NetWorker Usage Report 2016

© Preston de Guise

March 2017

Overview

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

The survey was conducted between December 1, 2016 and January 31, 2016, and follows previous surveys conducted in 2015, 2014, 2013, 2012, 2011 and 2010 (June and November). There were 159 respondents to the 2016 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's published works include:

- "Data Protection, Ensuring Data Availability"
  - CRC Press, 2017, 978-1482244151
  - https://www.crcpress.com/Data-Protection-Ensuring-Data-Availability/Guise/p/book/9781482244151
- "Enterprise Systems Backup and Recovery: A corporate insurance policy"
  - CRC Press, 2008, 978-1420076394
  - https://www.crcpress.com/Enterprise-Systems-Backup-and-Recovery-A-Corporate-Insurance-Policy/Guise/p/book/9781420076394

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 employed as a senior pre-sales systems engineer for Data Protection Solutions at DellEMC, and is based in Melbourne, Australia. This survey is conducted *independently* of that role.

© Preston de Guise

Table of Contents

© Preston de Guise

© Preston de Guise

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:

Version	*Number*
6.x or Lower	1
7.0.x to 7.4.x	2
7.5.x	2
7.6.x	14
8.0.x	9
8.1.x	32
8.2.x	124
9.x	39

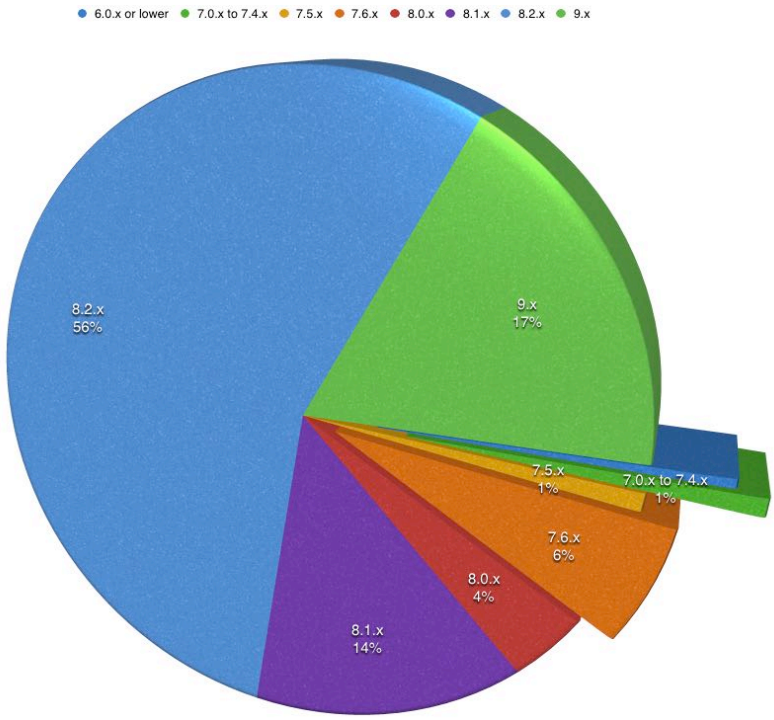

1.2 Findings

74% of respondents were running on the NetWorker 8.x tree, with 56% of respondents running on NetWorker 8.2.x. 17% of respondents had adopted NetWorker 9 into their environment. This survey did not split out NetWorker 9.0.x and NetWorker 9.1, since 9.1 was only released in late 2016, during the survey.

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

Region	*6.x or Lower*	*7.0.x-7.4.x*	*7.5.x*	*7.6.x*	*8.0.x*	*8.1.x*	*8.2.x*	*9.x*
Americas	0	0	1	1	1	1	19	4
EMEA	1	1	1	5	3	13	77	25
APJ	0	0	0	1	3	3	6	5
Two Regions	0	0	0	3	0	4	9	3
Global	0	1	0	4	2	11	13	2

As a stacked graph with 100% representing all versions installed within a single region, the per-region breakdown of NetWorker versions installed is as follows:

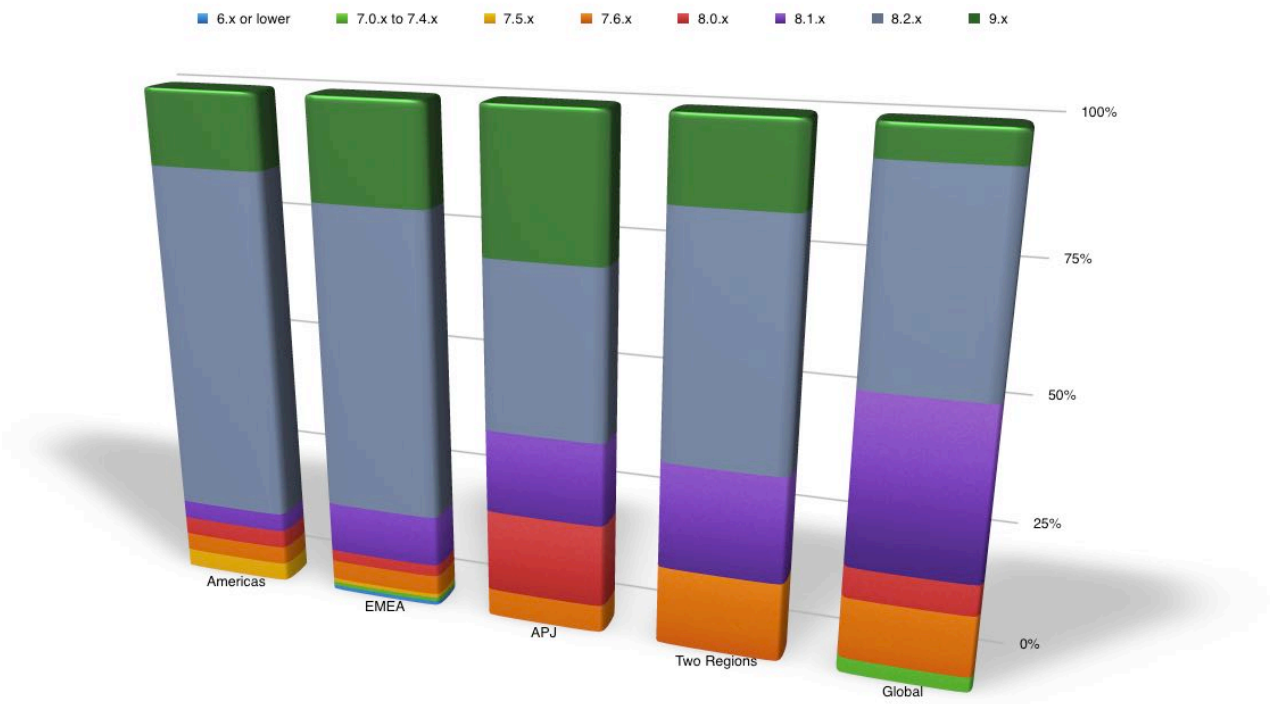

Looking at respondents *only* operating in EMEA (the largest group of respondents), this distribution looks like the following:

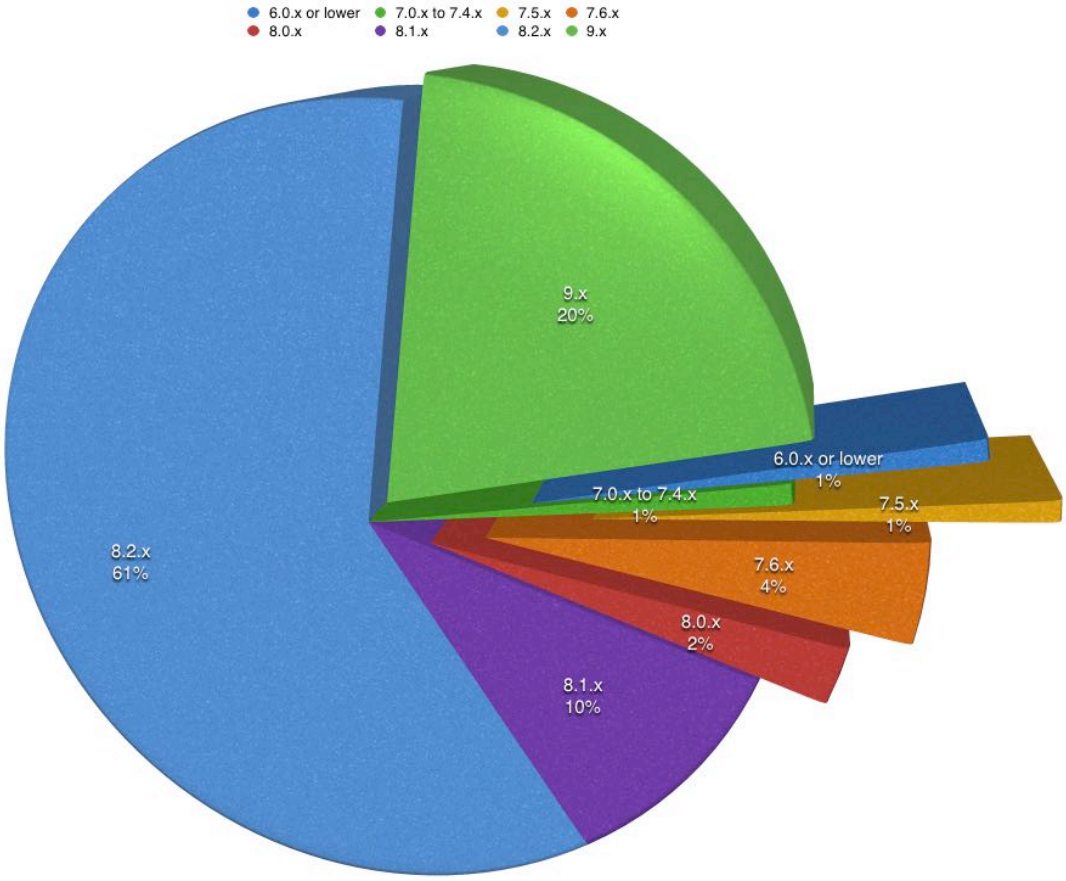

In the previous survey, NetWorker 9 adoption rate for EMEA was just 2%, having only been released around the time the survey was starting. In the time since then, NetWorker 9 adoption has grown to 20%. (Similarly, overall NetWorker 9 adoption has grown from 4% in the previous survey to 17% in this survey for all regions.)

© Preston de Guise

2 Number of NetWorker Datazones

2.1 Responses

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

# Datazones	Number
1	72
2	26
3	12
4	3
5	8
6-10	13
11-25	16
25-50	3
51-100	4
101+	2

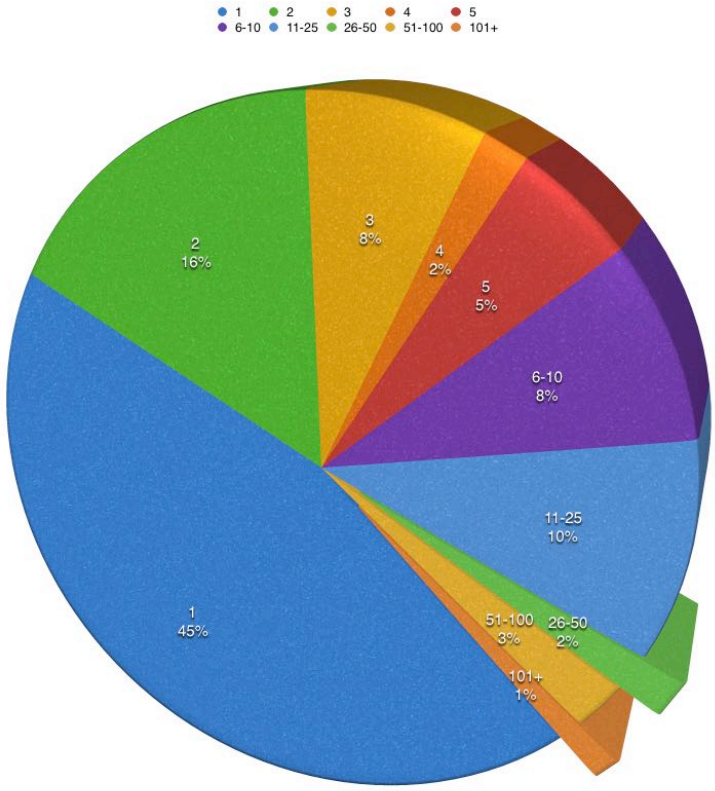

2.2 Findings

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.

Datazone numbers remain a reasonably static spread across surveys. The 2014 survey showed 46% of respondents had just one datazone, the 2015 survey reported 48%, and the 2016 survey shows 45%. Equally, the number of respondents with 6 or more datazones is remaining reasonably constant – 23% in 2013, 25% in 2014, 23% in 2015 and 24% in 2016.

Across the regions, the spread of datazone numbers resembles the following:

Region	1	2	3	4	5	6-10	11-25	26-50	51-100	101+
Americas	19	5	0	0	0	1	0	0	0	0
EMEA	46	17	11	1	1	6	10	2	3	0
APJ	4	2	0	0	2	0	1	0	0	1
Two Regions	2	1	1	0	1	5	1	0	0	0
Global	1	1	0	2	4	1	4	1	1	1

When viewed as a percentage of all datazones in each region, we see breakdowns of the likelihood of a region having a particular number of datazones. For these results, note how comparatively, respondents from the Americas were more likely to be running a single datazone, whereas all other regions were far more likely to have multiple datazones in use:

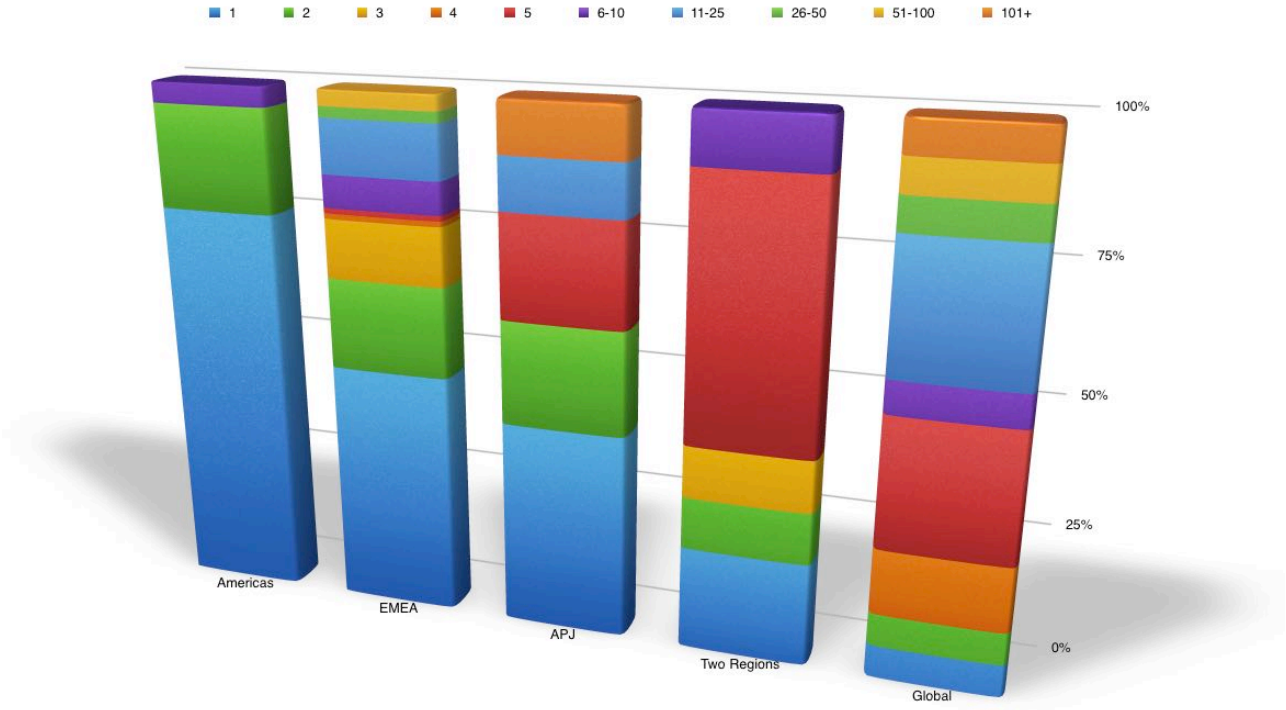

3 Total Client Count – All Datazones

3.1 Responses

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

# Clients	Number
1-10	2
11-25	7
26-50	8
51-100	17
101-250	28
251-500	3
501-1000	27
1001-2000	15
2001-5000	24
5001+	13

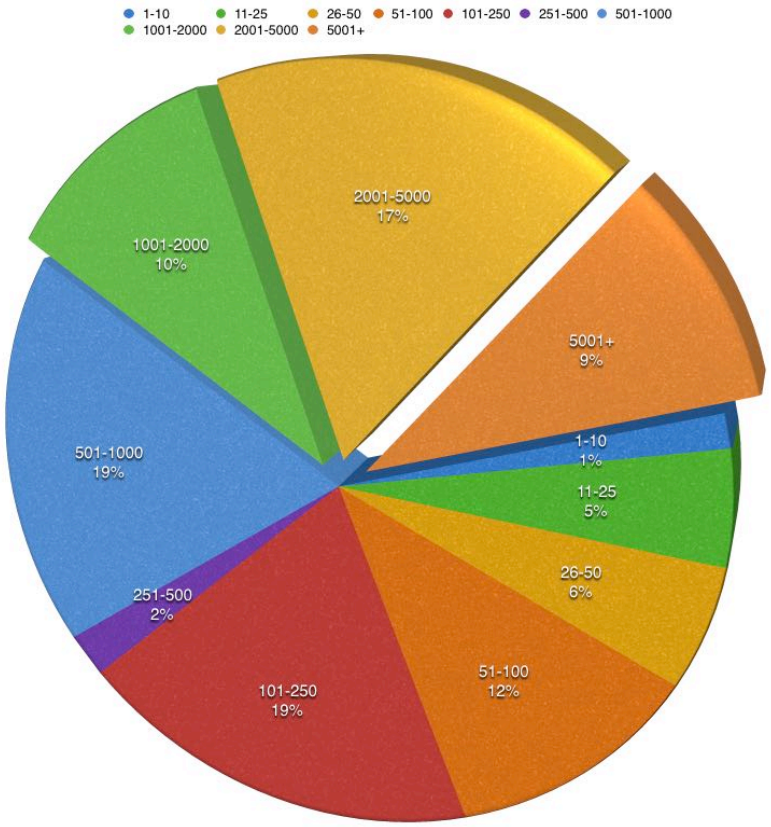

3.2 Findings

In the 2014 and 2015 surveys we saw 30% of respondents protecting 1,001 or more clients. With the 2016 survey this has risen to 36%, and may be indicative of larger numbers of virtual machines being deployed within most businesses.

The regional breakdown of client counts were as follows:

Region	1-10	11-25	26-50	51-100	101-250	251-500	501-1000	1001-2000	2001-5000	5001+
Americas	1	3	2	4	7	0	5	1	0	0
EMEA	1	4	5	11	13	3	16	13	13	7
APJ	0	0	0	0	5	0	2	1	0	1
Two Regions	0	0	0	2	3	0	2	0	3	1
Global	0	0	1	0	1	0	2	0	8	4

While the individual numbers of client counts per region become statistically quite small in places, a stacked graph showing the per-region breakdown is provided below.

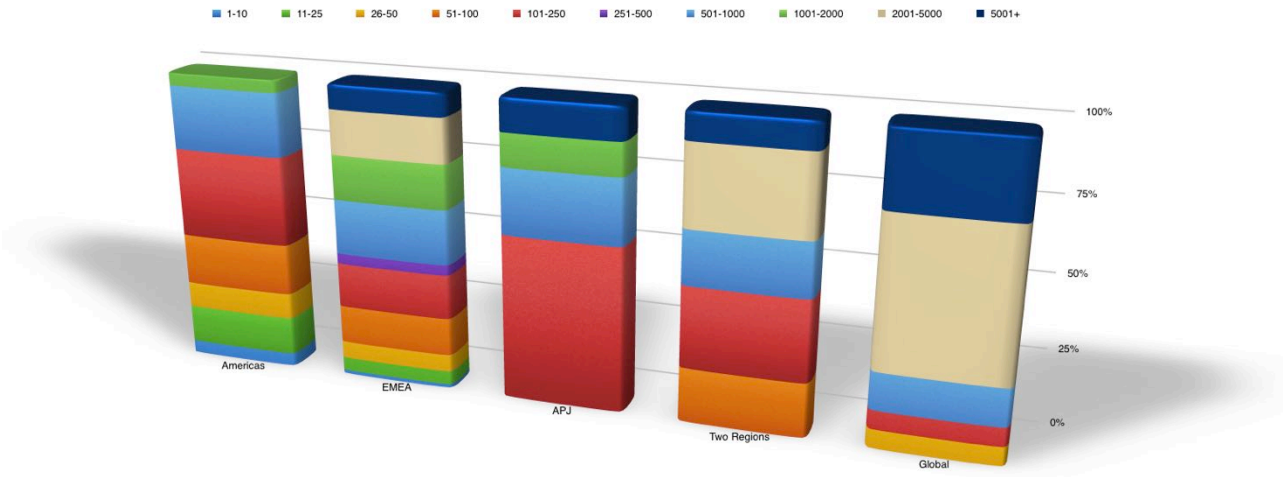

4 How big is a full backup of your environment?

4.1 Responses

This question focuses on the size of a single full backup of the entire environment. This is not always readily known by administrators and can fluctuate month by month depending on data growth, so we work on the principle of a guestimate. (This size effectively represents a "Front End TB" (FETB) capacity sizing.)

Size	*Number*
Unknown	23
<1 TB	2
1-50 TB	41
51-100 TB	23
101-500 TB	31
501 TB – 1 PB	17
1+ PB to 5 PB	14
5+ PB to 10 PB	3
10+ PB	5

("Unknown" sizes have been excluded.)

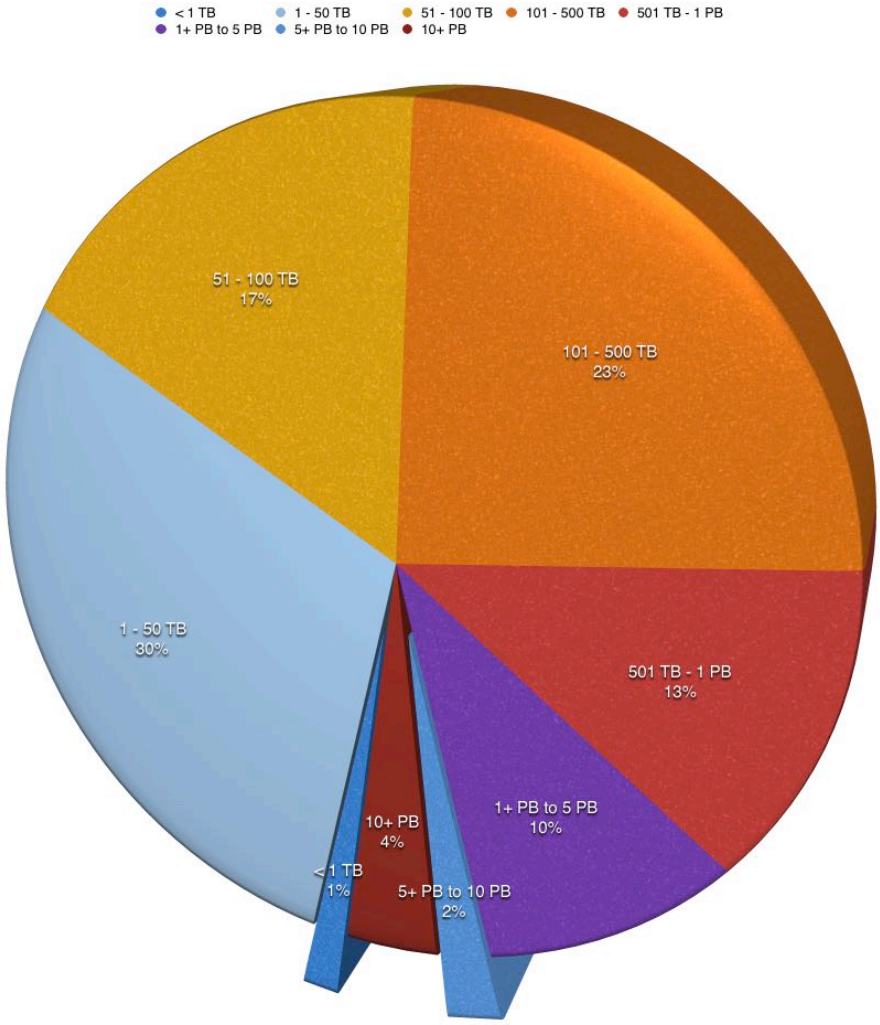

4.2 Findings

Per previous surveys, we continue to see broad mixes of data sizes being backed up with NetWorker. Last year's survey saw 17% of respondents protecting 1 PB or more, and this year that was slightly down at 16%, yet regardless it can be readily seen that NetWorker is capable of protecting significantly large environments.

On a per-regional basis, the breakdown of FETB sizes looks as follows:

Region	*Unknown*	*<1TB*	*1-50TB*	*51-100TB*	*101-500TB*	*501TB-1PB*	*1+PB-5 PB*	*5+ PB-10 PB*	*10+ PB*
Americas	5	1	9	4	2	3	1	0	0
EMEA	10	1	23	13	24	14	7	1	4
APJ	1	0	5	2	1	0	1	0	0

Region	Unknown	<1TB	1-50TB	51-100TB	101-500TB	501TB-1PB	1+PB-5 PB	5+ PB-10 PB	10+ PB
Two Regions	1	0	2	2	2	0	2	1	1
Global	6	0	2	2	2	0	3	1	0

In a stacked graph, we can see the per region distribution of FETB across respondents:

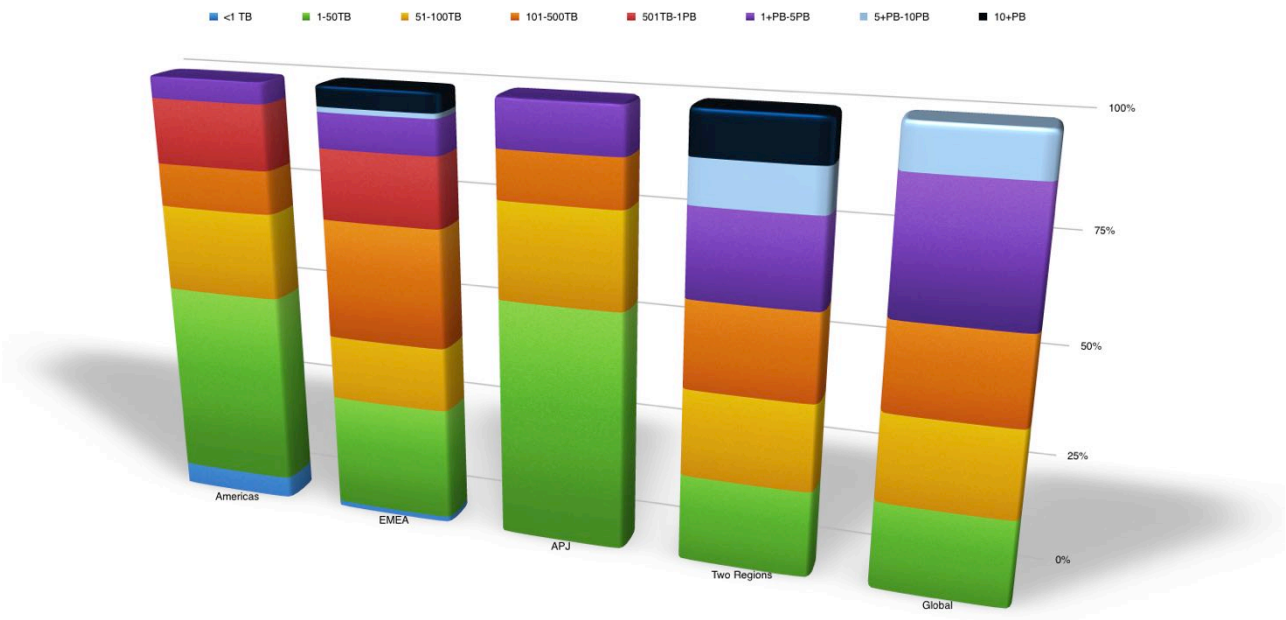

5 NetWorker Server Operating System

5.1 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.

