NetWorker Usage Survey

Overview

Between June 22 and July 16, 2011, a survey was run on the NetWorker blog (**http://nsrd.info/blog**) to gather a high level overview of the NetWorker usage of as many respondents as possible.

This survey aimed to review:

- Operating systems;
- NetWorker versions;
- Licensed modules;
- Cloning;
- Deduplication take-up;
- Backup to disk;
- Open source databases;
- Encryption strategies.

About the Author

Preston de Guise has been specialising in data protection services since 1996, and has provided consulting services to a diverse selection of companies ranging from small sites to companies in the Global Fortune 500.

Preston is the author of "Enterprise Systems Backup and Recovery: A corporate insurance policy" (ISBN-10 1420076396, ISBN-13 978-1420076394). Written for both technical and management users, "Enterprise Systems Backup and Recovery: A corporate insurance policy" provides insight into best practice approaches to designing policies and procedures for ensuring that data protection solutions installed form a cohesive and reliable *system* within an enterprise. Details of the book can be found at **http://www.enterprisesystemsbackup.com**.

Preston de Guise currently works for IDATA Resolutions, an Australian/New Zealand company that specialises in storage, archiving, data protection, virtualisation and high availability solutions. IDATA provides a wide range of services including installation and configuration, training, remote support, remote audits, on-site support, operational assistance and managed services. IDATA Resolutions can be found on the net at **http://www.idataresolutions.com**.

Table of Contents

Survey Introduction

Continuing the NetWorker Hub's periodic polling of NetWorker Usage within the community (effectively half-annually), this survey ran from June 22 through to July 16, 2011.

In addition to asking similar questions to those asked in the surveys March and November 2010, this survey also evaluated the the take-up of backup encryption within the NetWorker community – a topic that will continue to be evaluated in future surveys to note trends.

There were 108 responses to the survey. Since the survey needed to cover multiple data zones, and most questions allowed multiple selections, there are several questions where the total number of selected answers exceeds the number of individual survey responses.

Much gratitude is owed to all respondents.

How many datazones are you running?

The numbers for the responses were as follows:

- **1** – 56
- **2** – 12
- **3** – 5
- **4** – 5
- **5** – 3
- **6-10** – 13
- **More than 10** – 12

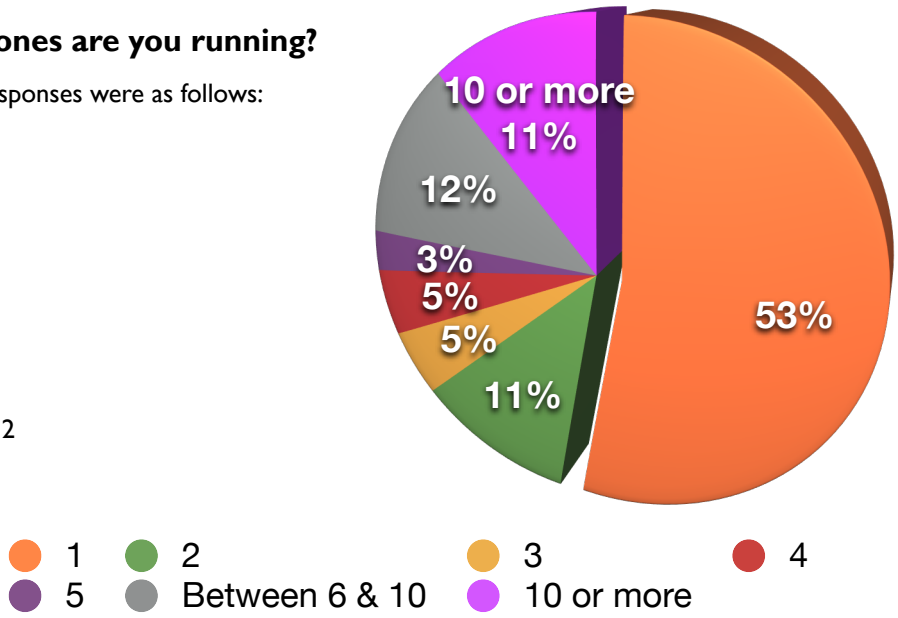

Comments and Conclusions

While the majority of NetWorker deployments feature a single datazone, there are significant numbers of environments with multiple datazones present.

With multiple datazones at so many sites, report aggregation systems such as EMC's Data Protection Advisor, and even NetWorker Management Console's reporting system become even more critical in providing a "big picture" view of the entire data protection environment.

