# NetWorker Usage Report 2015

© Preston de Guise

March 2016

## Overview

The purpose of the NetWorker Usage Survey is to gauge high level details of how EMC NetWorker is used within the community, and to report trends on the usage.

The survey was conducted between December 1, 2015 and January 31, 2015, and follows previous surveys conducted in 2014, 2013, 2012, 2011 and 2010 (June and November). There were 140 respondents to the 2015 survey.

The survey does not force answers for individual questions, so in some results there may be fewer answers than the number of respondents.

## About the Author

Preston de Guise is a long term data protection expert with a career focus on enterprise backup and recovery solutions. Preston is the author of "Enterprise Systems Backup and Recovery: A corporate insurance policy" (CRC Press, 2008, 978-1420076394) and the upcoming title, "Data Protection: Preventing Data Loss in the Age of Big Data, Cloud and Virtualization".

Preston has worked on and developed backup solutions in most industry verticals, covering the full range of businesses from SOHO through to Global Fortune 500 companies.

Preston is currently a pre-sales systems engineer for Data Protection Solutions at EMC, and is based in Melbourne, Australia. *This survey is conducted independently of that role*.

# **Table of Contents**

1	Ne	tWorker Server Version	6
	1.1	Responses	6
	1.2	Findings	6
2	Nu	mber of NetWorker Datazones	8
	2.1	Responses	8
	2.2	Findings	8
3	Tot	tal Client Count - All Datazones	9
	3.1	Responses	9
	3.2	Findings	9
4	Ne	tWorker Server Operating System	10
	4.1	Responses	10
	4.2	Findings	10
5	Ne	tWorker Client/Storage Node Operating Systems	13
	5.1	Responses	13
	5.2	Findings	13
6	Bu	sinesses Using Deduplication	15
	6.1	Responses	15
	6.2	Findings	15
7	Ne	tWorker and Data Domain Modules/Plugins	18
	7.1	Responses	18
	7.2	Findings	19
8	Vir	tualisation within the Environment	20

8.	1	Responses	20
8.	2	Findings	20
9	Bac	ckup to Disk Technology	21
9.	1	Responses	21
9.	2	Findings	21
10	D	Do you clone within your environment?	22
1	0.1	Responses	22
1	0.2	Findings	22
11	L	ongest Retention Time	23
1	1.1	Responses	23
1	1.2	Findings	23
12	В	Backup Encryption	24
1:	2.1	Responses	24
1:	2.2	Findings	24
13	L	ongevity of NetWorker Use	25
1	3.1	Responses	25
1	3.2	Findings	25
14	Ν	New NetWorker Feature Usage	26
1	4.1	Responses	26
1	4.2	Findings	26
15	D	Dedicated Backup Administrators?	28
1	5.1	Responses	28
1	5.2	Findings	

16	Hov	w are virtual machines being backed up?	29
16.	.1	Responses	29
16.	.2	Findings	29
17	Net	tWorker Licensing Model	30
17.	.1	Responses	30
17.	.2	Findings	30
18	ls ta	ape still in use?	32
18.	.1	Responses	32
18.	.2	Findings	32
19	Но	w big is a full backup of your environment?	34
19.	.1	Responses	34
19.	.2	Findings	34
20	Bus	siness Geographies	35
20.	.1	Responses	35
20.	.2	Findings	35
21	Clo	oudBoost	36
21.	.1	Responses	36
21.	.2	Findings	36
22	Cor	nclusions	37

## 1 NetWorker Server Version

### 1.1 Responses

Responses showed a majority of businesses are continuing to run a supported version of NetWorker. Since many organisations have more than one installation of NetWorker in their environment, this question allowed for multiple responses. The responses were as follows:



## 1.2 Findings

82% of respondents were running on the NetWorker 8.x tree, with an overwhelming 56% of respondents running on NetWorker 8.2.x. While many years ago it was common to find a larger percentage of respondents running older versions of NetWorker compared to the current release, this shows a continuing tendency<sup>1</sup> for businesses to be on a current release of NetWorker. It's worth noting even though NetWorker 9 had barely been released at the time the survey was running (becoming generally available at the end of December), there were already businesses evaluating or upgrading to that version.

