European Commission

Technical assistance study for the assessment of the feasibility of using "points system" methods in the implementation of Ecodesign Directive (2009/125/EC)

TASK 5

Extended Case study:
Data Storage System

Final report

June 2017

EUROPEAN COMMISSION

Directorate-General for Internal Market, Industry, Entrepreneurship and SMEs
Directorate Directorate C – Industrial Transformation and Advanced Value Chains
Unit Directorate C1 – Clean Technologies and Products

E-mail: grow-c1@ec.europa.eu

European Commission
B-1049 Brussels

EUROPEAN COMMISSION

Technical assistance study for the assessment of the feasibility of using "points system" methods in the implementation of Ecodesign Directive (2009/125/EC)

TASK 5

Extended Case study:
Data Storage System

Final report

LEGAL NOTICE

Luxembourg: Publications Office of the European Union, 2017

ISBN number 978-92-79-71411-5
doi:number 10.2873/211016

Printed in Belgium

Prepared by:
Viegand Maagøe A/S and Waide Strategic Efficiency

Project website: points-system.eu

Specific contract no.: 478/PP/GRO/IMA/15/1164 – SI2.722599
Implements Framework Contract: 409/PP/2014/FC Lot 2.

This study was ordered and paid for by the European Commission, Directorate-General for Growth.

Table of contents

List of tables

List of figures

Abbreviations

ASHRAE	American Society of Heating, Refrigerating and Air-conditioning Engineers
CFC	Chlorofluorocarbons
CO2-eq	carbon dioxide equivalent
COM	capacity optimisation method
DC	data centre
EPA	Environmental Protection Agency
EPEAT	Electronic Product Environmental Assessment Tool
GB	Gigabytes
GWP	global warming potential
HDD	hard disk drive
IOPS	input output operations per second
MiBPS	mebibytes per second
Mt	Megatonnes
OS	operating system
PSU	power supply unit
PUE	power usage effectiveness
RAID	redundant array of independent disks
Sb-eq	antimony equivalent
SNIA	Storage Networking Industry Association
SO2-eq	sulphur dioxide equivalent
SSD	solid state drive
TWh	terawatt hours

1. Introduction to Task 5

This objective of this report is to perform essential additional work regarding the Task 4 Case Study on "Data Storage Systems" to allow the European Commission to propose draft final metrics for the energy performance of data storage products as part of potential draft information requirements within a prospective Ecodesign regulation. It makes use of the work developed under Task 4 as its starting technical basis. The report is intended to enable the European Commission to directly consult stakeholders on these metrics. In addition, the work highlights any potential problem areas encountered.

Task 5 is a further development of the metrics and points system for data storage systems developed in Task 4. Through additional research, the technical foundation and evidence base of the metrics has been improved. The metrics are validated for application in possible future policies or regulations through a sensitivity analysis and the product efficiency compared for all available product data. This task does not address issues related to possible future regulations, in particular: which configurations to test per product, how products would demonstrate compliance, and at what energy performance levels.

The design goals of the metrics are established in Task 4 and it is only effective within these parameters i.e.:

- To allow comparison between different storage products and configurations suitable for the needs of the user, in terms of capacity, active IO performance and other functionality. This means the metrics cannot indicate what configurations are most suited to the use case and may provide confusing results for inappropriate configurations
- The energy performance test results do not give the absolute power consumption of the product, or the product active IO performance. This means it is not possible to analyse these factors and for the inability to do so to be taken into consideration when combining the test results to develop a metric
- There are three metrics for three distinct applications, capacity, transaction and streaming based on the energy performance tests.

The energy performance tests and applications are summarised below.

Table 1 Use of the energy performance tests for the applications considered (Y: Yes, N: No)

Application	Active energy performance tests		Idle energy performance tests
	Hot band test IOPS/W	Sequential test MiBPS/W	Ready idle capacity GB/W
Capacity applications	N	N	Y
Transactional applications	Y	N	Y
Streaming applications	N	Y	Y

2. Task 5A reassessment of all data and redefinition of the storage metric

Task 5A reassesses all the data and redefines the storage metric with the addition of existing research literature. The task is divided into the following steps:

Subtask 5A(a):

- The energy performance test results are analysed in relation to the storage product characteristics to understand the relationship between the factors and how they vary. This is the foundation of the metric
- A mathematical approach and formulae for the metrics are proposed based on the energy performance test analysis
- The weighting for the metric formulae are proposed based on real data centre usage to produce a more realistic representation
- Sensitivity analysis of the weighting to understand how changes affect the relative efficiency of the products.

Subtask 5A(b):

- Development of COM metric factor based on analysis of how common COM functionality is found in products and their impact on capacity savings and active mode performance.