What is the version of your NetWorker server(s)?

The numbers for the responses were as follows:

- **6.x or older** – 0
- **v7** – 0
- **v7.1.x** – 1
- **v7.2.x** – 11
- **v7.3.x** – 5
- **v7.4.x** – 20
- **v7.5.x** – 46
- **v7.6** – 9
- **v7.6 SP1** – 45
- **v7.6 SP2** – 23

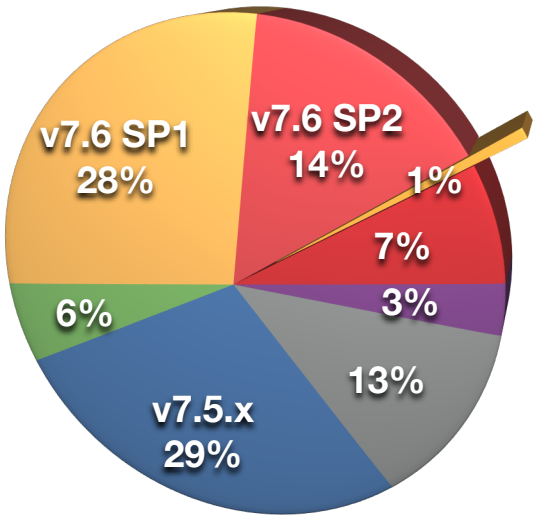

Comments and Conclusions

It's reassuring to see a continual drop-off across previous surveys of the numbers of sites still running older, unsupported versions of NetWorker.. In act, less than one quarter of respondents were running on an unsupported version of NetWorker, compared to the previous survey where close to a third of respondents were on unsupported versions.

EMC has been very successful in pushing customers to move to the latest release of NetWorker. In particular, the feature sets introduced into NetWorker version 7.6 SP1 and SP2 have encouraged a lot of customers to upgrade, despite a previous overall resistance amongst data protection customers to be on the latest version.

Indeed, over 40% of respondents are running current release servers. If we compare to the previous survey (November 2010), 39% of respondents were still running NetWorker 7.5.x, and 20% still running NetWorker 7.4.x. Both of these have significantly shrunk as customers have moved to take advantage of the features of the 7.6 service packs.

Number of clients being backed up

The numbers for the responses were as follows:

- **1-10** – 0
- **11-25** – 3
- **26-50**– 4
- **51-100** – 17
- **101-250**– 27
- **251-500** – 22
- **501-1000** – 9
- **1001 - 2000** – 14
- **2001 – 5000** – 10
- **5001+** – 2

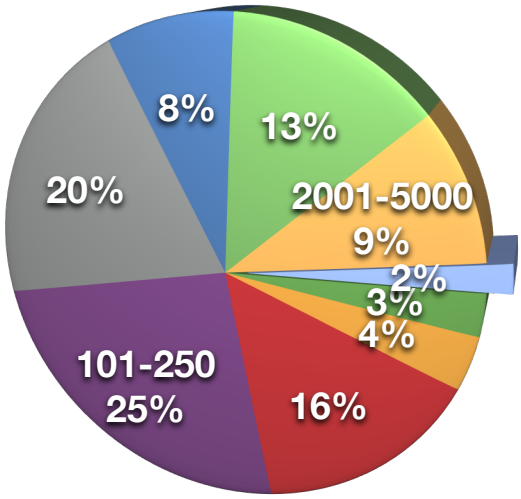

Comments and Conclusions

This survey question was expanded on from the November 2010 survey to further breakdown the number of sites running 2000+ clients. (The previous survey stopped at "2000+".)

11% of respondents are backing up more than 2000 clients in their NetWorker environments, continuing to show that NetWorker is a high end enterprise backup product. The range of hosts protected (from as low as 11-25 to 5000+) demonstrates NetWorker's design strength - an ability to scale up to extremely large environments, and down to the smallest of environments.