OS	Number
Windows 2016	2
Windows 2012/R2	57
Windows 2008/R2	53
Linux (Commercial)	77
Linux (Free)	20
Solaris	22
AIX	11
HPUX	8
Tru64	1

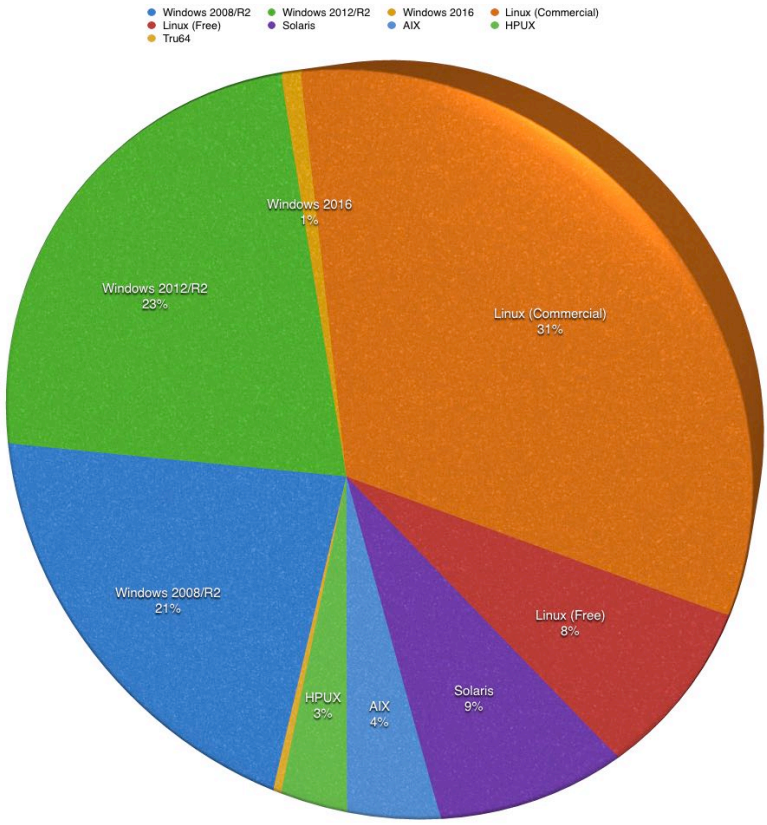

5.2 Findings

Windows continues to be the dominant NetWorker server platform when all versions are combined, with 41% of respondents running a NetWorker server on that platform. Linux however is closer this year at 40%. 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%
Dec 2016	41%	40%	10%	9%

© Preston de Guise

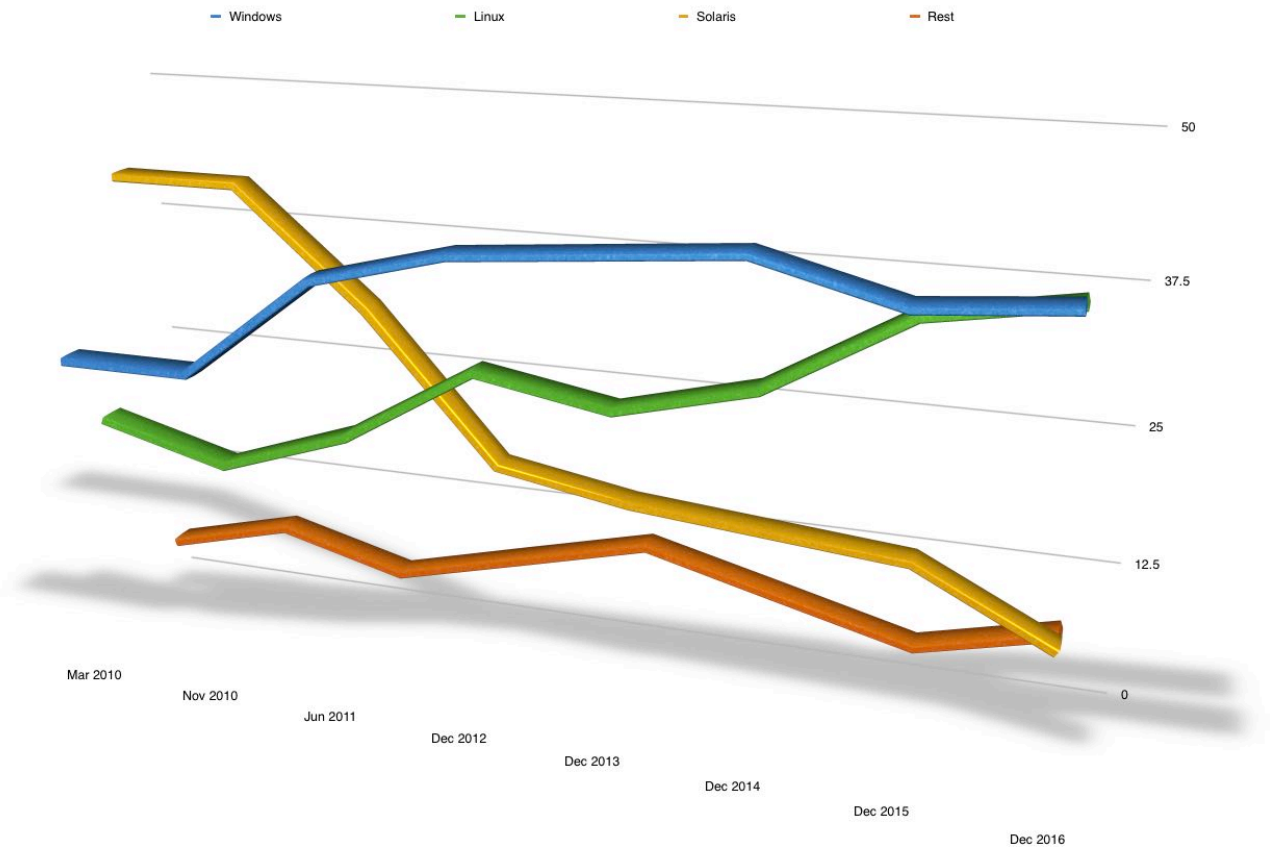

Oracle's purchase of Sun was completed in January 2010 and by the time of the June 2011 survey, the ascendency of Solaris as a NetWorker server platform was already slipping.

From NetWorker 9, only 64-bit Windows and Linux servers are supported as the operating system for the NetWorker server, and we will continue to see Solaris and "other" platforms (i.e., other Unix platforms) shrink as customers upgrade to current releases.

6 NetWorker Client/Storage Node Operating Systems

6.1 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. After requests from blog subscribers following last year's survey, Windows is now split out into all operating system variations.

In the graph on the following page, operating systems polling less than 1% have been removed. (This covers Windows NT4 and Lower, Tru64, Irix and NetWare.)

OS	Number
Win 2016	18
Win 2012/R2	137
Win 2008/R2	136
Win 2003	75
Win 2000	16
Win NT4 & Lower	3
Win Desktop	11
Linux (RPM)	124
Linux (non-RPM)	39
Linux (z/OS)	12
Solaris	58
AIX	34
HPUX	25
Tru64	1
Irix	0
Mac OS X	6
OpenVMS	7
NetWare	1

6.2 Findings

It's increasingly common that the majority of hosts being protected in a business are Windows or Linux.

Over the course of the NetWorker surveys, the client (and/or storage node) operating types in use (collapsing all types of Windows to a single OS, and all types of Linux to a single OS) has been as follows:

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%
Dec 2014	31%	33%	15%	21%
Dec 2015	31%	35%	16%	18%
Dec 2016	37%	31%	14%	18%

© Preston de Guise

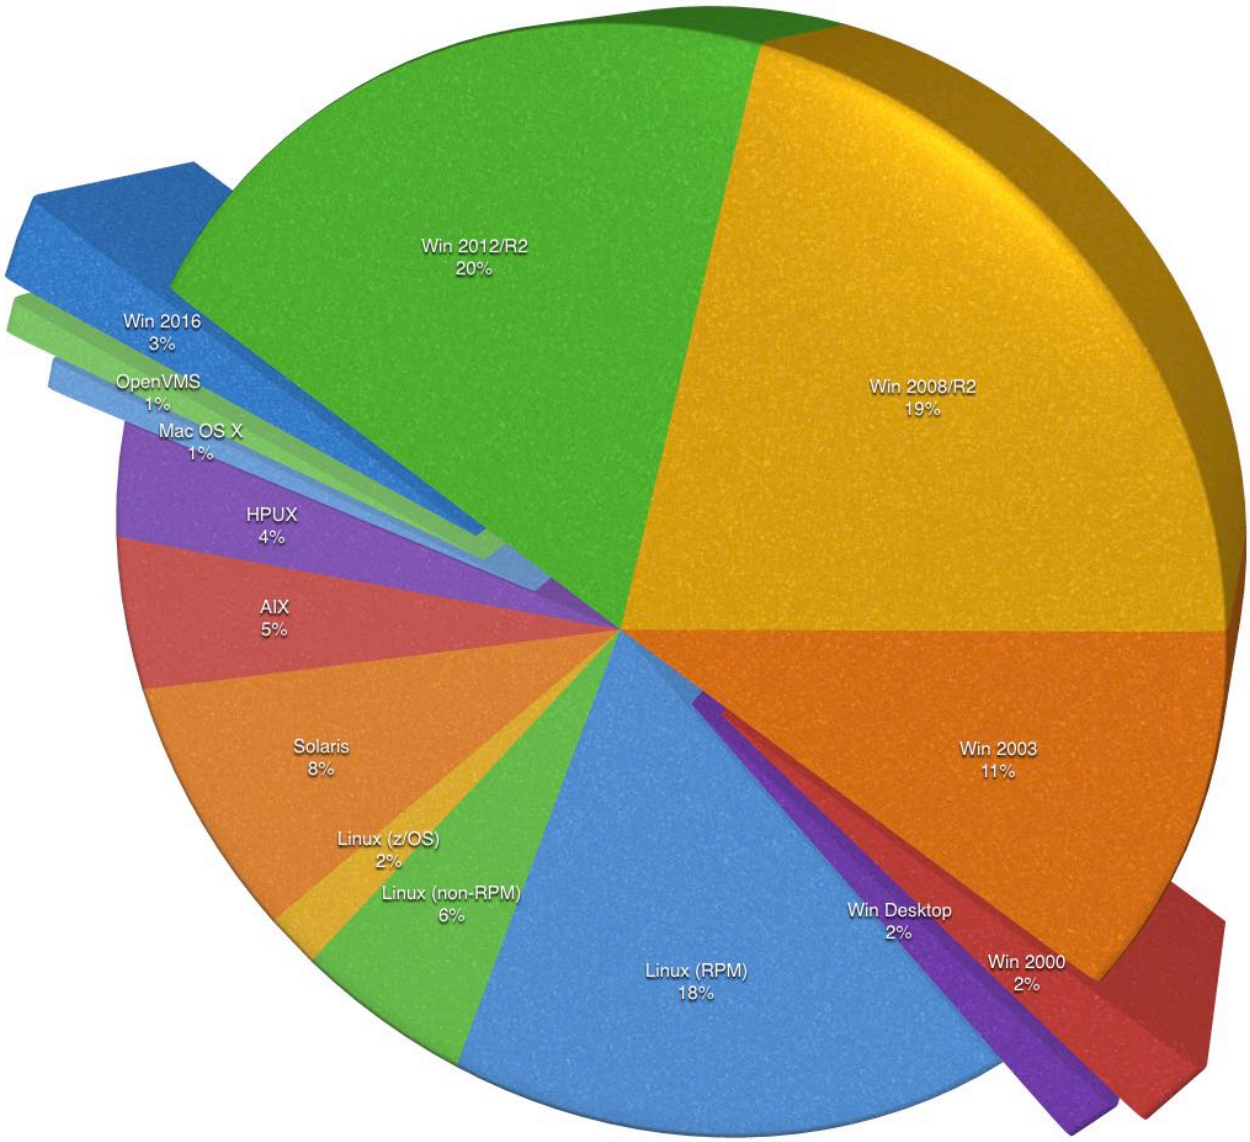

The ongoing trend for client and storage node operating systems is shown on the following page. There has been a marked increase in Windows client/storage node systems since the last survey, seemingly as a result of a decline in both Linux and Solaris clients. With Microsoft starting to focus further on Linux support (e.g., with their SQL Server system), that trend may start to settle back down over time.

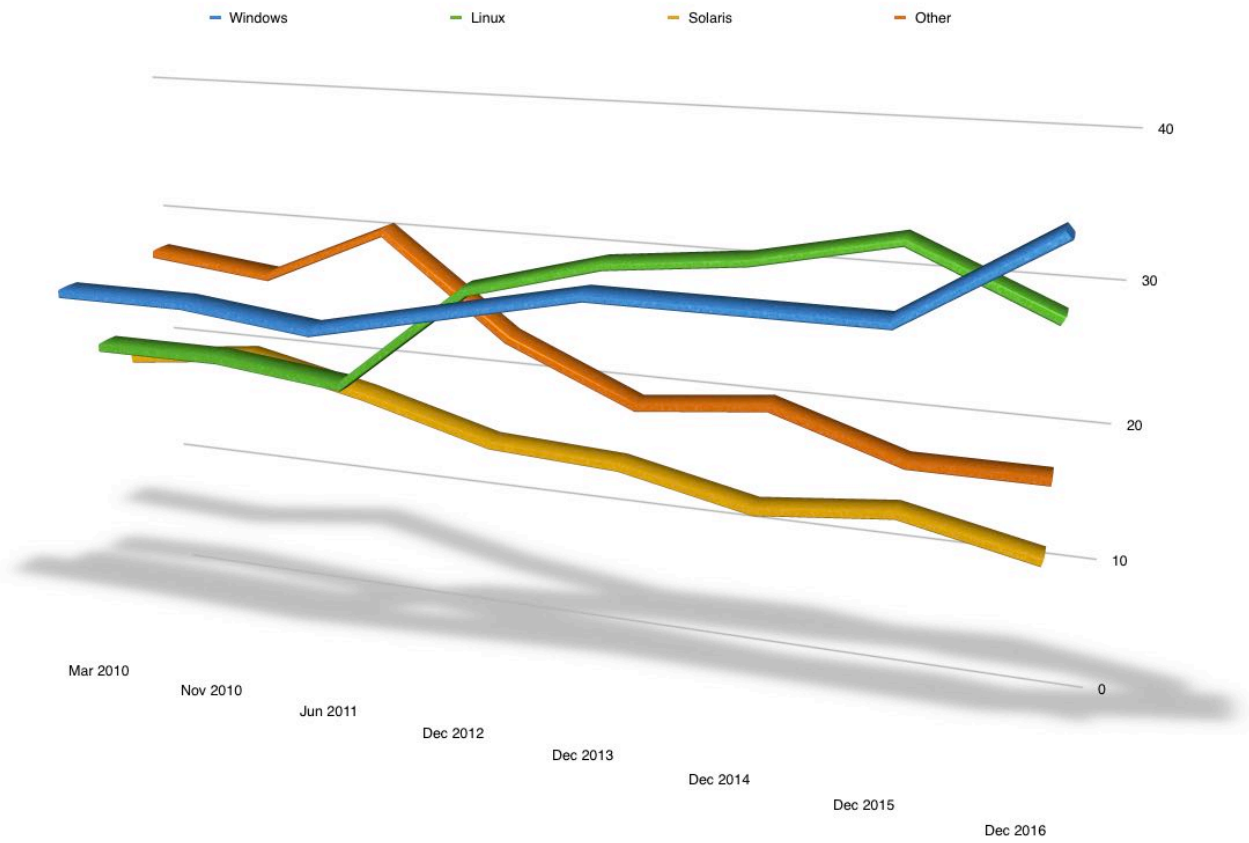
Windows
Linux
Solaris
Other
40
30
20
10
0
Mar 2010
Nov 2010
Jun 2011
Dec 2012
Dec 2013
Dec 2014
Dec 2015
Dec 2016

7 Businesses Using Deduplication

7.1 Responses

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

Dedupe?	*Number*
Yes Source Only	15
Yes Target Only	32
Yes Source & Target	90
No	21

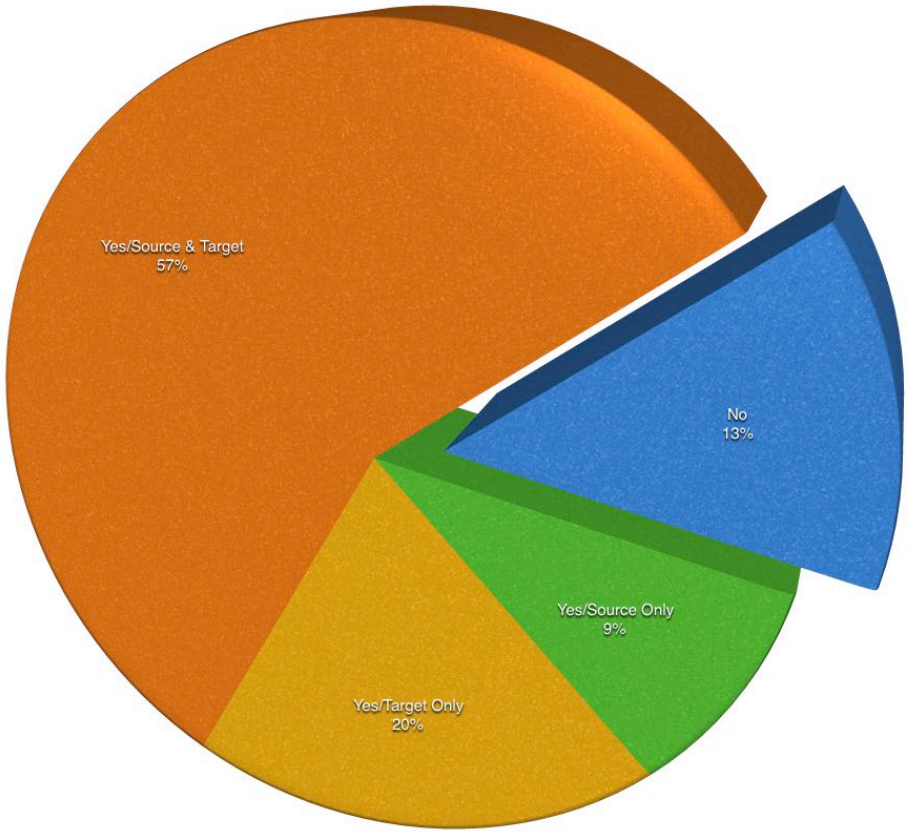

7.2 Findings

We are now seeing a marked reduction in the number of businesses not using any form of deduplication. This number was 22% in 2014, 24% in 2015, and now just 13% in 2016.