2.1 Analysis of performance test results trends

This section provides a basic analysis of the energy performance test results across all the test data taken from ENERGY STAR in May 2017. This is required to understand how the energy performance varies, the magnitude of the variations and what correlations exist with the product configuration. The products are categorised based on the following attributes:

- Product classification i.e. online 2, online 3, online 4 (abbreviated to OL2, OL3, OL4). A higher number generally indicates a higher number of drives and a greater number of features
- Type of drives used in the product, based on rotational speed and SSD. Some products are configured with a mix of different drives and are also identified as such, although the exact mix is not specified in the analysis. The rotational speed is measured in thousands of revolutions per minute (e.g. 7.2K indicates 7200 rpm) and varies between 7.2K to 15K. A higher rotational speed generally indicates higher IO performance but at increased energy consumption. In addition, the capacity of higher speed drives is lower
- Total raw capacity.

The drive format, i.e. whether it is 3.5″ or 2.5″ is not considered. We also do not categorise based on the type and number of controllers due to data reporting issues identified by The Green Grid.

2.1.1 Transactional performance test analysis

Figure 1 shows the variation of the active performance and ready idle with the drive count. Due to the large variation, the axes are both logarithmic. There are 69 different products and configurations which still represents a relatively small dataset given the very large number of configuraitons and products available.

The first observation is that there is no clear correlation between the active performance and ready idle tests across all products. In addition, there is significant variation in the performance, from 4-150 IOPS/W for the (active) hot band test and 1.5- 600 GB/W for the ready idle test.

There are some general trends and clustering of product categories which are expected:

- The SSD have a higher active efficiency with IOPs/W generally above 64
- 7.2K rotational hard drives have a higher ready idle performance, while 15K drives have the worst.

Within a category (a single colour and shade) the trends are not clear. The limited number of OL2 and OL4 datapoints also means no trends can be discerned.

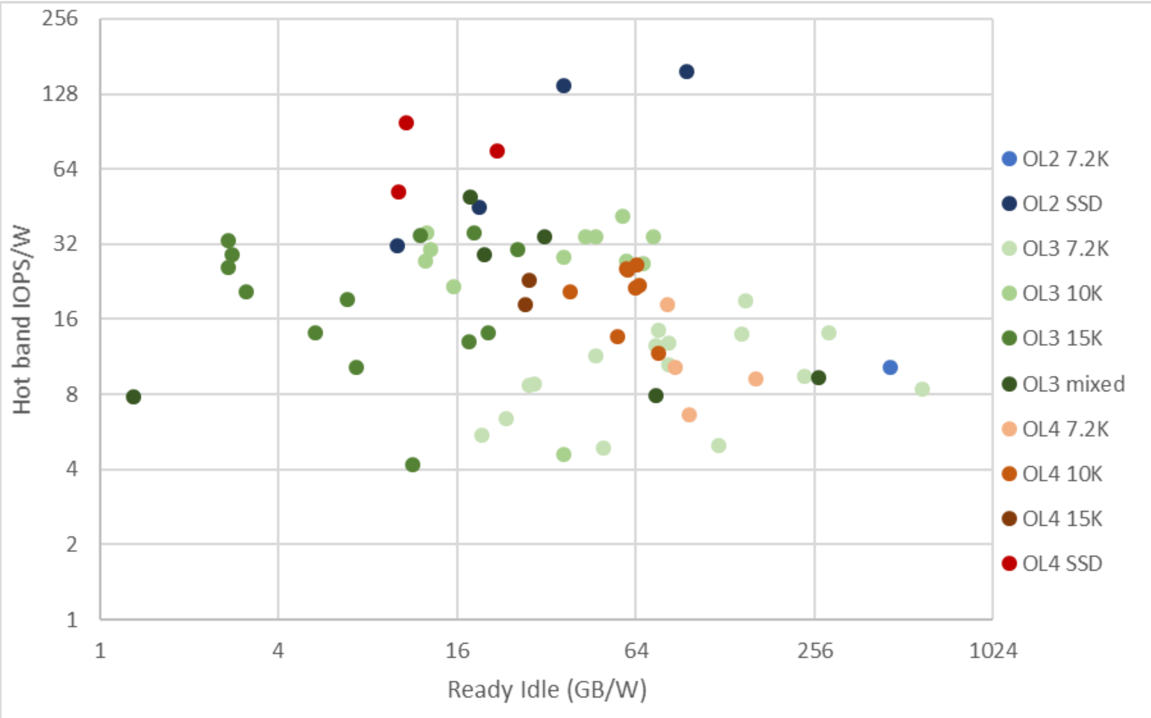

Figure 1 Relationship between active hot band performance and ready idle performance

Plotting the drive count against raw capacity (Figure 2) shows that in general, a greater number of drives results in higher total raw capacity. This is expected because a large number of products use similar capacity drives. However, the 7.2K drives generally have a higher capacity per drive and therefore tend to lie above the trendline. This will contribute to a higher ready idle efficiency.