NetWorker Server Operating Systems

The numbers for the responses were as follows:

- **Windows** – 52
- **Linux** – 32
- **Solaris** – 46
- **AIX** –78
- **HPUX** – 1

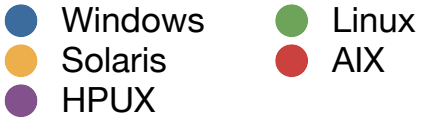

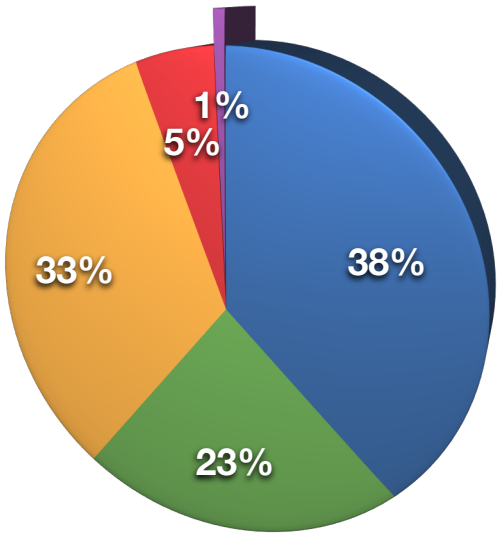

Comments and Conclusions

If we compare the major 3 server OS types over the three surveys completed:

Survey	Windows	Linux	Solaris
March 2010	29%	22%	43%
November 2010	29%	19%	43%
June 2011	38%	23%	33%

These numbers align with anecdotal evidence and the prediction made in the November 2010 survey results – companies are fleeing Sun under Oracle. Many businesses for instance have enacted policies that require Sun equipment to be swapped out for either Windows or Linux servers when it comes up for replacement, unless a rock solid business case can be made that the host functionality cannot be transferred to a different operating system. (For author opinion on why this will create significant long-term challenges for Oracle, refer to **http://nsrd.info/blog/2010/04/22/rip-solaris/**)

We will continue to track Solaris NetWorker usage over successive surveys, but this data would at least suggest the demise of Solaris as a popular operating system has started.

Client Operating Systems

The numbers for the responses were as follows:

- **Windows** – 104
- **Linux** – 85
- **Solaris** – 75
- **AIX** – 32
- **HPUX** – 28
- **Mac OS X** – 15
- **NetWare** – 14
- **OpenVMS** – 13
- **Tru64** – 12
- **Irix** – 3

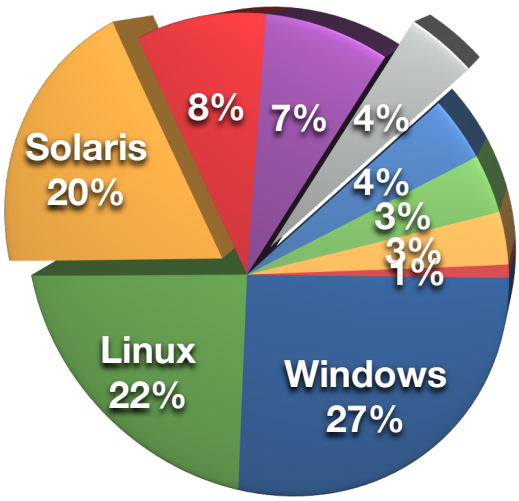

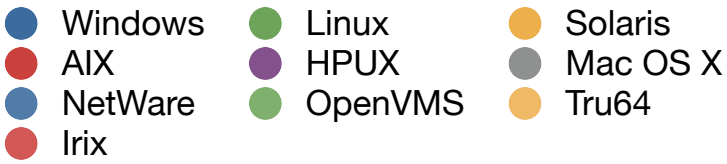

Comments and Conclusions

We are still seeing a reasonably comprehensive spread of client operating systems in NetWorker environments. Linux, Solaris and Windows operating systems retain the lions share again.

For the most part, client operating system spread and usage seems relatively unchanged across surveys conducted thus far.

Sites using Deduplication

Deduplication is not a binary activity; sites may choose not to use it, or use it at the target level, or the source level, or a mix of the two depending on the circumstances at hand.