If we look at NetWorker server versions in use by geography, we observed the following:

6.x 7.4.x 7.5.x 7.6.x 8.0.x 8.1.x 8.2.x 9.x Region 1 1 4 12 24 3 Americas Only 7 APJ Only 1 1 3 EMEA Only 2 3 1 1 20 26 103 3

<sup>1</sup> In the 2014 survey we saw 70% of respondents on the NetWorker 8.x tree.

Region	6.x	7.4.x	7.5.x	7.6.x	8.0.x	8.1.x	8.2.x	9.x
Americas/EMEA		1	1			1	2	1
EMEA/APJ								1
All Geos				2		4	7	

Looking at respondents *only* operating in EMEA (the largest group of respondents), this distribution looks like the following:



Figure 2: EMEA (only) NetWorker Server Versions

## 2 Number of NetWorker Datazones

#### 2.1 Responses

Less than 50% of respondents had a single NetWorker datazone in their business.





#### 2.2 Findings

It's reasonably common to encounter multi-datazone environments. Reasons can include:

- Internal vs DMZ layouts
- Lab environments
- Multiple geographic regions
- Different business groups
- Handling very large numbers of clients
- Any mix of the above.

The presence of multiple datazones within organisations is a product strength - businesses are able to manage the administration and operation of more than one datazone, and are working at a higher level than "one server only" in order to build a more flexible data protection system. This doesn't equate to decentralised backup management, but more to increased deployment flexibility and sophistication.

Overall it would appear the number of businesses running a single datazone is staying reasonably static - in the 2014 survey this was down to 46%, but in the 2013 survey it was 48% Similarly, the number of businesses running 6 or more datazones is remaining reasonably stable at 23% (2014 - 25%, 2013 - 23%).

## **3** Total Client Count - All Datazones

#### 3.1 Responses

This question used the same breakdown as previous years, and the responses were as follows:



Figure 4: Total client count, all datazones

### 3.2 Findings

Overall we're seeing reasonably similar percentages year on year in the number of clients being protected using NetWorker. The 2014 survey saw 30% of respondents backing up 1,001 or more clients, as did this year's survey.

Based on the distribution of respondents who were able to provide a region of operation, we saw a distribution of client numbers to operating region as follows:

Region	1- 10	11- 25	26- 50	51- 100	101- 250	251- 500	501- 1000	1001- 2000	2001- 5000	5001+
Americas	2	2	1	7	1	10	5	3	7	
Only										
APJ Only	1	1			1	3	3			
EMEA	2	3	3	11	15	7	13	7	8	5
Only										
Americas/					1			1	2	
EMEA										
EMEA/					1					
APJ										
All Geos						1	1	1	3	4

Given the available sample size, the largest numbers of client deployments for a business are likely to occur in businesses operating in *all* geographies. That being said, there's a healthy spread of client counts in each region based on the respondents.

## 4 NetWorker Server Operating System

#### 4.1 Responses

This question focuses on the operating system used for NetWorker servers. Keeping in mind respondents could indicate they had multiple datazones, this yields a higher number of individual responses than survey respondents.



# Windows continues to be the dominant NetWorker server platform when all versions are

Figure 5: NetWorker Server Operating System

evaluated, with 40% of respondents running a NetWorker server on that platform. Linux however is closer this year at 38%. Looking at our ongoing trends, we can see percentages of server platforms as follows:

Survey	Windows	Linux	Solaris	Rest
Mar 2010	29%	22%	43%	6%
Nov 2010	29%	19%	43%	9%
Jun 2011	38%	23%	33%	6%
Dec 2012	41%	30%	20%	9%
Dec 2013	42%	28%	18%	12%
Dec 2014	43%	31%	17%	9%
Dec 2015	40%	38%	16%	6%



Figure 6: NetWorker Server OS trends

With NetWorker servers predominantly running on Windows and Linux, with a shrinking Solaris footprint and a minimal "other" footprint (typically AIX, HPUX and formerly Tru64), the decision was made to move the NetWorker server platform to being dedicated to Windows 64-bit and Linux 64-bit platforms with NetWorker 9.

Taking a look at the server operating systems by respondent region, we can see the following distribution<sup>2</sup>:

Region	Windows	Linux	Solaris	AIX	HPUX
Americas Only	19	22	7	2	0
APJ Only	7	3	1	1	0
EMEA Only	30	44	20	6	2
Americas/EMEA	1	2	1	2	0
EMEA/APJ	0	1	0	0	0
All Geos	4	7	5	1	0
Not Specified	1	1	1	0	0

<sup>&</sup>lt;sup>2</sup> Sub-classifications of operating systems (e.g., Linux Commercial vs Linux Free) have been condensed to a single classification.