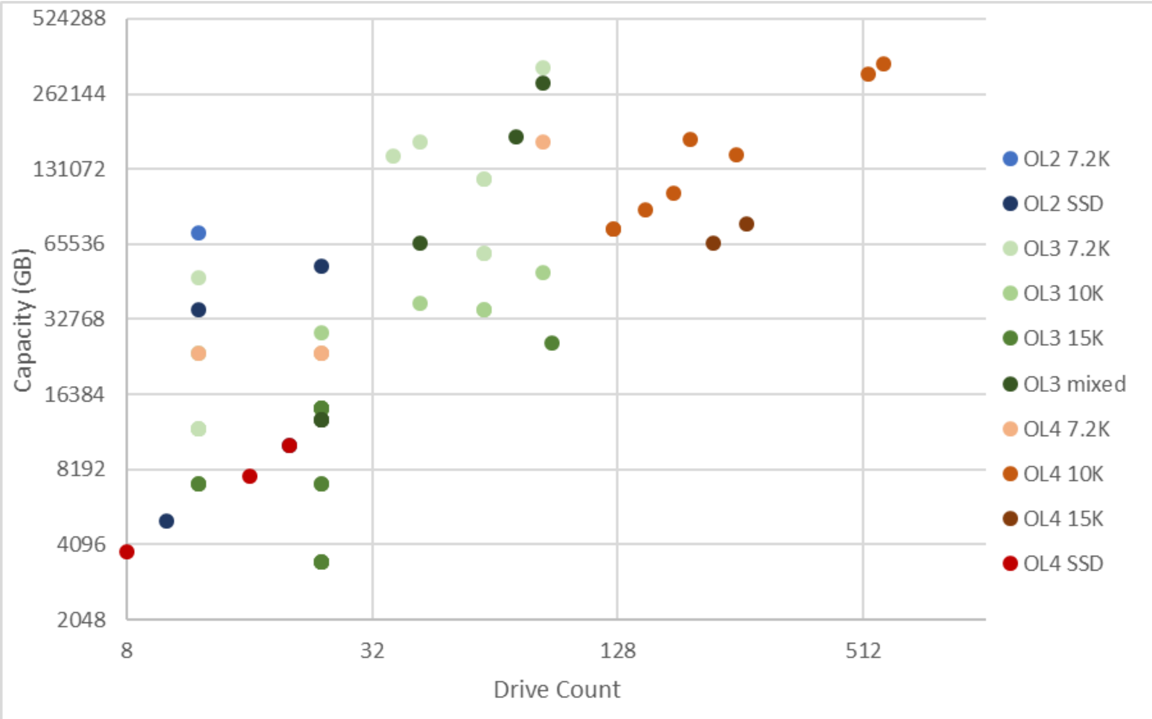

Figure 2 Relationship between drive count and raw storage capacity

Figure 3 and Figure 4 show that there is much more variation in performance for products with lower capacity and drive count. This is partly due to the higher number of data points but also indicates greater variation in products to which the metric must be applicable. SSD products also have a much higher active efficiency, due to the very high IO they can achieve compared to HDDs. However, the idle efficiency is similar to other products.