The 149 responses to this question covered:

- **No** – 68
- **Yes, Source** – 5
- **Yes, Target** – 29
- **Yes, Source & Target** – 4

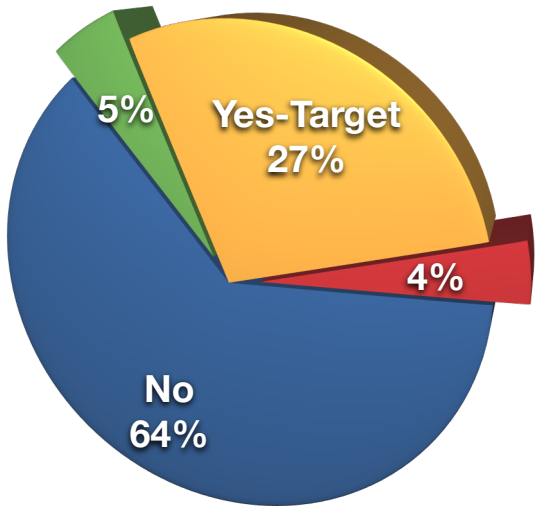

Comments and Conclusions

While it is a strong growth market, deduplication has not yet hit commodity status within the backup environment. Part of this would be attributed to the high cost, regardless of product in use, compared to traditional backup. While there are compelling arguments to say, deploy 10TB of target deduplication storage instead of 50TB of raw target storage, the cost of said deployments mean that deduplication will typically be integrated into an environment as part of a major refresh cycle, rather than an ad-hoc update.

Compared to the November 2010 survey results, there has been some shift in the deduplication take-up, but not a significant amount. For the most part, this is because unlike the transition to say, disk backup systems, the transition to deduplication requires the sort of planning only seen in infrastructure refresh operations. Since these typically occur on a 3-5 year cycle, we will undoubtedly continue to see deduplication rates increase, but at a slower pace than we saw the adoption of backup-to-disk.

Survey	No	Yes - Source	Yes - Target	Yes - Both
Nov 2010	68%	4%	20%	7%
June 2011	64%	5%	27%	4%

Modules in Use

Note that this question is not aimed at determining actual databases/mail servers/etc in use, but to find out what the layout of actual *module* use for those products were.

The numbers for the responses were as follows:

- **Oracle** – 45
- **Notes** – 12
- **NMM** – 25
- **Exchange** – 55
- **MSSQL** – 45
- **Documentum** – 4
- **SAP on Oracle** – 12
- **SnapImage** – 1
- **PowerSnap** – 4
- **DB2** – 7
- **Sybase** – 4
- **EDM** – 0
- **Meditech** – 0
- **Informix** – 4
- **NMDA** – 12

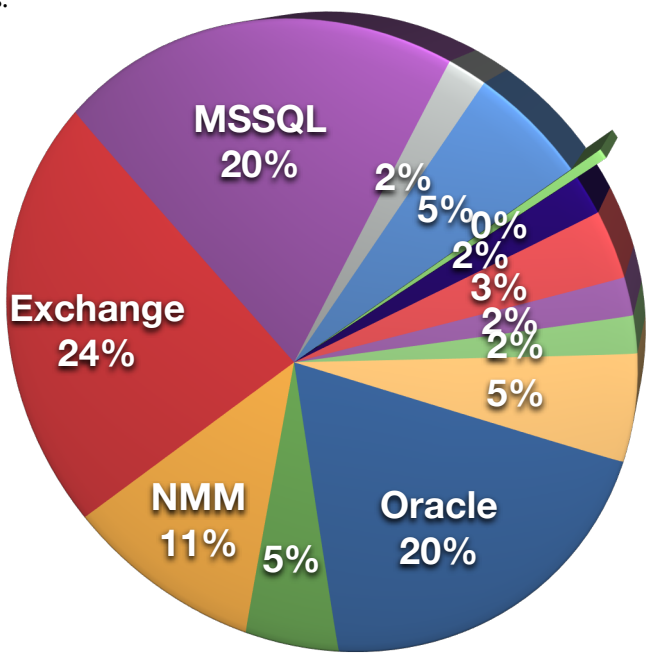

Comments and Conclusions

The continued dominance of particular modules in use (Oracle, Exchange and Microsoft SQL Server) reflect both anecdotal evidence from social networking environments and this author's personal experience that these are the most commonly deployed/used modules within NetWorker datazones of any size.

As expected in previous surveys, we are now seeing an upswing in the usage of NMM – this has grown from 5% in March 2010, to 9% in November 2010, and is now at 11% in June 2011. However, it should be noted that with Exchange and MSSQL percentages relatively unchanged from the November 2010 survey, it would seem likely that a significant number of NMM deployments may be happening around Sharepoint, with customers sticking to traditional module backups of Exchange and MSSQL for the time being. We will continue to track this over time.

Open Source Database Usage

We continue to track the popularity and NetWorker customer thoughts on open source database modules in this survey.