With this breakdown we see starkly how little AIX and HPUX server adoption remains on a perregion basis. Solaris still has a more substantial following than the other two platforms combined in all regions, but given the position it came from originally as the ascendant NetWorker server operating system, this is not surprising.

Per-region this also helps us to see the growing popularity of Linux as a NetWorker server platform. While Windows beat Linux past the post for first position again this year when we split out the various operating system categories and count each category individually, it was by a margin that has continued to dwindle (86 to 83, 40% to 38%). If we count just by overall OS platform regardless of the *number* of server installs, Linux has jumped past Windows as the most popular NetWorker platform. This is further backed up by Figure 6: NetWorker Server OS trends.

Since one of the options in NetWorker 9 is NetWorker Virtual Edition (NVE), which runs on Linux, I expect we'll continue to see growth in the Linux/NetWorker server platform for a while yet to come.

#### **NetWorker Client/Storage Node Operating Systems** 5

#### 5.1 Responses

While not all operating systems can run NetWorker as a storage node, the two types were combined to avoid confusion between NetWorker servers and storage nodes. Also note that while we evaluate whether businesses are running a NetWorker server on a free vs a commercial version of Linux, we do not make the distinction for clients and storage nodes - instead the emphasis is on the broader type of Linux being run. This was the first year we included Linux on z/OS as a specific option. For ease of consideration however, Windows is treated as a single OS.



#### 5.2 Findings

We're still seeing a broad distribution of operating systems in the NetWorker ecosystem. Overall we're seeing a reasonably stable client deployment rate on Windows, and a slowly climbing deployment rate on Linux at the expense of all other operating system types. Over the course of the surveys the overall distribution has been:

Survey	Windows	Linux	Solaris	Other
Mar 2010	28%	23%	21%	28%
Nov 2010	28%	23%	22%	27%
Jun 2011	27%	22%	20%	31%
Dec 2012	29%	29.5%	17.5%	24%
Dec 2013	31%	32%	17%	20%



Figure 8: NetWorker Client and Storage Node OS Types

## 6 Businesses Using Deduplication

#### 6.1 Responses

This question confirms whether or not businesses are using deduplication, and if so, what type(s) of deduplication they are using.





### 6.2 Findings

This year saw a slight increase (22% in 2014, 24% in 2015) in the number of businesses not using deduplication. This does not necessarily represent an actual plateauing of deduplication adoption, particularly when we consider the refresh cycles many businesses have on data protection environments



Figure 9: Deduplication use in the environment

(5+ years). On top of a poor economic forecast in many geographies many businesses have used slower refresh cycles across a lot of their infrastructure, and some smaller businesses are stretching that even further by moving some of their workloads into public cloud services. This year however we also saw a shrinking of both source only deduplication *and* target only deduplication, with those respondents (mostly) moving into a combined source and target deduplication.

The year on year findings have been as follows:

Survey	None	Source Only	Target Only	Source & Target
Nov 2010 <sup>3</sup>	68%	4%	20%	7%
Jun 2011	64%	5%	27%	4%
Dec 2012	37%	5%	31%	27%

<sup>3</sup> Question was not asked in the March 2010 survey.

Survey	None	Source Only	Target Only	Source & Target
Dec 2013	27%	8%	31%	34%
Dec 2014	22%	5%	28%	45%
Dec 2015	24%	3%	18%	56%

The follow-up question asks respondents to provide details of the *type* of deduplication technology they're using – differentiating between the various options available around Data Domain, as well as Avamar and non-EMC deduplication technology.

#### Deduplication Type Number



Non-EMC deduplication continues to rank low - 7% this year down from 8% last year. This is undoubtedly due to the exceptional levels of integration offered by Data Domain within a NetWorker (and data protection) environment. Last year's survey saw 80% of deduplication users leveraging Data Domain within within their environment; this year that figure is at 81%.