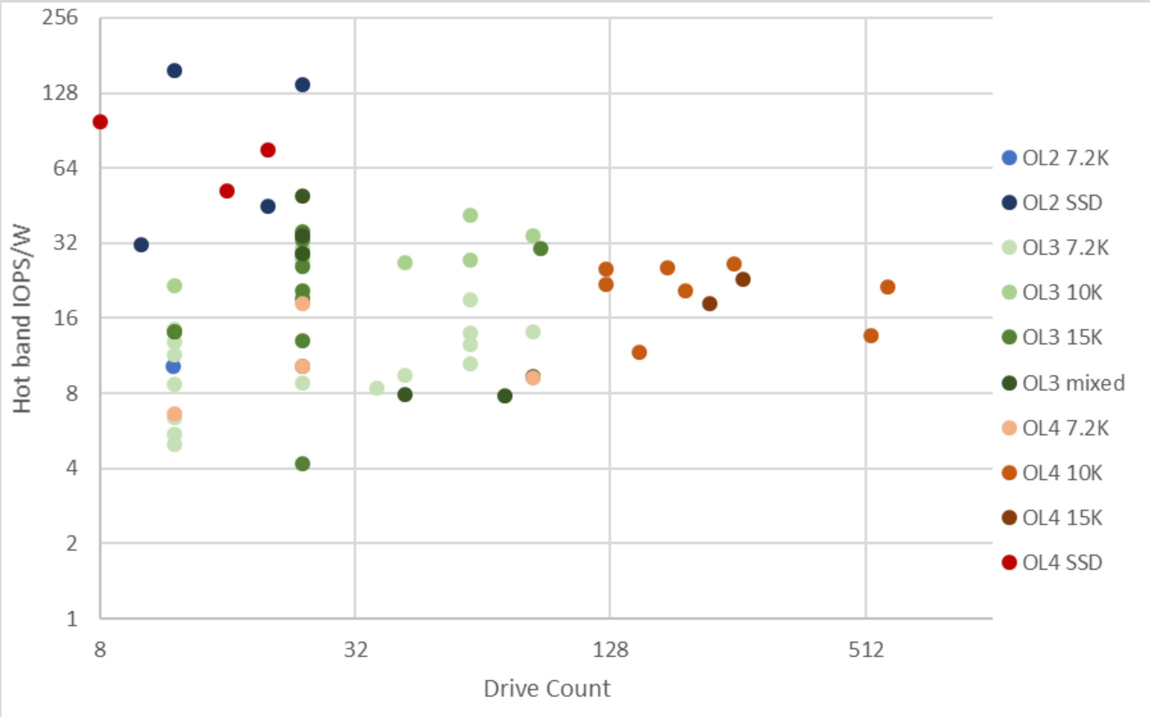

Figure 3 Relationship between drive count and active hot band energy performance

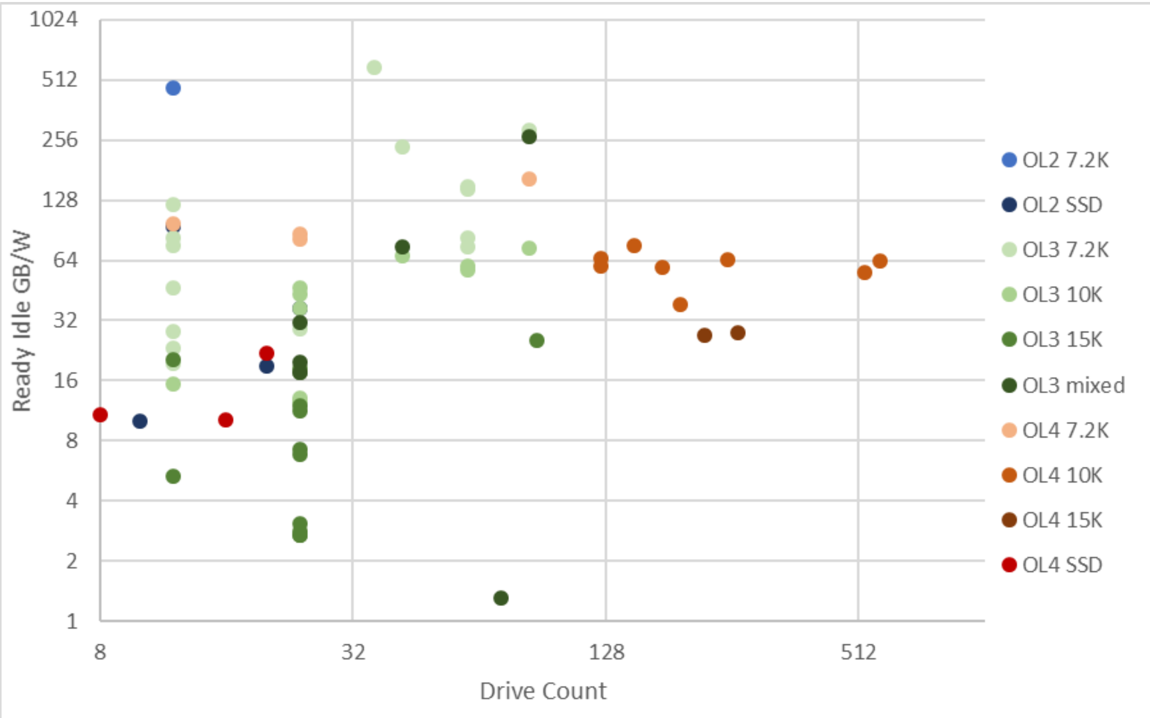

Figure 4 Relationship between drive count and ready idle energy performance

The ready idle performance and capacity shows one of the clearest correlations (Figure 5). However, the scales are logarithmic, which means that the variance in this graph is still very large.