Usage numbers for the were as follows:

- **PostgreSQL** – 19
- **MySQL** – 46
- **Firebird** – 4
- **Others** – 8

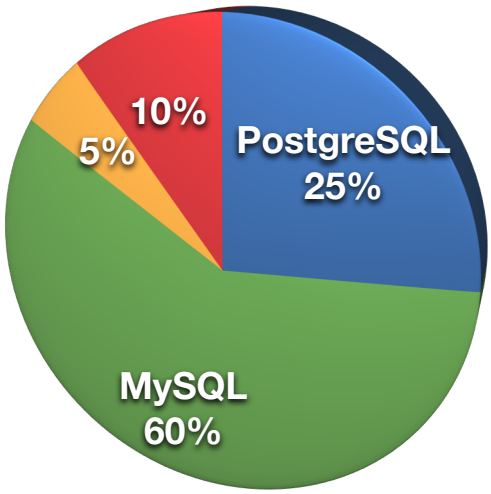

Comments and Conclusions

The survey shows that MySQL database deployments remain relatively popular in sites, with it clearly being an obvious target for a first pass of an "Open Source Backup Module".

Accounting for a skew created in the previous survey from sites not using Open Source databases answering questions on the "how much would you pay (USD) for a module?", the question relating to pricing was expanded to give those sites a voice in the results.

The responses to this question were:

- **Would not use open source DB** – 19
- **Would not use Modul**e – 26
- **<$500** – 20
- **Between $500 and $1000** – 15
- **Between $1000 and $2000** – 2
- **Over $2000** – 2

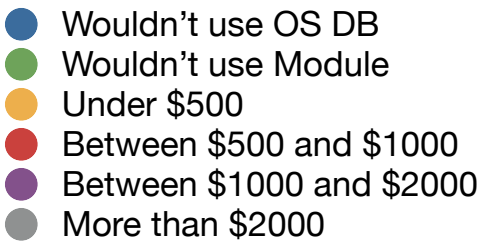

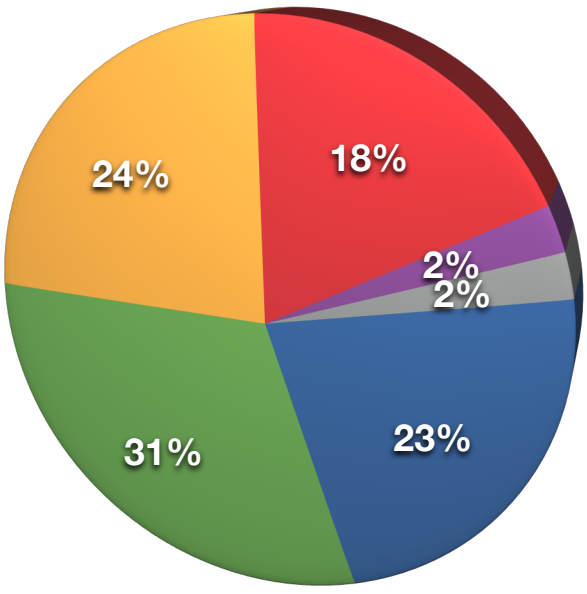

Even taking into account those who would not use open source databases, and those who would not pay for a database module, there remains strong enough numbers to indicate that this an option EMC still needs to consider.

Backup to Disk

The responses were as follows:

- **No** – 19
- **Yes, ADV_FILE** – 61
- **Yes, VTL** – 42
- **Yes, DD Boost** – 7

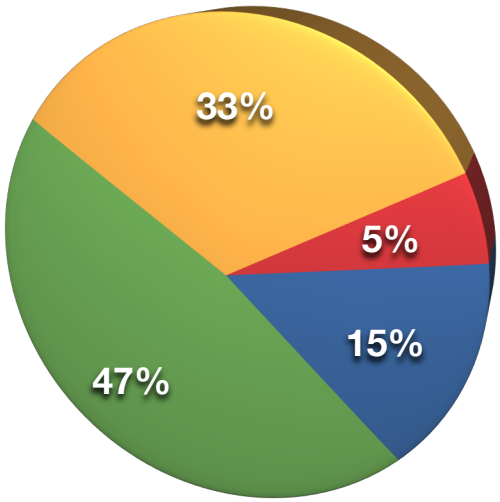

Comments and Conclusions

First, compare the results of this survey and the November 2010 survey:

Survey	No	Yes - ADV_FILE	Yes - VTL	Yes - DD Boost
Nov 2010	16%	52%	32%	N/A
June 2011	15%	47%	33%	5%

With the release of NetWorker 7.6 SP1 and integration of Data Domain boost functionality, it seemed pertinent to include a query on the take-up of that option. This seems to have directly eaten into traditional ADV_FILE disk backup usage – or rather, it seems likely that those sites previously using ADV_FILE style backups are most likely to consider Boost.