The year on year findings have been as follows:

Survey	*None*	*Source Only*	*Target Only*	*Source & Target*
Nov 2010[1]	68%	4%	20%	7%
Jun 2011	64%	5%	27%	4%
Dec 2012	37%	5%	31%	27%
Dec 2013	27%	8%	31%	34%
Dec 2014	22%	5%	28%	45%
Dec 2015	24%	3%	18%	56%
Dec 2016	13%	9%	20%	57%

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.

[1] Question was not asked in the March 2010 survey.

© Preston de Guise

Deduplication Type	Number
Non-EMC	73
Avamar	31
DD – VTL	23
DD – CIFS/NFS	35
DD – Boost	131

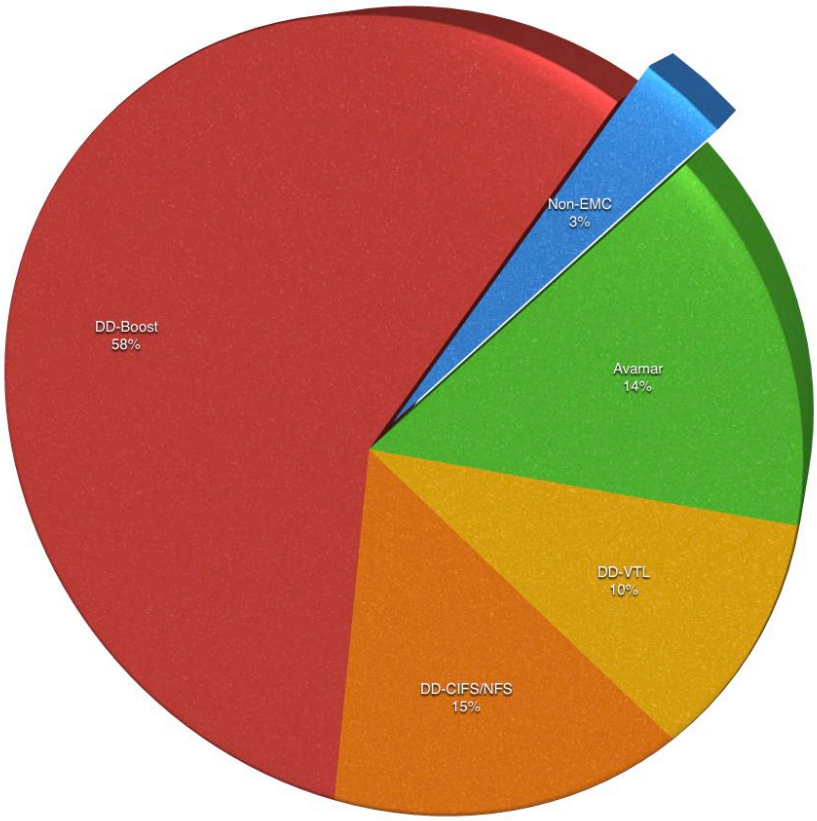

Non-EMC deduplication continues to slump (2014 – 8%, 2015 – 7%, 2016 – 3%). This is undoubtedly due to the exceptional levels of integration offered by Data Domain within a NetWorker (and data protection) environment. The use of Data Domain as the deduplication technology remains steady at 81% year on year.

(Avamar has shown a slight rise: this is likely due to the general awareness that the VBA technology included in the NetWorker 8.x tree has been borrowed from Avamar's virtual backup appliance systems.)

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

Region	Source and Target	Target Only	Source Only	No Deduplication
Americas	8	6	5	4
EMEA	57	20	8	12
APJ	7	1	0	1
Two Regions	7	2	1	1
Global	11	3	1	1

© Preston de Guise

8 NetWorker and Data Domain Modules/Plugins

8.1 Responses

Continuing from last year, this survey asked users for both NetWorker and Data Domain plugin usage. This year included NAS and other new snapshot related functionality into the list of plugins.

Module/Plugin	*Number*	*Module/Plugin*	*Number*
NDMP	64	NMM	122
NMDA	83	Boost Plugin for SQL	29
MSSQL Server (Legacy)	36	Exchange (Legacy)	25
Informix (Legacy)	3	DB2 (Legacy)	9
NAS Snapshot	8	ProtectPoint	2
Boost Plugin for RMAN	23	Boost FS	15
Meditech	3	Documentum	0
Oracle (Legacy)	30	SAP on Oracle	24
Sybase (Legacy)	2	PowerSnap	4
SnapImage	0	EDM	1
MySQL Enterprise	6	Boost for Ent. Apps	24

© Preston de Guise

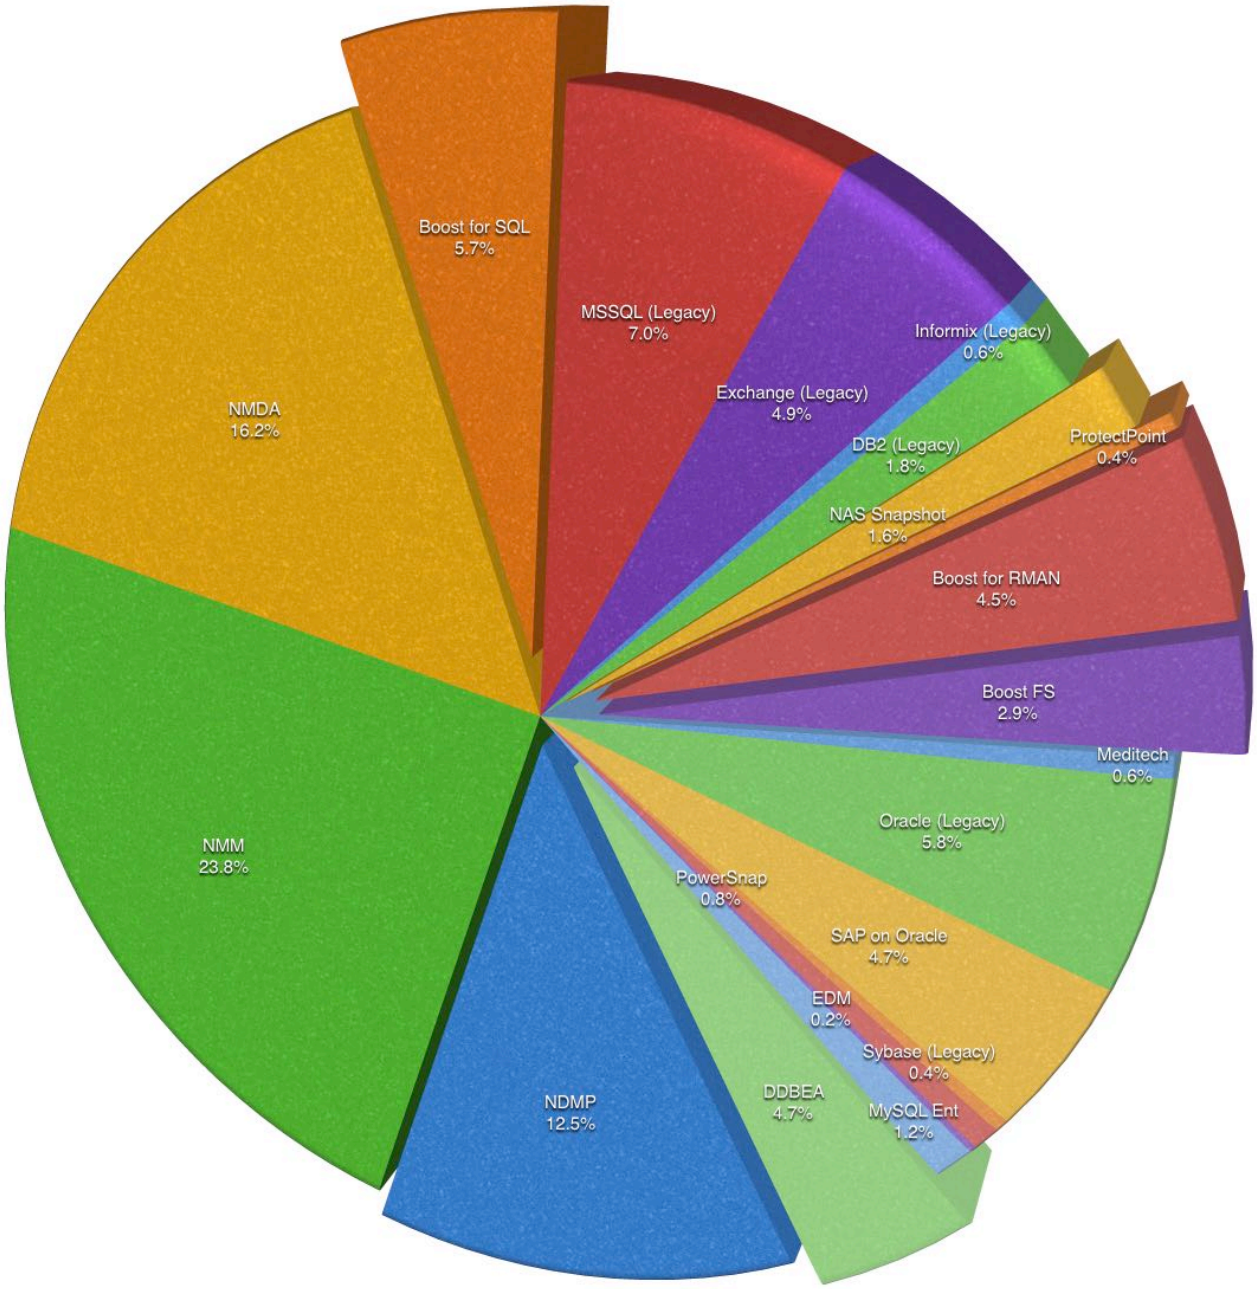

8.2 Findings

Documentum and SnapImage both recorded zero active users in the respondents. Given both of these plugins are end of life future versions of the survey may drop them.

Low ProtectPoint usage in NetWorker does not correspond to ProtectPoint usage generally, since the plugin can operate independently of NetWorker, and indeed is often deployed in environments with other backup problems to solve the performance of large database protection actions.

18.2% of respondents were using plugins primarily designed to operate *without* a backup product (Boost for SQL, Boost for RMAN, etc.) This reflects a changing attitude in many businesses, that being the importance of keeping backup and recovery control with subject

matter experts (SMEs) rather than requiring the NetWorker backup administrator(s) to become multi-disciplinary SMEs.

At the same time, given the costs associated with primary storage (even factoring in storage efficiencies found in all-flash primary storage), maintaining a large number of database 'dumps' on primary storage is an impractical solution unsuited for modern environments. The best of breed mixing of Data Domain as protection storage ("storage of last resort") coupled with plugins that allow database and application administrators direct access through native tools circumvents the "dump and sweep" approach found in many organisations where DBAs maintain full control of their backup processes. This evolving trend is discussed in more detail in "Data Protection: Ensuring Data Availability". Even in environments with full enterprise backup software such as NetWorker, we will likely see continued adoption of "database direct" plugins to spread operational activities for data protection amongst the critical SMEs of the business.

9 Virtualisation within the Environment

9.1 Responses

Virtualization is consuming 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.

Host Type	*Number*
No virtualization	4
Some clients	120
All clients	21
Some storage nodes	30
All storage nodes	19
Some NSR servers	31
All NSR servers	51

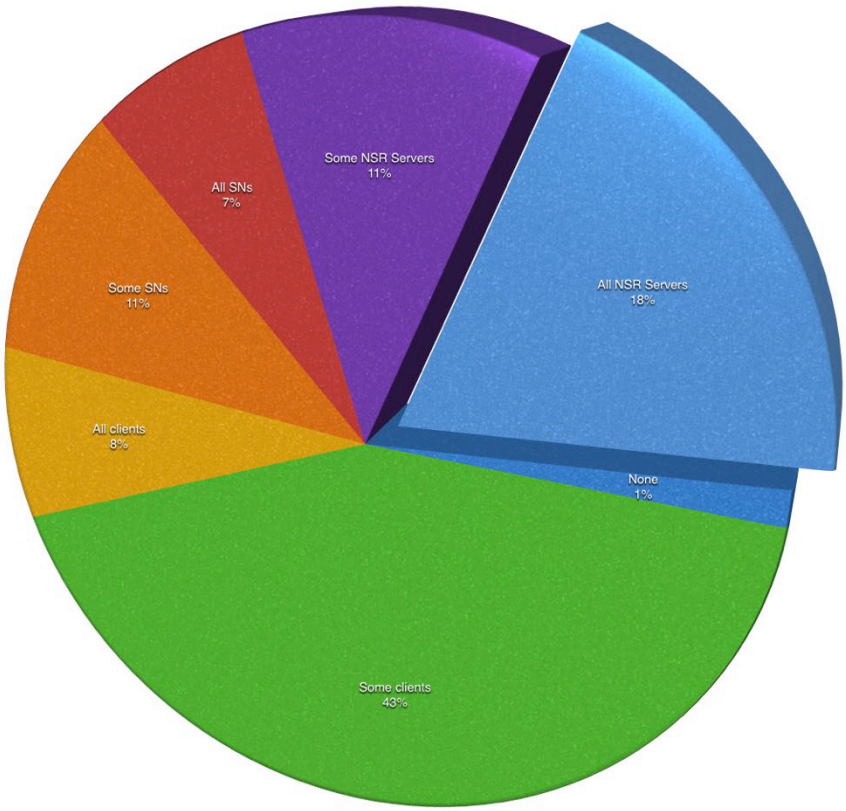

9.2 Findings

We've seen decreasing numbers of respondents in the surveys who say they have no virtualisation within their environment. This has shrunk from 6% to 5% to now, 1% (2014, 2015 and 2016 surveys respectively). Given the pervasiveness of virtualisation in business and the growing popularity of HyperV in addition to VMware, there is now a very high likelihood that backup administrators will need to work with virtualisation as part of their protection processes.