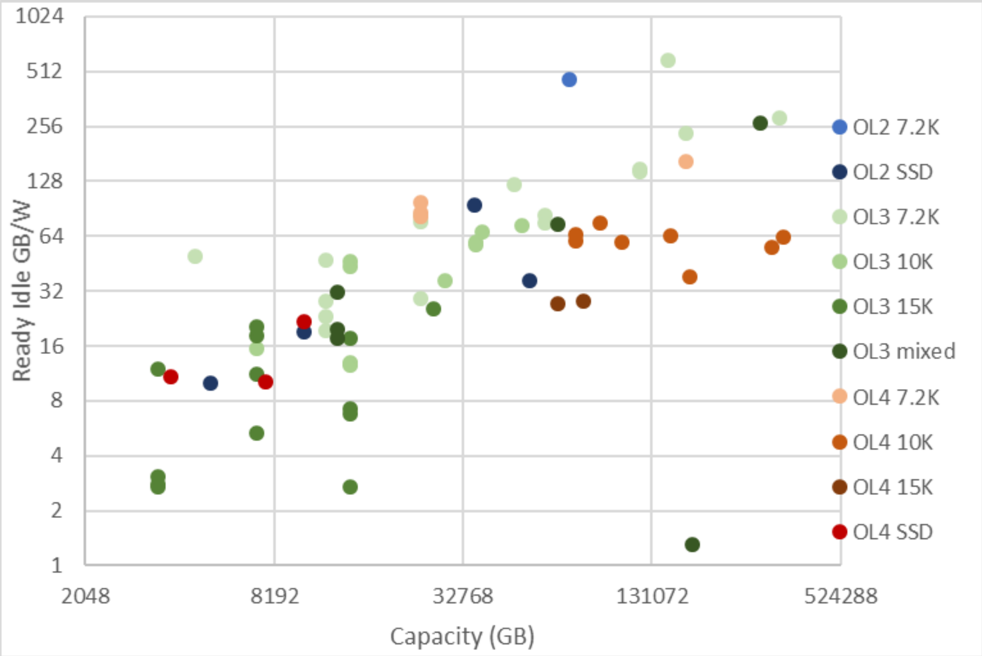

Figure 5 Relationship between capacity and ready idle energy performance

Based on this analysis, it is clear that the active and idle performances are mostly independent of each other and that the correlation between any single factor is weak. The numerical values can vary widely between products which appear relatively similar in this analysis.

The exact performance is the result of a combination of factors, which it has not been possible to quantify such as the number and types of controller and other product design factors. These factors are also likely to influence the suitability and buying decisions of the user. The metric must therefore treat the active and idle performance as two separate and important aspects in the overall metric. A metric which results in overemphasis on any particular test is unlikely to be applicable to a broad range of situations and will fail to provide a useful indicator of efficiency to consumers.

2.1.2 Streaming energy performance test analysis

There are only 33 datapoints for streaming optimised storage products, and the lack of data makes correlation trends harder to identify.

The read performance efficiency (Figure 7) shows that efficiency drops as the drive count increases, although low performance products exist at every level of drive count. The 7.2K drives tend to have a lower performance efficiency but 15K drives do not have a clear performance advantage.

