / (total raw storage capacity – snapshot reserve); where, *total raw storage capacity* is defined as: (total number of DB and Log drives) \* (total capacity per drive), and *snapshot reserve* in this case was 20% of total raw storage capacity. The achieved 53% storage efficiency is noteworthy because any Exchange storage solution based on RAID 10 would have a theoretic limit of storage efficiency (or space efficiency) of 50% [7].

[7] RAID

<http://en.wikipedia.org/wiki/RAID>