That being said, we're seeing a relatively unchanged backup to disk environment in terms of the percentage of sites actually using *some* form of it. This would seem to be consistent with standard product bell curves – we're now in the situation where it's only the late adopters who haven't taken the plunge. Whether they will or not remains to be seen – it may be that a small percentage of companies will stay tape-only.

Cloning policies

The answers were:

- **No, no budget** – 9
- **No, not interested** – 16
- **No, no time** – 12
- **Yes (all)** – 32
- **Yes (just prod)** – 9
- **Yes (very selectively)** – 29

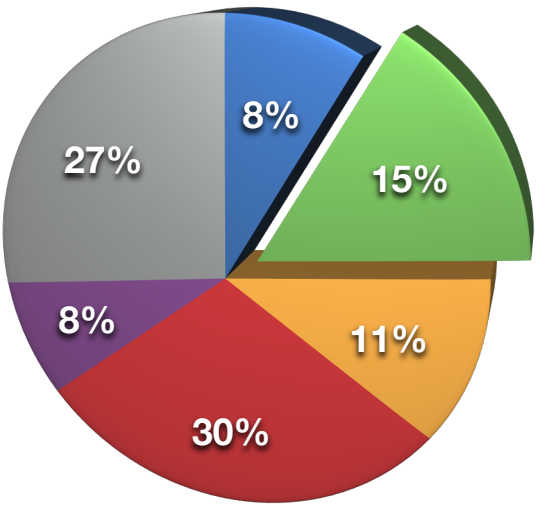

- No (No budget)
- No (not interested)
- No (no time)
- Yes (all)
- Yes (just production)
- Yes (very selectively)

Comments and Conclusions

In the previous survey, 26% of respondents indicated that for one reason or another, they chose not to clone backups. This has climbed somewhat in this survey to 34%, a somewhat disheartening trend.

Companies need to keep in mind the importance of ensuring that a backup system does not represent a single point of failure within the environment, and by not having duplicates of backups, that's exactly what it *is*.

One hopes to see this trend reverse in the next survey.

Longest Retention Period

The purpose of this question was to determine the longest length of time that backups are being kept for.

Available answers were:

- **Less than a month** – 2
- **1-3 months** – 8
- **3-6 months** – 6
- **6-12 months** – 19
- **1-3 years** – 22
- **3-7 years** – 15
- **7-15 years** – 30
- **15+ years** – 6

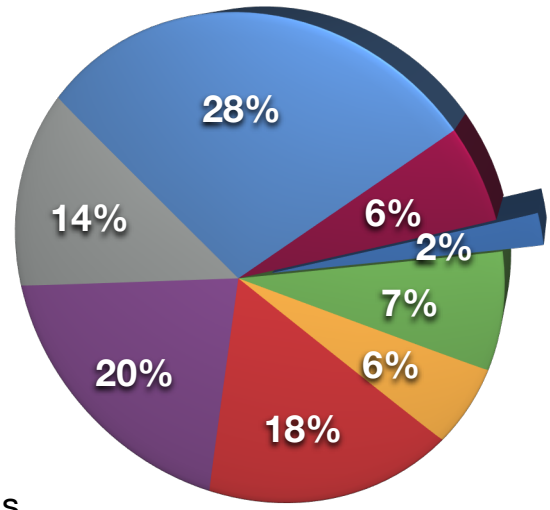

Comments and Conclusions

With almost 50% of respondents having retention periods of 7 or more years, we see compelling evidence of how inaccurate the "tape is dead" meme actually is. While the purpose of tape may be changing in a lot of datacentres (with an increased focus on using disk backup for short term retention cycles, and only pushing out to tape for longer term cycles or when disk backup is full), very few of those businesses with 7+ year retention cycles are going to be able to retain that data on disk only – either safely, or in a cost-justifiable way.

Encryption – Attitude Towards and Usage

New in this survey were questions about attitude towards encryption, and actual adoption within environments. Across the two questions, some interesting results came out:

Question 1 – What is the company attitude towards encryption?