We are continuing to see an increase in the percentage of servers and storage nodes that are being virtualised:

Year	*Some Storage Nodes*	*All Storage Nodes*	*Some Servers*	*All Servers*
2014	10%	5%	10%	9%
2015	14%	3%	10%	12%
2016	11%	7%	11%	18%

The small drop in 'some storage nodes' between 2015 and 2016 is balanced by the increase in 'all storage nodes' being virtualised in 2016. Increasingly the only reason to deploy a physical storage node in a NetWorker environment is when tape-out functionality is required. Particularly when NetWorker is coupled with Data Domain for storage, the actual systems infrastructure (server and storage nodes) required to run a NetWorker environment are trivial compared to

some of the other competitors on the marketplace. (This is often quite surprising, anecdotally, with prospective customers, to hear that they can protect their entire infrastructure using a relatively few light-weight virtual machines, compared to spending sometimes millions of dollars on high performance physical hosts just to run media servers and media agents for other backup products.)

There is definitely a move for NetWorker servers to be increasingly virtualised. This gives a variety of additional recovery options, particularly when larger organisations use stretched networks across datacentres, allowing a server to be migrated seamlessly between two physical locations while still running.

Looking to the virtual machines themselves, the protection methods are evolving:

Protection Method	*Number*
In guest	98
VCB	0
VADP	36
VBA	73
Custom Scripts	9
3rd Party Products	10

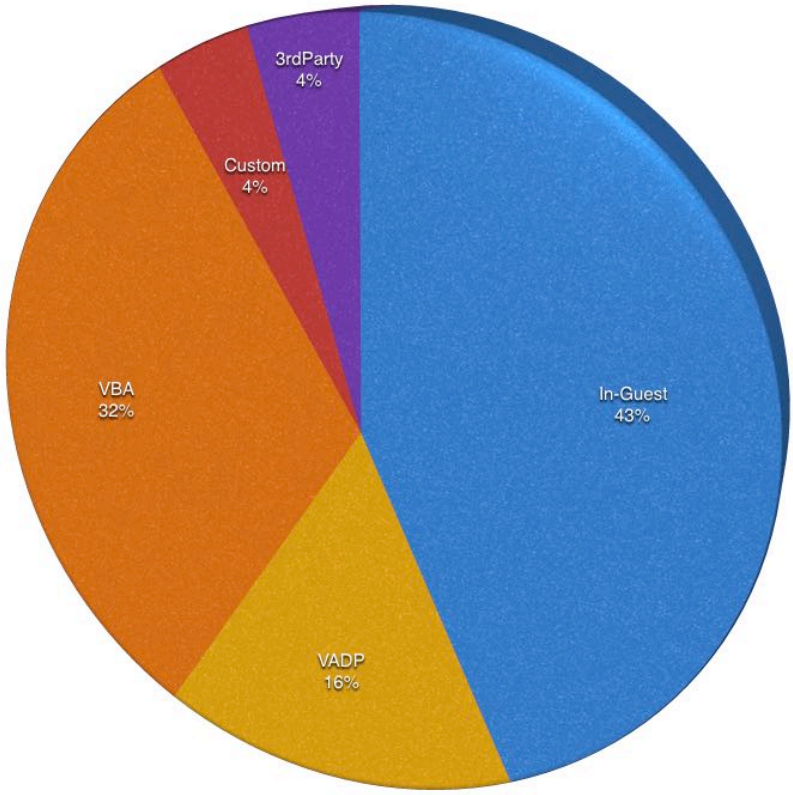

It should be noted that multiple options could be chosen – it is quite common for instance for a business to perform image based backups using say, VBA, and then in-guest backups for databases or applications running on the virtual systems to ensure full application recoverability.

The year-on-year comparison between key backup functions for virtual machines is as follows:

Year	*In-Guest*	*VCB*	*VADP*	*VBA*	*Custom*	*3rd Party*
2014	42%	3%	24%	20%	4%	8%
2015	45%	2%	16%	22%	6%	10%
2016	43%	0%	16%	32%	6%	4%

The next survey will also poll NVP (NetWorker Virtual Proxy) usage, the next generation, high-performance replacement to VBA.

© Preston de Guise

10 Backup to Disk Technology

10.1 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*
No	15
AFTD (dedupe)	21
AFTD (non-dedupe)	28
VTL (dedupe)	22
VTL (non-dedupe)	6
Yes – Boost	121

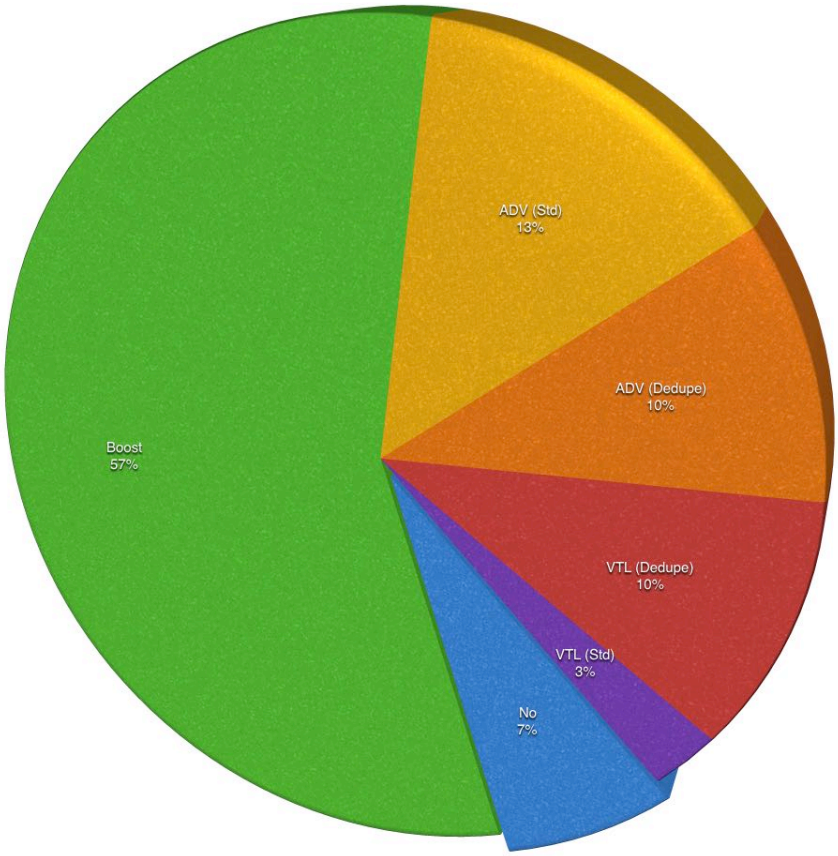

10.2 Findings

We are seeing a continuing trend towards the majority of backup to disk accomplished in NetWorker via Data Domain. The pairing of the two technologies is significantly complimentary and achieves great results for customers. This will only continue given NetWorker's deep integration into Data Domain new features, such as Cloud Tiering.

Survey	*None*	*ADV*	*VTL*	*Boost*
Nov 2010[2]	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%
Dec 2016	7%	23%	13%	57%

The following graph shows the overall trends over the course of the past 7 surveys:

[2] Question was not asked in the May 2010 survey.

© Preston de Guise

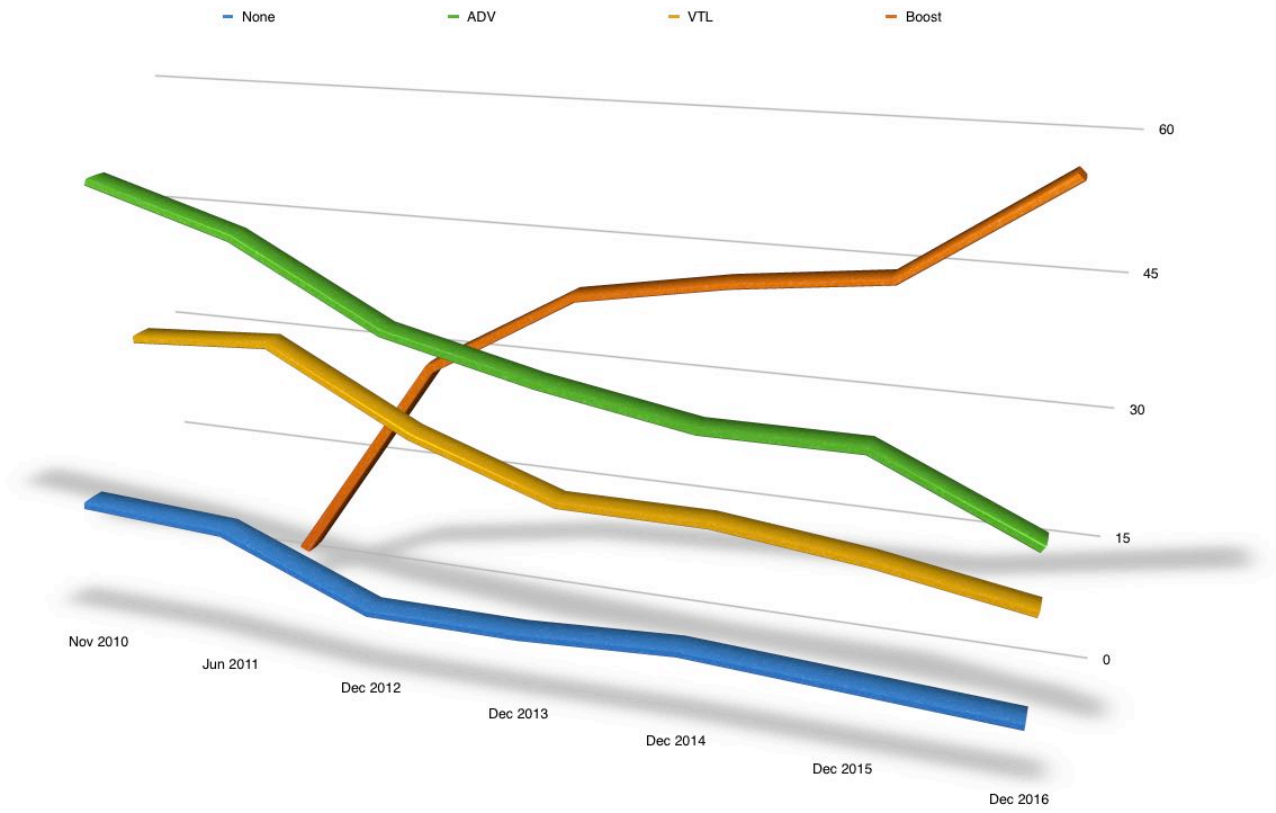

In little more than 5 years, Data Domain Boost has become the primary option for backup to disk technology with NetWorker. The storage efficiencies of Data Domain, combined with distributed segment processing at the client side to substantially reduce the amount data needing to be streamed across for the backup has definitely resulted in rapid adoption within the marketplace.

If we evaluate the estimated FETB sizes for a backup environment vs whether or not a form of backup to disk will be used, there is a higher likelihood for environments with 101 or more TB of data, but there is still a high propensity for backup to disk being used even in those environments with smaller FETB sizes:

FETB Size	*%Using Backup to Disk*	*%Not using Backup to Disk*
No idea	83%	17%
< 1 TB	100%	0%
1-50 TB	88%	12%
51 – 100 TB	87%	13%
Up to 100 TB	**87.9%**	**12.1%**
101 – 500 TB	97%	3%
501 TB – 1 PB	88%	12%
1+ PB to 5 PB	100%	0%
5+ to 10 PB	100%	0%

FETB Size	*%Using Backup to Disk*	*%Not using Backup to Disk*
10+ PB	100%	0%
101+ TB	**95.7%**	**4.3%**

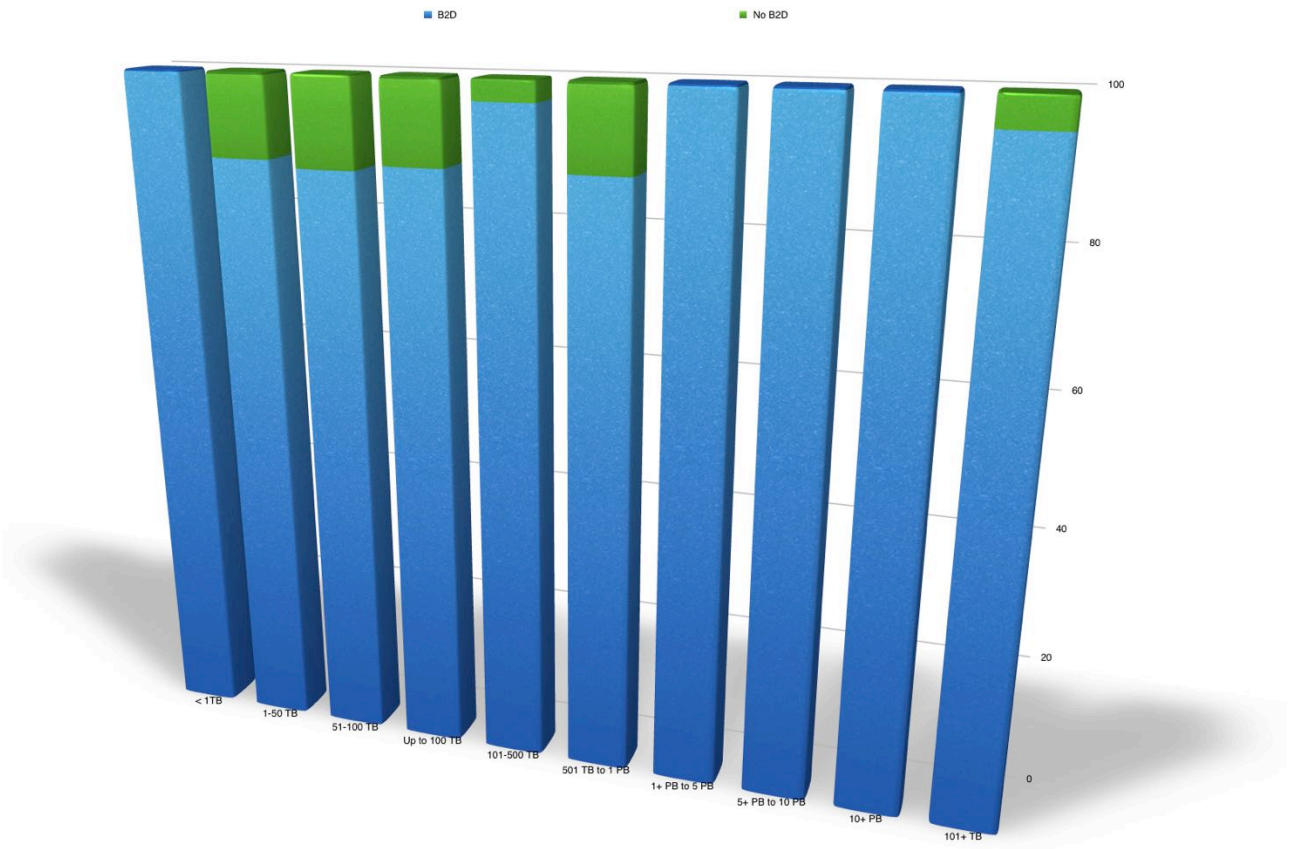

In short, it is now practically unheard of based on ongoing surveys to find environments where there is *no* backup to disk technology involved.