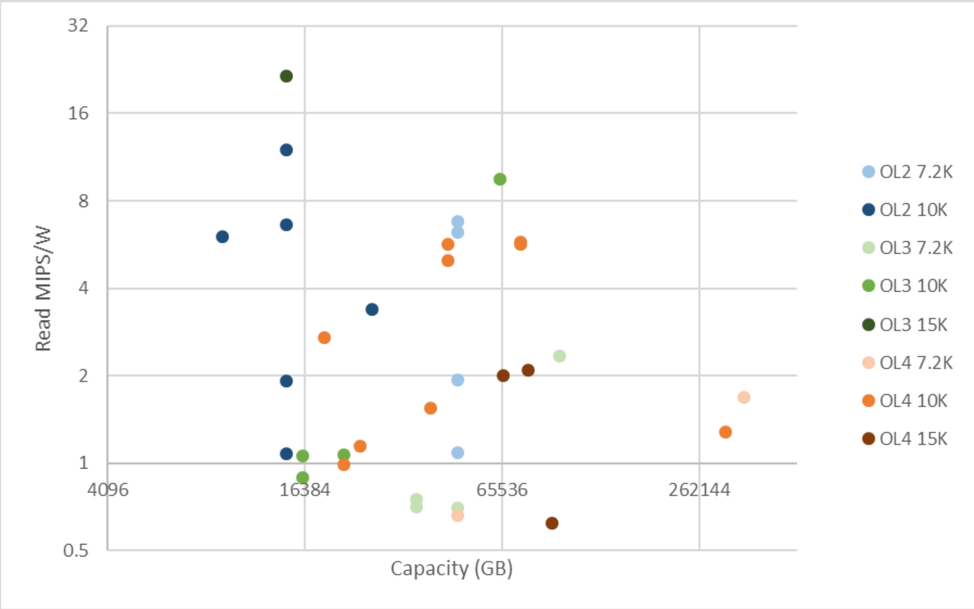

Figure 6 Relationship between storage capacity and data read energy performance

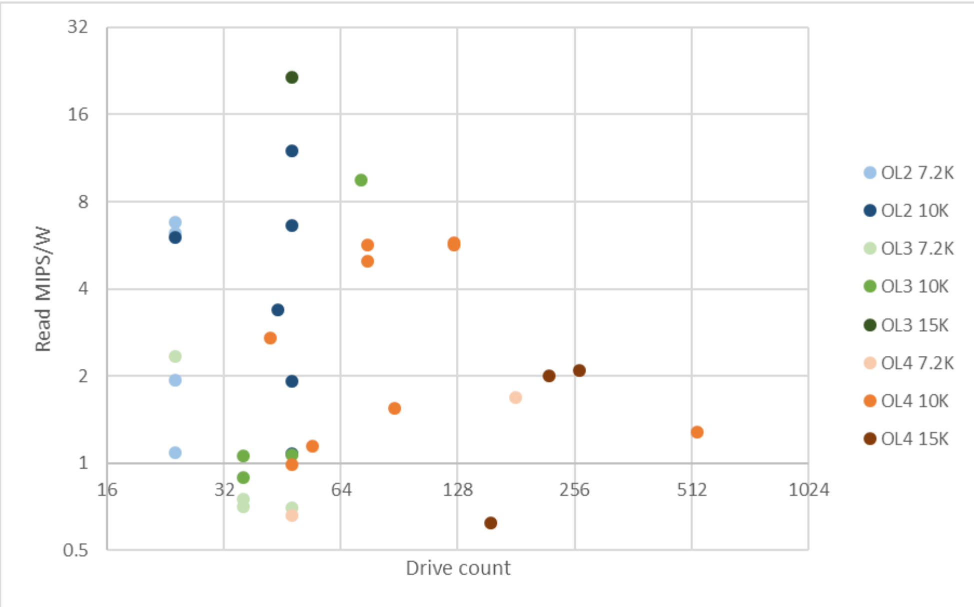

Figure 7 Relationship between drive count and data read energy performance

Ready idle performance is relatively similar for all products, and the majority of products lie between 32 and 64 GB/W. The 7.2K drives tend to have a higher ready idle performance, similar to the transactional analysis. There are a couple of outliers with ready idle levels of around 512 GB/W and the reason for this is not clear and may be a data error. The full data set has been used in this Task 5 analysis without any removal of potential faulty data. When compared to the analysis presented in Task 4, in the present task it should also be emphasised that the highest performing 10% of products has *not* been excluded.

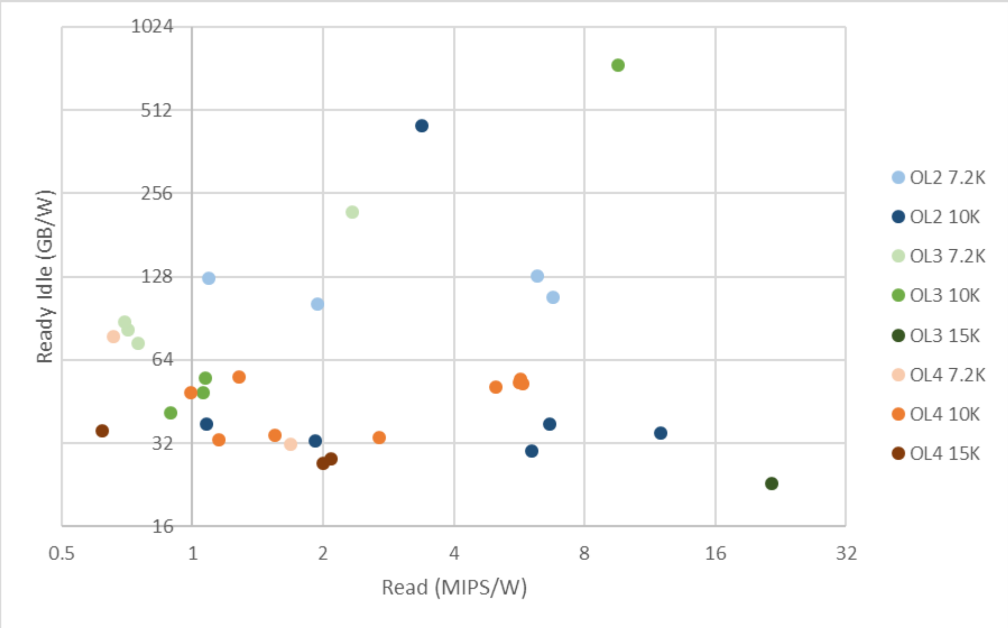

Figure 8 Relationship between data read energy performance and ready idle energy performance