- **Not interested** – 37
- **Desirable, not implemented** – 53
- **Required for offsite backups** – 10
- **Required for all backups** – 5

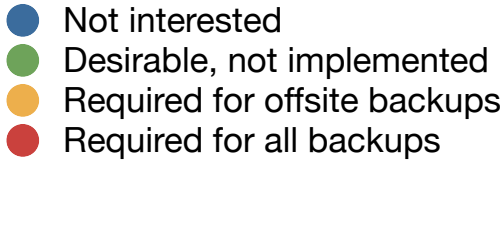

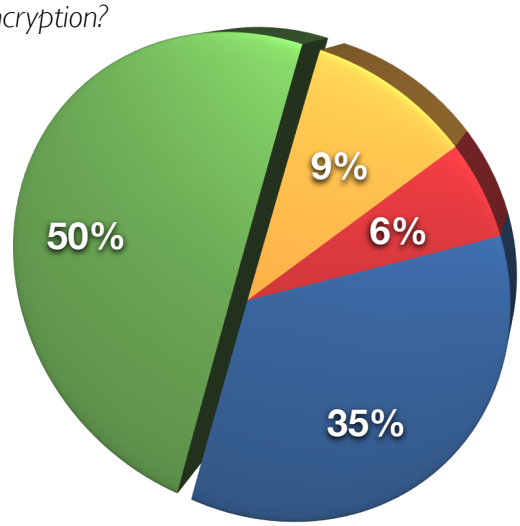

Question 2 – Do you encrypt your backups?

- **No** – 87
- **Yes, in software, on client** – 7
- **Yes, native tape encryption** – 12
- **Yes, FC/SCSI Appliance** – 2

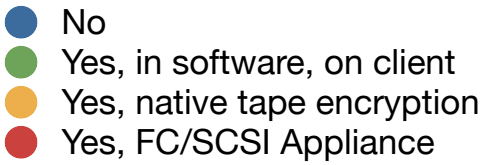

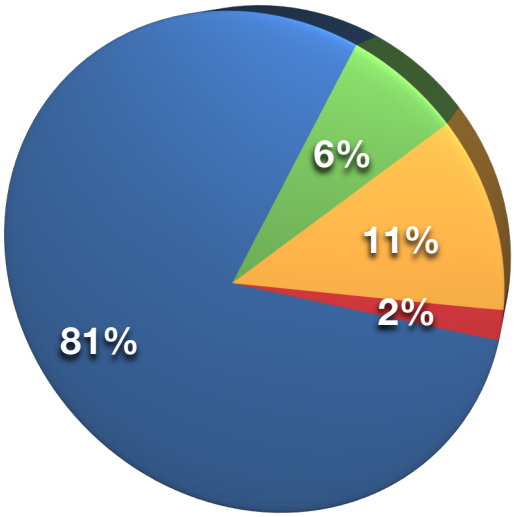

Comments and Conclusions

Tellingly, we see a significant disparity between the number of organisations that would *like* to encrypt their backups and those who actually do. While only 65% of businesses either wanted to, or had a requirement to encrypt backups, 81% of businesses who responded currently *do not* encrypt their backups.

This suggests that the data protection industry still has some way to go in order to make encryption a readily approachable process. One might suggest that deduplication will actually add complexity to this problem for a while, since most forms of encryption will not integrate at all with deduplication systems.

In Conclusion

Without a doubt, NetWorker has a strong user following, and a strong usage base in the enterprise environment. With deployments featuring multiple datazones, thousands of clients with a plethora of operating systems and databases/applications being protected, many organisations, regardless of size, rely on NetWorker daily to ensure successful continued operations.

A continued strong up-take in the current releases of NetWorker (7.6 SP1 and 7.6 SP2) demonstrate that EMC's strong focus on new feature integration in these releases are striking a chord with customers.

Deduplication remains a slowly growing focus in NetWorker customers, and based on the direction that many EMC competitors are going in this area, EMC are likely going to need to aggressively pursue a stronger level of integration between NetWorker and Avamar. That is, it is not unreasonable to suggest that an end-goal should be native source-based deduplication capability in the NetWorker client, rather than separate products.

With the primary focus of the NetWorker 7.6 SP2 release having been to offer full vStorage API integration, it will be interesting to see the focus of product releases over the coming 12 months – and to see whether we'll finally see a NetWorker version 8.