Figure 10: Deduplication types

Examining deduplication adoption per region, we see the following results:

Region	Source and Target	Source Only	Target Only	No Deduplication
Americas Only	17	2	8	12
APJ Only	6	0	0	3
EMEA Only	46	2	12	15
Americas/EMEA	3	0	0	1
EMEA/APJ	1	0	0	0
All Geos	4	0	5	1
Not Specified	1	0	0	1

If we consider deduplication adoption based on the number of overall clients being protected by NetWorker, we can see the following:

# Clients	Source and Target	Source Only	Target Only	No Deduplication
1-10	1	1	0	3
11-25	3	0	1	2
26-50	1	0	1	2
51-100	8	0	3	9
101-250	9	0	3	7
251-500	13	1	4	4
501-1000	13	0	5	4
1001-2000	8	2	1	1
2001-5000	15	0	4	1
5001+	7	0	3	0

. . . . .

Viewing deduplication adoption against client counts allows us to draw a couple of distinct inferences:

- The more clients a business is protecting, the more likely they are to be using ٠ deduplication. Looking at businesses protecting 501 or more clients:
  - 90.6% indicated they were using deduplication
  - 0 70% were using source and target or source deduplication
- Even the smallest of businesses can see some benefit from (and therefore reason to deploy) deduplication

What I interpret out of this is simple: deduplication substantially solves scaling issues when it comes to backup and recovery, which is the storage node problem. Other backup products scale by having businesses deploy media servers by the dozens - sometimes even the hundreds - and require large network funnels into those storage nodes for data processing<sup>4</sup>.

Source-side deduplication solves the conventional scaling issue and costly investment challenge by distributing the deduplication workload throughout the entire protected environment. This lets businesses scale where they need most: out for the size of the protected environment, and up for the size of the protection storage (thereby reaping maximum benefit from a global deduplication pool), eliminating or substantially reducing a largely unnecessary middle layer in the backup topology.

<sup>&</sup>lt;sup>4</sup> While NetWorker still completely supports storage node models, I feel the use cases for storage nodes as "workhorses" in a backup environment are substantially reducing when the environment uses fully integrated data protection storage appliances.

## 7 NetWorker and Data Domain Modules/Plugins

#### 7.1 Responses

Continuing from last year, this survey asked users for both NetWorker and Data Domain plugin usage.



Figure 11: NetWorker Module and Data Domain Plugin Usage

## 7.2 Findings

The modules for Enterprise MySQL, Documentum, Meditech and ProtectPoint were all excluded from the graph as they contributed less than 1% of adoption. This is not surprising – Meditech is a highly specialised application used by a single industry. The Documentum module has been phased out in place of alternatives from specialist third party providers, and the MySQL module is limited to only those businesses running Oracle's Enterprise version of MySQL.

Given ProtectPoint support requires a VMAX3 and preferably NetWorker 9, a single-respondent adoption rate within a month (at most) of NetWorker 9 being released amongst 140 survey respondents is unsurprising, though this will likely grow in the coming year as functionality and platforms evolve.



Figure 12: Use of open source databases

Only 11% of respondents indicated they were using Enterprise MySQL. The number of respondents *not* using any open source database within their environment remains steady at 40% compared to last year's survey, though is likely to be reflective of the *believed* vs *actual* percentage: many packages and tools making their way into IT departments leverage one open source database or another (MySQL and PostgreSQL being the dominant two).

(As was noted in previous surveys when the question was asked, a large number of businesses that use open source databases within their environments indicated a lack of desire to pay for a backup module for the database.)

Zero users of SnapImage is not surprising - the module has been depreciated for some time and Block Based Backup options now exist natively for Windows and Linux NetWorker clients. SnapImage will be removed from future surveys.

## 8 Virtualisation within the Environment

#### 8.1 Responses

Virtualization is subsuming the datacentre. We've ceased talking about *just* virtual hosts; now the entire stack from networks through to the datacentre itself can be virtualized. The pervasiveness of virtualization as it applies to data protection is tested in this question.



systems within their environments; this year that was 5%.