11 Do you clone within your environment?

11.1 Responses

This question focuses on understanding whether businesses are duplicating their backups to ensure they do not represent a single point of failure within the data protection environment.

Answer	*Responses*
No No Budget	8
No Not Interested	7
No No Time	5
Yes All	62
Yes Prod Only	29
Yes Selectively	48

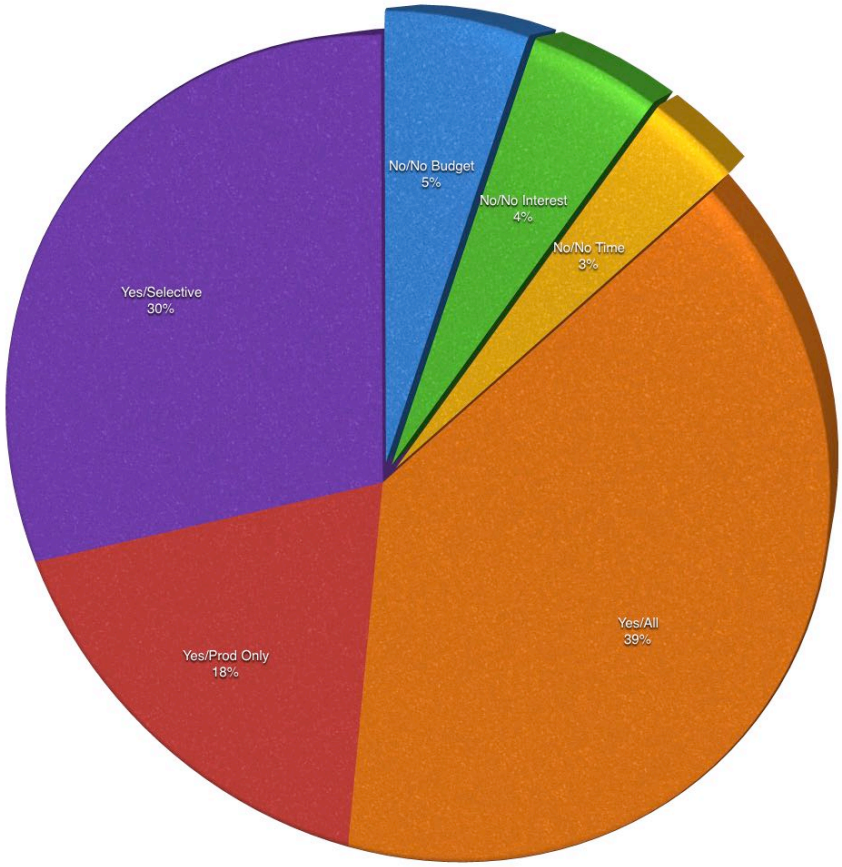

11.2 Findings

The surveys continue to show high percentages of businesses choosing to clone at least some of their backups:

Survey	*Some cloning*	*No cloning*
Dec 2012	76%	24%
Dec 2013	78%	22%
Dec 2014	80%	20%
Dec 2015	83%	17%
Dec 2016	87.4%	12.6%

The continuing trend remains an increase in the number of businesses choosing to clone at least *some* of their backups. Time saving technologies such as Data Domain deduplicated replication, and increased awareness of the hazards of backups being a single point of failure are likely both contributing to the increasing protection maturity. We are also seeing increasing

percentages of businesses cloning all their backups: 23% in 2013, 28% in 2014, 34% in 2015 and 39% in 2016.

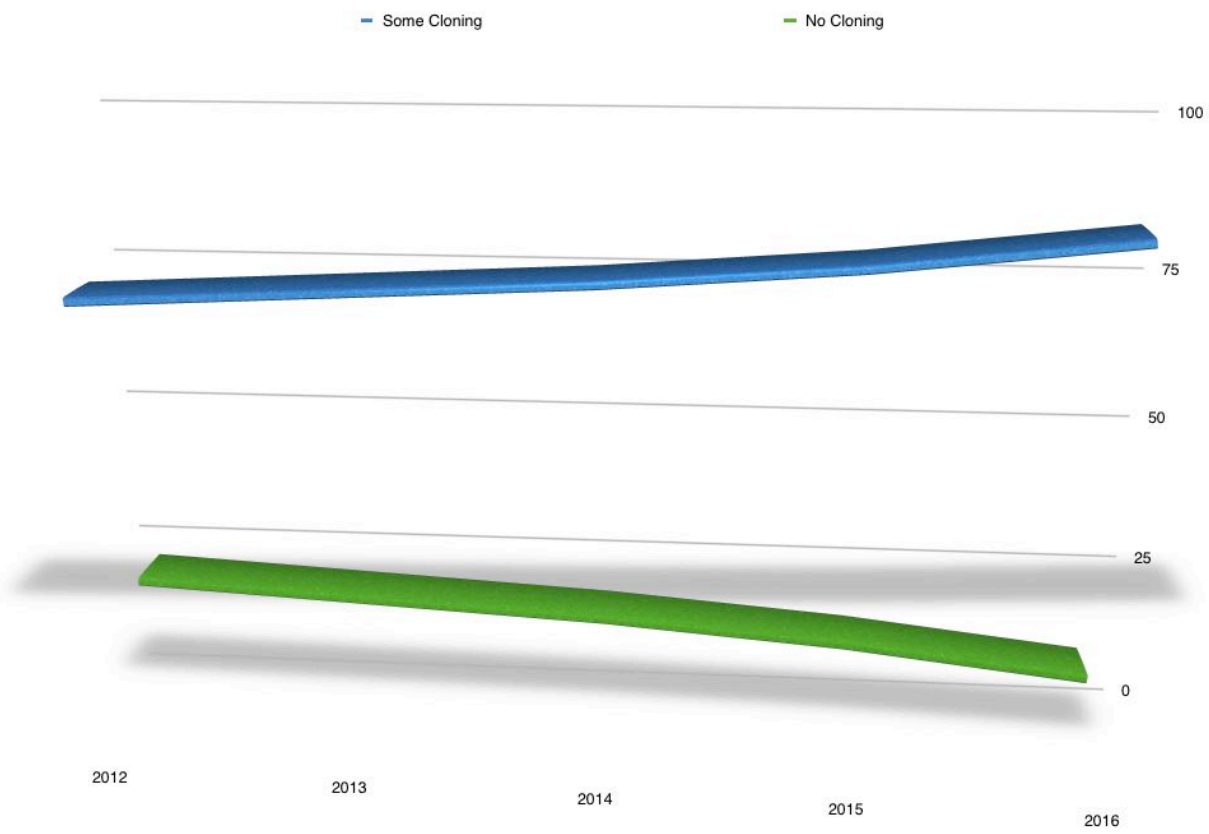

If we compare the use of cloning within environments to the FETB size of the environment, we do see increased likelihood that cloning will be used for at least some backups as the environment size increases:

FETB Size	*Some cloning*	*No cloning*
No Idea	*100%*	*0%*
< 1 TB	50%	50%
1-50 TB	80%	20%
51-100 TB	83%	17%
Up to 100 TB	**80%**	**20%**
101-500 TB	84%	16%
501 TB – 1 PB	94%	6%
1+ PB to 5 PB	93%	7%
5+ PB to 10 PB	100%	0%
10+ PB	100%	0%
101+ TB	**90%**	**10%**

Stacked, these percentages show as follows:

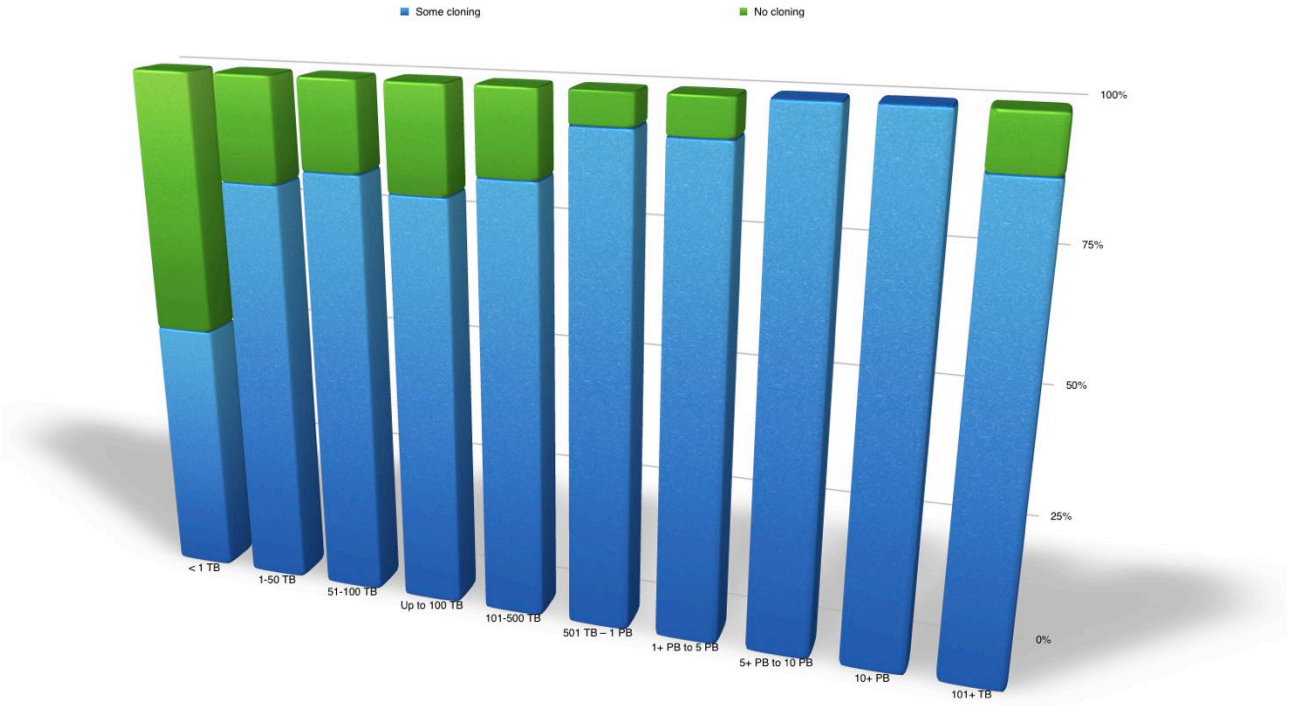

This would suggest that (by and large) as the amount of data to be protected by an environment increases, the more likelihood there is that the business will take backup duplication seriously.

12 Longest Retention Time

12.1 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.

Retention Time	*Number*
<= 31 days	9
1+ to 3 months	17
3+ to 6 months	11
6+ to 12 months	23
1+ to 3 years	26
3+ to 7 years	23
7+ to 15 years	39
15+ years	11

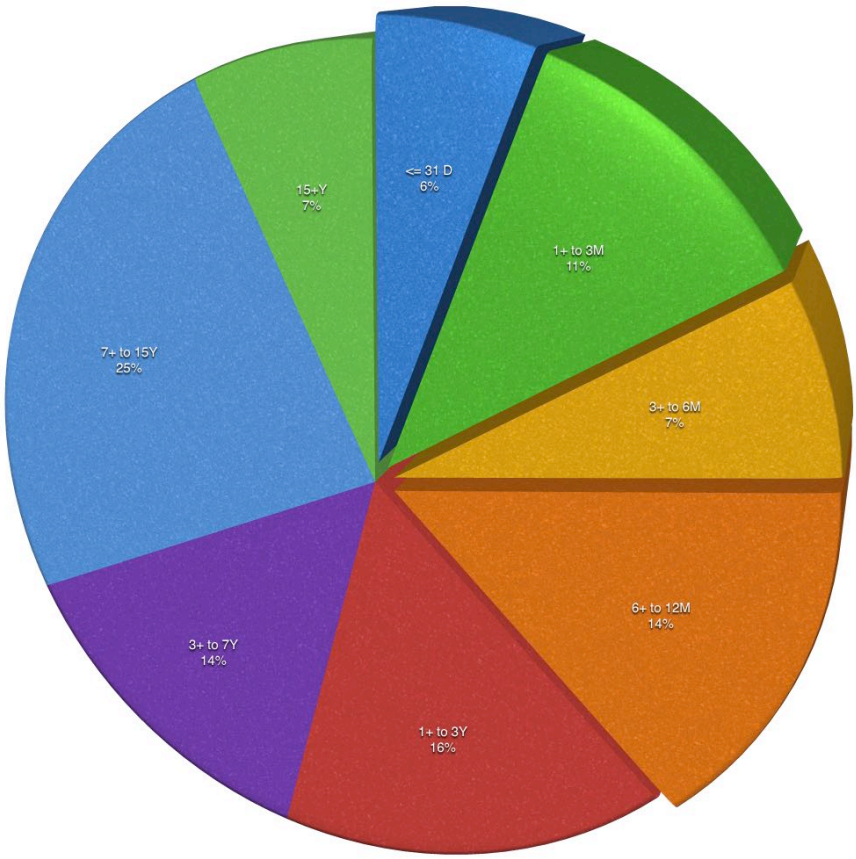

12.2 Findings

In similar results to last year (37%), this year saw 38% of respondents having a longest retention time of 1 year or less.

27% of respondents are using NetWorker to retain backups for 7+ years. 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.

It is interesting to see the longest retention times by regional use of NetWorker[3]:

Region	*<= 31D*	*1+ to 3M*	*3+ to 6M*	*6+ to 12M*	*1+ to 3Y*	*3+ to 7Y*	*7+ to 15Y*	*15+ Y*
Americas	0%	4%	8%	8%	44%	24%	12%	0%
EMEA	7.2%	13.4%	7.2%	19.6%	12.4%	6.2%	29.9%	4.1%
APJ	0%	10%	10%	0%	0%	40%	20%	20%

[3] Slight rounding errors noted and maintained to avoid longer numbers.