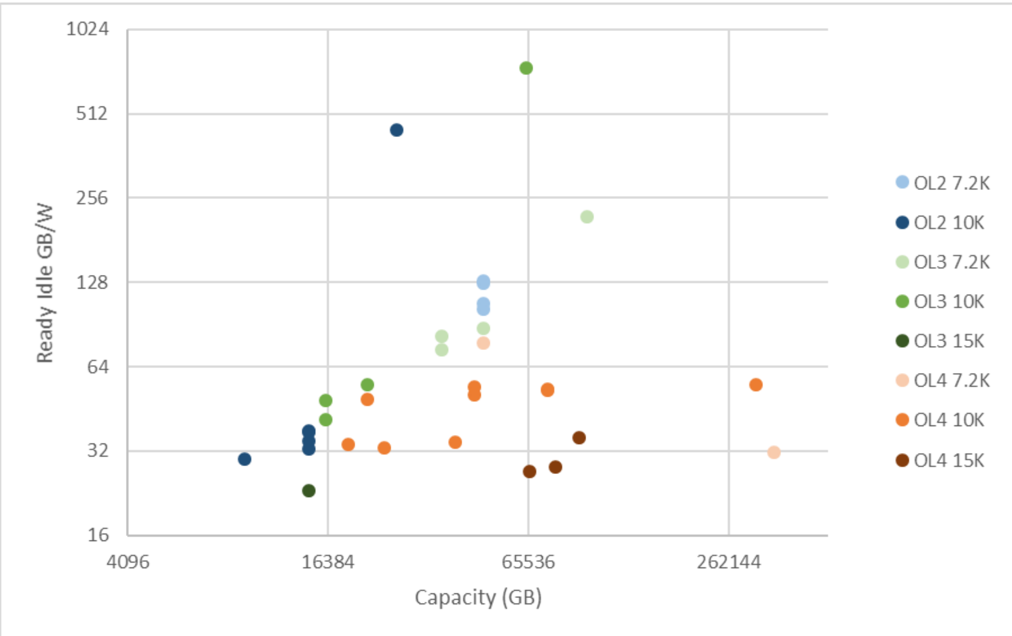

Figure 9 Relationship between storage capacity and ready idle energy performance

Between the read and write test (Figure 10), there is a good correlation which is expected since the read and write speed of drives are linked. This means that the sensitivity between read and write weightings for the metric is low.

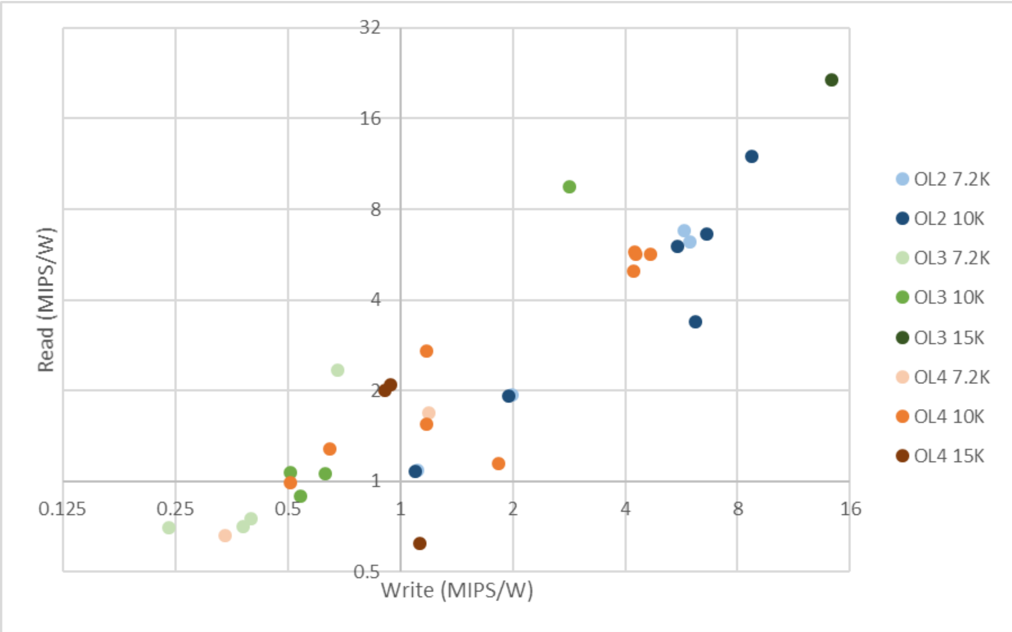

Figure 10 Relationship between data write and read energy performance

The capacity and drive count (Figure 11) shows the same trends as the transactional analysis and shows a good correlation.