Figure 13: Virtualization within the data protection environment

The percentage of respondents

indicating "some clients" remained static between the two surveys at 47%. We have seen an increase in the percentage of servers and storage nodes virtualized however:

Yoar	Some Storage Nodes	All Storage Nodes	Some Servers	All Sarvars
rear	Some Storage Nodes	All Stolage Nodes	Some Servers	All Servers

2014	10%	5%	10%	9%
2015	14%	3%	10%	12%

While a small drop can be seen in the 'All Storage Nodes' option, both 'Some Storage Nodes' and 'All Servers' increased year on year. The increased virtualization levels in particular of NetWorker servers is not surprising; with increasing adoption of backup to disk and a move to 'director mode' only NetWorker servers, there are little benefits to having a physical NetWorker server - and in fact, virtualization of the NetWorker server results in better options for disaster recoverability and higher availability for the datazone.

## 9 Backup to Disk Technology

#### 9.1 Responses

This question is a little broader than previous questions on deduplication, and looks at the three major types of backup to disk technology in use with NetWorker today, viz.:

- Advanced File Type Devices
- Virtual Tape Libraries
- Data Domain Boost

Technology Number



### 9.2 Findings

There are slight variations between the 2014 and 2015 survey, showing a decrease in the use of non-deduplicating backup to disk (2014: 21%, 2015: 19%), a decrease in VTL use (2014: 18%,



#### Figure 14: Backup to disk technology

2015: 16%), and a decrease in the number of environments not using backup to disk. The percentage of environments not using backup to disk has now remained reasonably stable at the 8-9% mark since 2012.

Survey	None	ADV	VTL	Boost
Nov 2010 <sup>5</sup>	16%	52%	32%	N/A
Jun 2011	15%	47%	33%	5%
Dec 2012	8%	38%	24%	30%
Dec 2013	8%	34%	18%	40%
Dec 2014	9%	31%	18%	43%
Dec 2015	8%	31%	16%	45%

Boost adoption as the preferred backup to disk technology continues to rise year on year. Use of ADV\_FILE technology has remained static between 2014 and 2015.

<sup>&</sup>lt;sup>5</sup> Question was not asked in the May 2010 survey.
© Preston de Guise

## 10 Do you clone within your environment?

#### 10.1 Responses

As per last year's survey, this question explicitly refers to both cloning and *clone controlled replication* in case newer NetWorker sites might otherwise see them as different operations.



trend in the number of businesses cloning their

backups - 76% in 2012, 78%

Figure 15: Cloning within the environment

in 2013, 80% in 2014 and now 83% in 2015. This is an excellent trend - backups are a critical part of data protection infrastructure, and should never represent a single point of failure.

Equally, the number of businesses cloning *all* backups continues to rise: 23% in 2013, 28% in 2014 and 34% in 2015. Hopefully this trend will continue: while cloning only production or selective backups may seem cost effective, it still introduces a degree of risk into the overall data protection environment. (The growth in all-backup cloning may also be reflective of the effectiveness of Data Domain deduplicated replication, too.)

#### **Longest Retention Time** 11

#### 11.1 Responses

This question focuses on the longest retention period for any backups within the organization, and is not necessarily representative of the overall average retention times. This is useful to understand where backup is being used to achieve compliance retention.



### 11.2

a longest retention period of 1 year or less. Just over 50% of businesses have a longest retention period of 3 years or more.

Long retention times are usually an indication of the

backup environment being used for compliance retention - e.g., financial or other regulatory requirements. This speaks to a continuing need for businesses to consider adopting information lifecycle management: data archive should by and large be the enabler of compliance retention rather than backup software.

14%

Figure 16: Longest retention time

## **12 Backup Encryption**

#### 12.1 Responses

This topic covers both whether or not a business has any *requirement* to use encryption, and so, whether it's been implemented – and how.

nber

104	Not required
8	Required (in-flight) Not implemented
9	Required (at-rest) Not implemented
8	Implemented (in-flight)
23	Implemented (at-rest)

Answer Number

7

13

10

13

3

Client side

(software)

(in-flight/Boost)

At rest (disk)

At rest (tape)

Routers

**Findings** 

We continue to see environments where some form of encryption (in-flight or at-rest) is required, but not implemented – just over 10% of respondents gave this

requirement. In-flight encryption

FC/IP Encryption

12.2

Client







Figure 18: Encryption in use

via Data Domain boost remains a popular option for those businesses implementing encryption in transit - 28% for 2015 (31% for 2014).

## 13 Longevity of NetWorker Use

#### 13.1 Responses

This question gauged how long NetWorker had been installed within environments.