© Preston de Guise

Region	<= 31D	1+ to 3M	3+ to 6M	6+ to 12M	1+ to 3Y	3+ to 7Y	7+ to 15Y	15+ Y
Two Regions	9.1%	0%	9.1%	18.2%	0%	36.4%	18.2%	9.1%
Global	6.3%	12.5%	0%	0%	18.8%	18.8%	18.8%	25%

13 Backup Encryption

13.1 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.

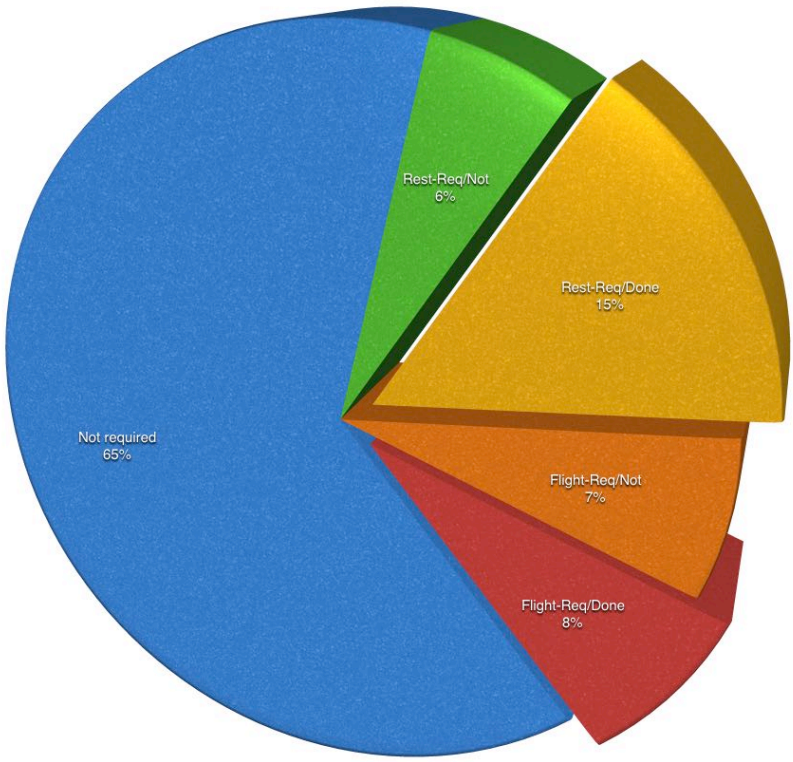

Answer	*Number*
Not required	111
At-rest required, not implemented	10
At-rest required, implemented	25
In-flight required, not implemented	12
In-flight required, implemented	14

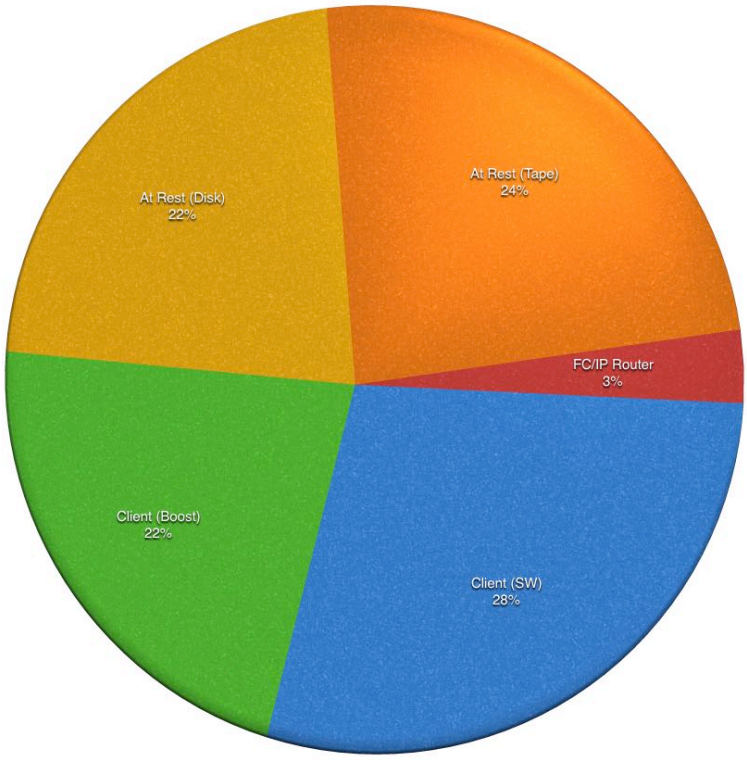

Answer	*Number*
Client side (software)	19
Client (in-flight/Boost)	15
At rest (disk)	15
At rest (tape)	16
FC/IP Encryption Routers	2

13.2 Findings

We have seen a slight increase this year in the number of respondents that require some form of encryption within their backup environment, regardless of whether they've actually implemented it yet. (65% state "not required" in the 2016 survey, down from 68% in the 2015 survey.)

In-flight encryption has grown from last year, with 15% of respondents indicating it's required (regardless of whether it is implemented), vs just 10% last year. This could indicate a growing

understanding in business that the internal network is not guaranteed to be secure. Increasing numbers of countries now also require the encryption of any potentially sensitive data not only when stored at rest, but when it is in transit, as well.

The regional breakdown on encryption requirements was as follows:

Region	*Not Req*	*At-Rest Req, Not Impl*	*At-Rest Req, Impl*	*In-Flight Req, Not Impl*	*In-Flight Req, Impl*
Americas	14	3	6	3	2
EMEA	77	2	10	6	8
APJ	6	1	2	0	1
Two Regions	4	3	3	2	2
Global	10	1	4	1	1

© Preston de Guise

14 Longevity of NetWorker Use

14.1 Responses

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

Time	*Number*
<1 year	9
1-2 years	10
2+ to 5 years	39
5+ to 10 years	47
10+ to 15 years	30
15+ years	24

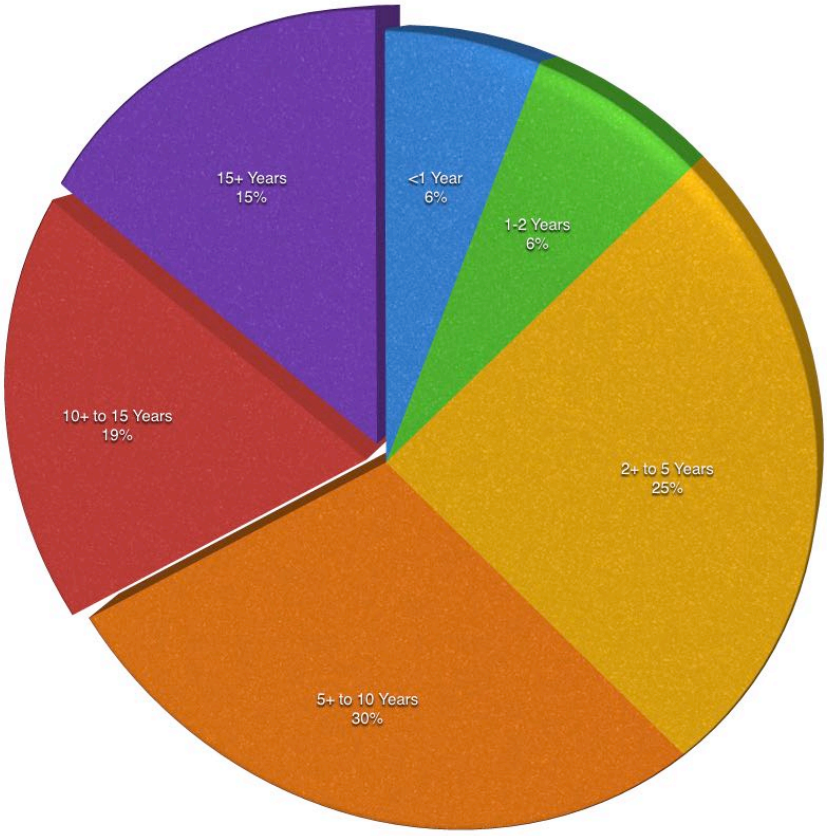

14.2 Findings

We continue to see a large number of respondents who have been using NetWorker within their environments for a considerable period of time. 15% of respondents have had it installed within their environment for 15 or more years, and all up, 34% of respondents have had the product installed in and protecting their environment for more than a decade.

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

By region, the NetWorker longevity responses were as follows:

Region	*<1 Y*	*1-2 Y*	*2+ to 5Y*	*5+ to 10Y*	*10+ to 15Y*	*15+ Y*
Americas	0	4%	36%	24%	24%	12%
EMEA	5%	8%	18%	30%	20%	19%
APJ	10%	0%	70%	10%	0%	10%
Two Regions	9%	0%	37%	45%	0%	9%
Global	13%	6%	12%	38%	25%	6%

15 Dedicated Backup Administrators?

15.1 Responses

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

Response	*Number*
Yes	84
No	75

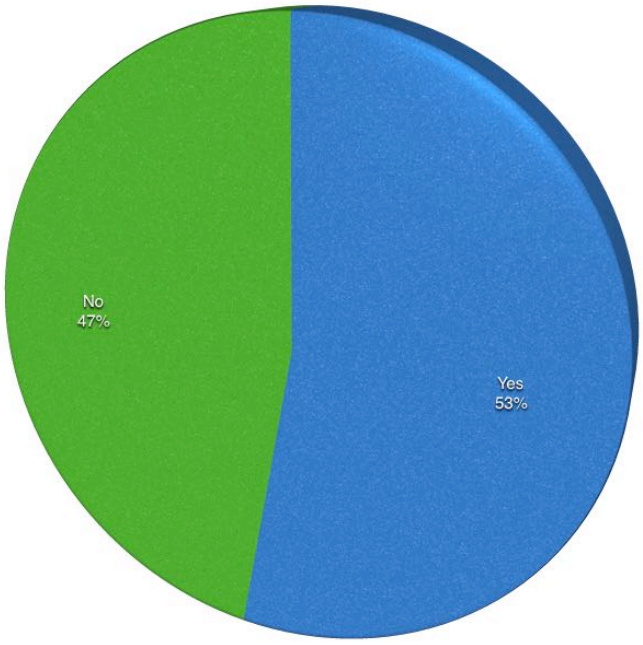

15.2 Findings

We have seen a decline in the number of environments featuring dedicated backup administrators. In the previous survey (2015), 66% of respondents indicated they had dedicated backup administrators, whereas in this survey just 53% of respondents said they did.

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:

- 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.

There is, however, an increased likelihood of there being dedicated backup administrators as the number of datazones within an environment increases:

#Datazones	*Dedicated Backup Admins (%)*	*No Dedicated Backup Admins (%)*
1	26%	74%
2	54%	46%
3	75%	25%
4	67%	33%
5	75%	25%
1-5 Datazones	**41%**	**49%**
6-10	85%	15%
11-25	88%	13%
26-50	100%	0%
51-100	100%	0%
101+	100%	0%
6+ Datazones	**89%**	**11%**

It might be expected that as businesses move more towards running their IT environments as hybrid or private clouds, there is the potential for the chances of there being dedicated backup administrators to shrink, even in larger environments. As backup functionality gets integrated into enterprise service catalogues via REST APIs, etc., backup administrators may very well evolve into backup *architects*, or more correctly, *data protection architects* within the business. For the time being at least, there is a considerably higher chance of there being dedicated backup administrators within an environment when there are 6 or more NetWorker datazones.

© Preston de Guise

16 Has your business suffered a ransomware attack?

16.1 Responses

This was a new question introduced into the 2016 survey: "Has your business suffered a ransomware or other data destructive attack in the past 12 months?" The purpose of this question was to gauge businesses that are being struck by this new generation of malware.

Response	*Number*
Yes	49
No	77
Don't know	18
Prefer not to say	15

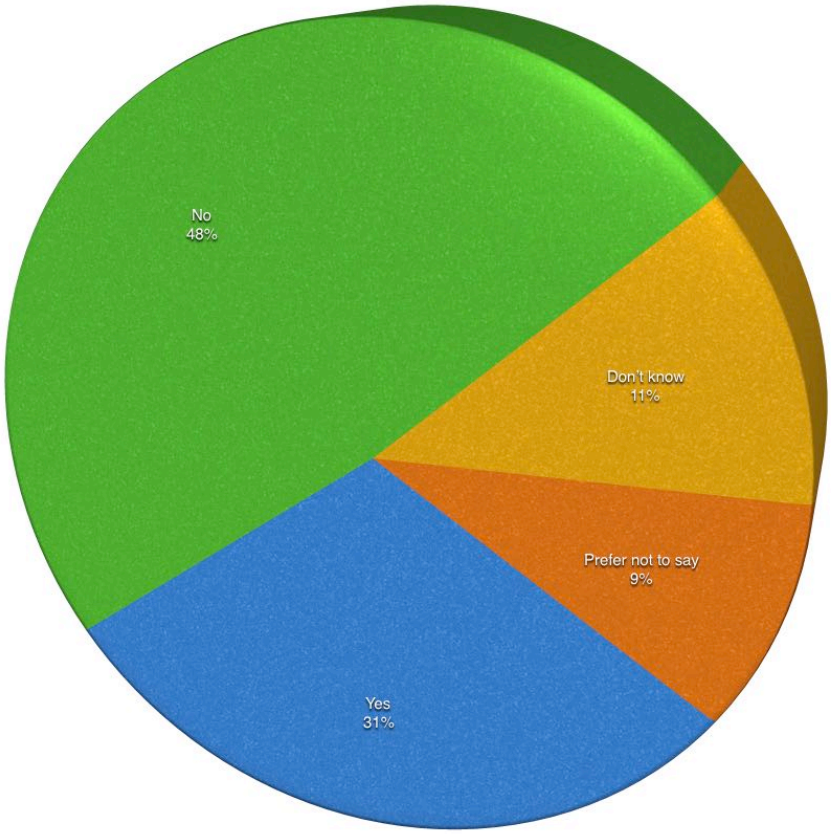

16.2 Findings

"Prefer not to say" was used to give respondents the option of not directly answering the question. Based on personal as well as anecdotal experience, there is a reasonably high percentage of staff at businesses who will, when a ransomware discussion starts, say they have not been affected or would prefer not to say, then throughout the conversation admit to penetration. It is understandable for businesses to be reluctant to speak about data destructive events that have taken place within their environments, but what these results should show is that these new types of malware attacks are reasonably common already.

If we assume that at least half of the respondents who answered "Prefer not to say" had actually been impacted by ransomware style attacks, the number of "Yes" respondents grows to 35% or higher.

It is unsurprising given these numbers that the US FBI had estimated ransomware would be a $1 billion industry by the end of 2016.

Looking at ransomware vs the presence of dedicated backup administrators, as you would imagine this does not immunise a business:

Dedicated Backup Administrators?	*Hit*	*Not Hit*	*Don't Know*	*Prefer not to Say*
Yes	25%	51%	13%	11%
No	38%	45%	9%	8%

© Preston de Guise

The regional breakdowns on ransomware penetration were as follows:

Region	*Hit*	*Not Hit*	*Don't Know*	*Prefer not to Say*
Americas	28%	56%	12%	4%
EMEA	36%	39%	14%	11%
APJ	10%	70%	0%	20%
Two Regions	9%	82%	9%	0%
Global	32%	56%	6%	6%

Finally, the breakdown of FETB size of an environment vs ransomware penetration[4]:

FETB	*Hit*	*Not Hit*	*Don't Know*	*Prefer not to Say*
No idea	31%	39%	17%	13%
<1 TB	0%	100%	0%	0%
1-50 TB	25%	63%	5%	7%
51-100 TB	39%	43%	9%	9%
Up to 100 TB	**28.8%**	**57.6%**	**6.1%**	**7.6%**
101-500 TB	42%	38%	10%	10%
501 TB – 1 PB	29%	24%	29%	18%
1+ PB to 5 PB	29%	57%	14%	0%
5+ PB to 10 PB	0%	100%	0%	0%
10+ PB	20%	60%	0%	20%
101 TB or More	**33%**	**43%**	**14%**	**10%**

[4] Small rounding errors will exist to keep numbers to 1 decimal place.

© Preston de Guise

17 Cloud

This section is broken into three key questions:

1. Do you have public cloud workloads?
2. If you have public cloud workloads, how are you protecting them?
3. Are you using, or considering using cloud/object storage for backups?

17.1 Responses

Public Cloud Workloads?	*Responses*
No Public	94
Don't Know	20
IaaS	17
SaaS	36
PaaS	3

Of those using public cloud workloads:

Cloud Protection Mechanism	*Responses*
Don't know	24
Rely on Cloud Provider	14
Data Dumps	2
Backup Infrastructure in the Cloud	11
SaaS Backup	3
Backup Cloud to On-Premises DC	4
No Persistent Data in Cloud	7

Public cloud workload aside, object strategies were as follows:

Using/Intending to use Object?	*Responses*
Not under consideration	116
Cloud Boost	26
Cloud Tier	13
Third Party	11

17.2 Findings

Public cloud does not yet rule the business:

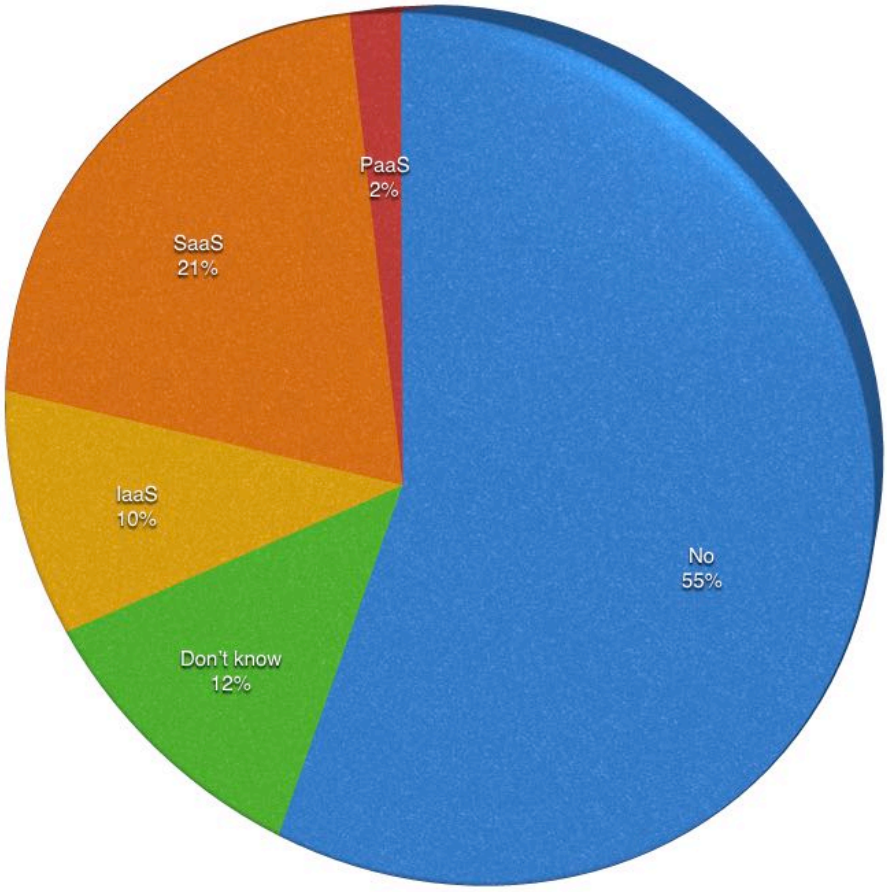

55% of respondents indicated that there was currently no public cloud workloads within the business. (Though we cannot rule out the presence of shadow IT – that is, after all, what it means). 12% of respondents admitted to not knowing whether there were any public cloud workloads, and this is not surprising – often public cloud workloads are initially trialled or spun up out of core business units rather than IT divisions. (Or at least in IT areas outside of datacentre infrastructure.)

SaaS is a more popular public cloud workload at the moment, and this too is understandable. Businesses wanting to test the water of public cloud see SaaS systems such as Office 365 or Salesforce as good "quick wins" for the business – they outsource the management of key business systems to providers who specialise in that function.

The breakdown of public cloud workloads vs FETB sizes within an environment is worth reviewing:

FETB Size	*No Public*	*Don't Know*	*IaaS*	*SaaS*	*PaaS*
No Idea	*50%*	*17%*	*4%*	*25%*	*4%*
< 1TB	0%	0%	50%	50%	0%
1-50 TB	58%	12%	10%	20%	0%
51-100 TB	63%	8%	12%	17%	0%
Up to 100 TB	**58%**	**10.5%**	**12%**	**19.5%**	**0%**
101 – 500 TB	60%	5%	9%	23%	3%
501 TB – 1 PB	68%	0%	11%	16%	5%

FETB Size	*No Public*	*Don't Know*	*IaaS*	*SaaS*	*PaaS*
1+ PB to 5 PB	44%	19%	12%	25%	0%
5+ to 10 PB	0%	25%	25%	50%	0%
10+ PB	40%	60%	0%	0%	0%
101 or more TB	**36%**	**28%**	**12%**	**24%**	**0%**

It is concerning to note a relatively high percentage of respondents who know there are public cloud workloads, but don't know how (or if) those workloads are protected (37% of respondents). This will be something to monitor in future surveys.

The idea of utilising high density, low cost object storage, regardless of whether it's on-premises (e.g., via ECS) or in a public cloud is still relatively new to businesses. As discussed in "Data Protection: Ensuring Data Availability", this topic will likely get increasing focus over the coming years as businesses realise the cost of maintaining tape is rarely, if ever, correctly calculated.

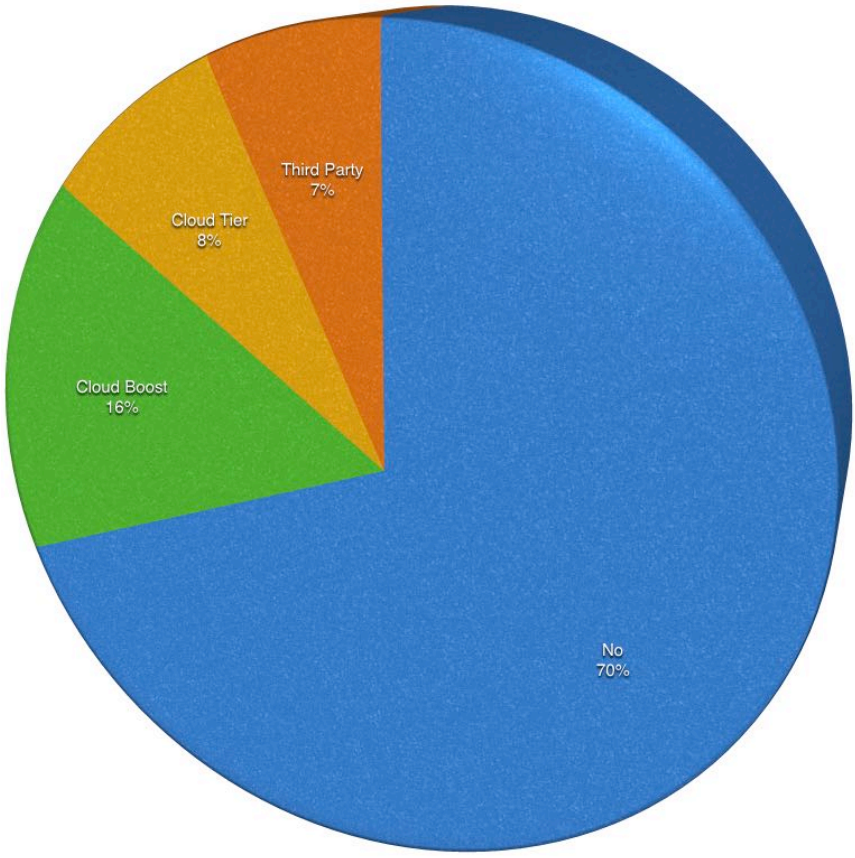

70% of businesses are not currently considering object storage as part of their backup strategy. However, object storage integrated into the backup environment is still a new topic, so 30% of respondents either using or considering the use of it indicates that businesses are eager to gain additional efficiencies in their data protection environment. This will be an area to watch over coming surveys.

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.

Use	*Number*
No	55
Daily Backups	33
Weekly Backups	18
Selective Backups	22
Monthly/Longer Backups	33
Daily Clones	20
Weekly Clones	21
Selective Clones	19
Monthly/Longer Clones	39

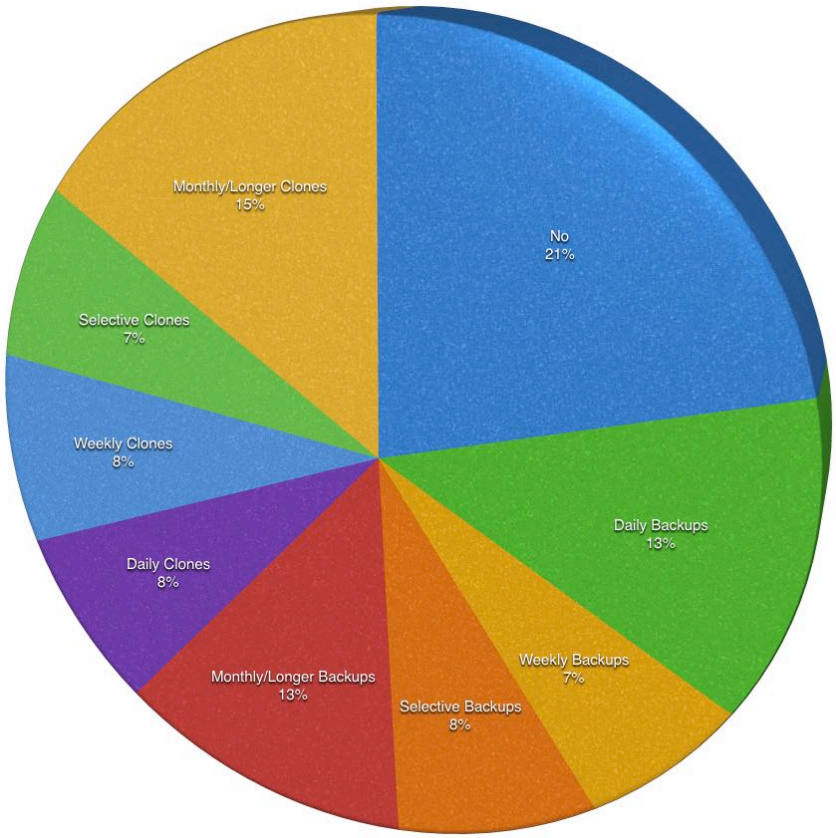

18.2 Findings

21% of respondents are not using tape at all within their environments, up from 20% in last year's survey, the first survey to poll tape utilisation.

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:

Region	*No Tape*	*Daily Backups*	*Weekly Backups*	*Selective Backups*	*Monthly Longer Backups*	*Daily Clones*	*Weekly Clones*	*Selective Clones*	*Monthly Longer Clones*
Americas	5	6	5	4	5	4	5	2	4
EMEA	39	19	8	9	16	10	13	12	20
APJ	3	2	1	3	1	3	1	2	3
Two Regions	2	3	0	4	5	1	0	3	5
Global	6	3	4	2	6	2	2	0	7

If we review the FETB size of an environment vs the likelihood of tape being used, we can see the following details:

FETB Size	*% Yes Backup*	*% Yes Clone*	*% No Tape*
No Idea	*32%*	*16%*	*52%*
< 1TB	0%	0%	100%
1-50 TB	31%	44%	25%
51-100 TB	39%	39%	22%
Up to 100 TB	**33%**	**41%**	**26%**
101 – 500 TB	24%	33%	43%
501 TB – 1 PB	32%	36%	32%
1+ PB to 5 PB	50%	45%	5%
5+ to 10 PB	50%	25%	25%
10+ PB	44.4%	44.4%	11.2%
101 or more TB	**35%**	**37%**	**28%**

© Preston de Guise

19 Conclusions

This year's survey expanded on previous years surveys by also evaluating cloud adoption and ransomware attacks.

From a trending perspective, we see:

- Increasing numbers of clients protected
- Higher adoption rates of more modern approaches to backing up virtual machines
- Decreasing use of Unix systems within backup environments
- Higher focus on Data Domain integration
- Increased cloning within environments
- Increased use of backup to disk within environments.

Many backup environments tend to have slower adoption rates for new technologies, but in some areas we've seen that is not the case – Data Domain Boost for instance rapidly went from a new product to something that was heavily used in NetWorker environments. Business use of IT is changing rapidly, and this means that backup and recovery environments no longer have the luxury of holding off implementing new technology until it is "old hat". This is why for instance we are already seeing 30% of businesses using or planning to use object storage to supplement their backup environments.

Ransomware is a serious issue within the IT environment. A high number of businesses have suffered ransomware or data destructive attacks, and this is only likely to increase. We would expect to see an increasing focus, particularly from financial and defence environments on the use of Isolated Recovery Sites (IRS). However, 'hacktivism' does not just strike these particular industries. Airline industries, medical businesses, media companies and even retailers all represent areas where "hacktivists" might take umbrage at a product, service or offering of the business and seek to permanently destroy data. The high percentage of ransomware attacks shows this next level of intrusion is something businesses need to be concerned about.

A full "push to public cloud" is not happening. Anecdotally, many customers and businesses who have previously adopted "cloud first" strategies are now rethinking those strategies to be "cloud fit" – i.e., making sure the workload they want to move to the cloud actually makes sense in the cloud.

At least some businesses looking at moving workloads into public cloud are seriously looking at how they can protect those workloads, though there is still some failings on that front. NetWorker customers can, at least, breathe easy as when they move traditional systems into cloud environments. The NetWorker/Cloud Boost architecture is highly optimised for cloud environments – a single CloudBoost appliance for instance deployed in the Cloud can address up to 6PB of object storage using just 32GB of RAM[5]. This will undoubtedly gain significant attention from businesses as public cloud workload adoption increases.

Thanks to everyone who participated in the survey.

[5] By comparison, other competitive options have stuck with the "lift and shift" model, leaving customers stuck with either not backing up in public cloud at all, or forced to deploy large numbers of resource-hungry backup servers to protect the infrastructure.

© Preston de Guise