Time	Number
<1 year	5
1-2 years	8
2+ to 5 years	30
5+ to 10 years	34
10+ to 15 years	29
15+ years	34

### 13.2 Findings

2015 saw NetWorker turn 25 years old<sup>6</sup>, and the survey shows us many businesses have had a long association with it - almost 25% of respondents have been using NetWorker for 15 or more years - and 45% of respondents have been using NetWorker for 10 years or more.



Figure 19: Longevity of NetWorker in Businesses

NetWorker proves itself year-in, year-out to be a trusted part of critical business infrastructure and data protection solutions.

<sup>&</sup>lt;sup>6</sup> http://nsrd.info/blog/2015/06/30/celebrating-25-years-of-networker/

## 14 New NetWorker Feature Usage

#### 14.1 Responses

This question probes the use of advanced new features in NetWorker - some introduced in the 8.x tree, and others introduces in the 9 tree.

Feature	Number
Unix Parallel Streaming	45
Windows Parallel Streaming	43
Windows BBB	43
Linux BBB	4
Synthetic Full	17
Virtual Synthetic Full	25
VBA	45
Storage node balancing	14
Optimised backup of windows deduplicated filesystems	4
Boost over FC	13
NMC recovery interface	50
Maintenance mode	11
NMM BBB	6
NSM	13
ProtectPoint	1

#### 14.2 Findings

A broad spread of new backup options are getting adoption within the NetWorker community note in particular 4 respondents already using Linux Block Based Backups (requiring NetWorker 9) - per the first response, 6 respondents had adopted NetWorker 9 into their environment.



Figure 20: NetWorker feature usage

## **15 Dedicated Backup Administrators?**

#### 15.1 Responses

This question covered whether or not respondents had dedicated backup administrators in their environment.

Response Number



#### 15.2 Findings

This is reasonably similar to last year's results: just 1 per cent higher on the Yes and 1 per cent lower on the No.

Stepping back from this question directly it's worth considering changes happening within the IT environment.

There are two directions backup administration is taking that we should be aware of:



Figure 21: Dedicated backup administrators?

- A hybrid approach to backup topology is giving application and virtual machine control backup to their respective administrators while having backup administrators set broad policy direction and be responsible for *protection storage*
- The appeal of converged and hyper-converged infrastructure within IT environments during refresh cycles is changing the nature of administration.

For the first, we'll still see people calling themselves *backup administrators* for some time to come. But the converged and hyper-converged market is leading to a new breed of administrator within the IT environment: the *infrastructure* administrator. In CI and HCI environments, with each component closely aligned (particularly so in HCI), having administrators that can deal with the entire infrastructure stack - the hardware blocks, virtualization *and* data protection thereof is critical. The biggest mistake a business can make, after all, is deploying CI or HCI and still divvying up tasks per "classic" technical team - network, storage, virtualization, data protection.

The adoption curve for CI/HCI is not instantaneous, though the drive towards cloud-like agility even in on premise environments will create pushes here. The market is a growing one and that growth is likely to continue for some time to come.

Even those businesses *not* directly adopting CI or HCI will look towards the agility offered by those platforms and seek to standardise, in some way or another, their infrastructure stacks to highly automated, low-maintenance systems, and allocate FTE resources aligned to the *entire* stacks.

#### How are virtual machines being backed up? 16

#### 16.1 Responses

While a previous question dealt with the extent of virtualisation within environments, this guestion focused on how virtual environments are being backed up. Remembering that a mix of options can be used depending on NetWorker versions and the presence (or lack thereof) of virtualised databases, this was a multiple-choice question. Since this is focused on the ways virtual machines are backed up, the option of "no VMs" is removed from the results as it was supplied to encourage all respondents to answer.



#### 16.2

VBA adoption has grown between this year and the last (20% 2014, 22% 2015), and the older backup integration options - VCB and VADP have both shrunk year on year (VCB: 3% 2014, 2% 2015; VADP: 24%



2014, 16% 2015). We have seen a slight increase in the percentage of third party backup tool usage - from 8% to 10%.

Perhaps representative of the increasing levels of virtualization in many businesses, the number of client/in-guest backup options have increased - these are most usually deployed in virtualized environments to support databases.

We're not there yet: virtual machine backups as image level backups at the hypervisor while not new, are new enough that they require tighter integration between different infrastructure teams - particularly backup and virtualization. As these groups integrate in IT environments and virtual backup options increase, expect to see higher growth in the VBA rates over successive surveys.

(With an increasing number of virtualization technologies entering the marketplace each year, we also see a catch-up game between new hypervisors and backup product support for those hypervisors.)

## 17 NetWorker Licensing Model

#### 17.1 Responses

There are three specific licensing models now available to NetWorker environments.



### 17.2 Findings

This was the first year this question was asked, so there are no trends or correlations to be drawn.

It is gratifying to see more businesses on NetWorker Capacity licensing than NetWorker Traditional Licensing. While the old per-feature traditional licensing model allowed a high level of granularity, it came at the cost of agility - if the business needed to implement a new database,



Figure 23: NetWorker Licensing Model

application or storage system, data protection had to wait until license budget and then licenses were obtained.

The capacity models - both NetWorker and DPS - represent considerably more agile options: a licensed size based on the amount of front end TB protected by the environment, with *all* features enabled. This instantly grants agility an order of magnitude higher than traditional licensing.

Given the rapidity of which infrastructure can be deployed within virtualised environments, this sort of capacity based licensing model makes more sense – both for the business and for the backup administrators who would otherwise need to manage large numbers of per-feature licenses. (Businesses have literally shrunk their licenses in NetWorker from *hundreds* to less than 5 by transitioning to capacity licensing.)

DPS for Backup is, compared to NetWorker Capacity licensing, still a newer option, and we have to keep in mind it includes Avamar licensing as well (so there would be owners of DPS for Backup licensing not actually making use of NetWorker). That being said, as businesses enter refresh or maintenance renewal periods, seriously evaluating the benefits of DPS for Backup over traditional or NetWorker Capacity licensing is worthwhile. The current feature list includes:

• Avamar

© Preston de Guise

- NetWorker
- Data Domain Boost for Enterprise/Microsoft Applications
- Data Protection Search
- Data Protection Advisor
- CloudBoost

Businesses need a flexible licensing structure that gives them not only backup and recovery, but *also* monitoring, reporting, search, Cloud extension *and* options to shift application backup/recovery to subject matter experts. DPS for Backup does give that flexibility within a single licensing option.

## 18 Is tape still in use?

#### 18.1 Responses

This question focused on whether tape was still in use *at all* within environments, and if so, what it was used for. Multiple-choice was used to provide better granularity on use cases.



Figure 24: Tape usage

While the number of respondents using backup to disk is high (92%), tape is still being used for workloads in many organizations, both for backup and for cloning. As backup to disk technology continues to fall in price and the performance requirements for tape technology continues to increase, we would expect to see over successive surveys an increase in the "No" option above. The shift to fully tape-less backups has successfully happened for many organizations, but if we consider NetWorker's legacy – coming from a tape based background – it's a technology many administrators feel comfortable with. If we consider the number of environments using NetWorker for long term retention (50% of respondents hold backups for 3 years or more), we can assume that a lot of tape utilisation within NetWorker is driven by using the product for compliance/regulatory required retention periods.

This will be something to track in successive surveys.

Turning our attention to a regional perspective, the use of tape in NetWorker environments can be noted as follows:

drawn.

Region	Daily/Weekly Clones	Monthly/Longer Clones	Selective Clones Only	Daily/Weekly Backups	Monthly/Longer Backups	Selective Backups Only	No
Americas Only	9	9	0	12	7	2	14
APJ Only	2	1	1	3	2	3	3
EMEA Only	14	20	12	13	19	8	19
Americas/ EMEA	0	2	0	0	0	0	1
EMEA/ APJ	0	0	0	0	0	0	1
All Geos	4	6	1	4	4	2	2
None Specified	1	0	0	0	0	0	0

Adoption of tape-less backup solutions can be seen across all geographic regions. If we examine each of Americas, APJ and EMEA we see tapeless adoption rates of 26.4%, 20% and 18.1% respectively.

The transition to fully tapeless is still an ongoing process, driven by refresh cycles and retention periods requirements. Ironically, the true cost of actually using tape for long term retention is rarely counted, as evidenced by the number of businesses that have legacy tapes they keep for retention which are in a format they can no longer read<sup>7</sup>. Speaking simply: if a business still has long-term retention backups on tapes that can no longer be read, then retention is being "solved" only as a box-ticking exercise rather than a functional reality: this should be seen as a clear mark *against* tape for long term retention rather than a positive.

<sup>&</sup>lt;sup>7</sup> Anecdotally, every time I discuss long term retention with businesses who are using tape currently for the process they will admit at least one, if not two generations of tape media being held for which they no longer have the ability to read, short of sourcing parts through second-hand markets. *This is regardless of backup product used.* 

#### How big is a full backup of your environment? 19

#### 19.1 Responses

This focuses on the size of a single full backup of the entire environment. This is often a tricky question - calculating it is not always straight forward, dependent on fluctuations and system changes, so it is premised on being a guestimate.



#### 19.2 Findings

We're continuing to see a broad mix of data sizes being backed up to NetWorker. While 36% of respondents were backing up 50TB or less,

Figure 25: Estimate of front end capacity (TB)

34%

17% of respondents were using NetWorker to backup a petabyte or more. This goes to show that NetWorker is a flexible product able to be utilised in a large variety of businesses from the smallest to the biggest.

## 20 Business Geographies

#### 20.1 Responses

This was a new question designed to help build details *by geography* on other questions. Given many businesses operate in more than one geographic region, this was a multiple-choice question. (This question was added approximately 24 hours after the survey started.)



### 20.2 Findings

There were 138 respondents who indicates which geographic region they were operating in. Most of the purpose behind this question is to provide cross-referencing between region of operation and other data, but we can note the following composition of regions in respondents:



Figure 26: Geographic region (including multi-regional businesses)

Region	Respondents	Percentage
Americas Only	39	28.3%
Americas/EMEA	4	2.9%
Worldwide	10	7.2%
APJ Only	9	6.5%
EMEA Only	75	54.3%
EMEA/APJ	1	0.7%

## 21 CloudBoost

### 21.1 Responses

A few days after the survey kicked off, EMC product management requested I put in a question about CloudBoost, and wanting to

come up with a question quickly I used the following options:

Answer	Number
Don't know about it	49
Don't plan on using it	46
Planning on/currently trialling it	11
Using it for some long term retention	2
Using it for all long term retention	0

### 21.2 Findings

Given CloudBoost is primarily targeted (at this point) at DPS for Backup customers and 10% of respondents were using DPS, these numbers currently make sense given the overall number of



Figure 27: CloudBoost

respondents. It does show we need to demonstrate and promote CloudBoost. (I also believe CloudBoost v2, released only *during* the survey, will resonate more readily than CloudBoost v1.)

Using Cloud storage for long term retention *is* however becoming a regular and genuine conversation with customers, but the many businesses are having this discussion on the basis of *testing out* Cloud strategy options, which means they're taking a careful approach and spending some time planning the process. Additionally, it's usually being considered in conjunction with upcoming tape infrastructure maintenance refreshes, so the combined requirements for LTR to the Cloud is seeing a stretched out early adoption period. Based on the conversations I'm having almost every day with customers about using Cloud storage for long term retention, I firmly believe we'll see these figures change next year.

## 22 Conclusions

This year saw a few additional questions in the survey, covering use of tape within environments, licensing, geographic region of operation and CloudBoost operations.

When we look at overall usage patterns with NetWorker, we see strong continuing trends, including:

- Large options in the use cases for:
  - Number of clients protected
  - Amount of data protected
- Decreasing use of Solaris, AIX and HPUX for the NetWorker server platform
- Broad operating system types for client and storage node systems
- Increased adoption in the combined Source and Target deduplication workloads, undoubtedly as a result of the distributed segment processing available in the embedded Boost API for NetWorker clients
- Increasing use of backup to disk
- Increasing use of cloning

With additional analysis we could also see this year that the more clients there are being protected in a NetWorker environment, the more likelihood of deduplication being used; reiterating the earlier comments on this: deduplication – particularly when distributed to the entire environment, including clients – substantially addresses scaling within backup environments.

While it was the first year the question was asked, we can see that a shift has happened towards capacity-based licensing for NetWorker, which is actually reflected across almost the entire backup industry. Businesses are no longer interested in minutely managing licenses on a perfeature basis, and instead what a flexible, adaptable license based on a simple metric: how much they backup.

NetWorker remains an extremely popular product, with more than 45% of respondents having used it for 10 or more years within their environments. This speaks considerably to the trust businesses place in it year in, year out.

As always, thanks to everyone who participated in the survey.