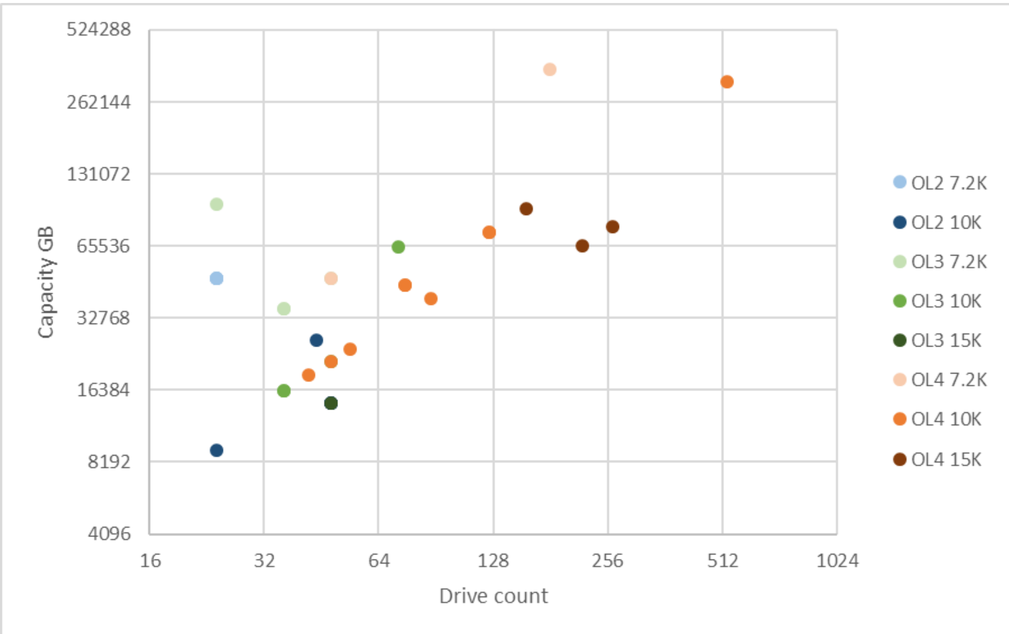

Figure 11 Relationship between drive count and capacity for streaming optimised products

2.2 Mathematical metric approach

The mathematical approach needs to assess the active and idle performance independently without overemphasis on either. Task 4 is an adaptation of the duty profile and uses the arithmetic mean. However, a duty profile is calculated from the time spent in each mode, and the power consumed in each mode, measured in *Watts*. The total energy is therefore the arithmetic sum.

The performance tests are measures of active and idle mode, and are not representative of power consumption when compared against each other since they are measured *per Watt*. In addition, they have very wide ranges and different magnitudes. For example, the OL2 7.2K product has a IOPS/W of 11.4 and ready idle of 460 GB/W. This means the ready idle performance indicator will have a much more significant impact on the mean, and will favour products with very high ready idle energy performance regardless of the importance of the active performance. This could encourage manufacturers to chase ready idle efficiency at the cost of active performance and be unrepresentative of real use and therefore uninformative. Conversely, other products have a ready idle of less than 5 GB/W and any weighting to reduce the impact of the high ready idle will reduce the low scores to insignificance.

If the IOPS and GB were known, the sum or arithmetic mean could be used since it would effectively give the power consumption. However, without this, the arithmetic calculation does not give a meaningful value. Therefore, the geometric mean is used since the two factors are two different, unrelated measures of performance:

$$\bar{x} = \left(\prod_{i=1}^{n} x_i^{w_i}\right)^{1/\sum_{i=1}^{n} w_i} = \exp\left(\frac{\sum_{i=1}^{n} w_i \ln x_i}{\sum_{i=1}^{n} w_i}\right)$$

This can then be weighted based on the relative importance of the active and idle factors.

By ensuring the weighting factors sum to 1, the equation can be simplified to:

$$\bar{x} = \left(\prod_{i=1}^{n} x_i^{w_i}\right)$$

2.2.1 Transactional metric

$$transactional\ metric = hot\ band\ performance^{active\ utilisation} \times ready\ idle\ performance^{idle\ utilisation}$$

2.2.2 Streaming metric

$$streaming\ metric = write\ perf^{write\ utilisation} \times read\ perf^{read\ utilisation} \times ready\ idle\ perf^{idle\ utilisation}$$

2.2.3 Capacity metric

$$capacity\ metric = ready\ idle\ performance$$

2.3 Weighting of performance tests in metric

This section proposes the weighting of the performance test based on real use data. The relative weighting is based on the utilisation level of the product. A storage product which is not being accessed will be in ready idle and therefore a higher ready idle test results in greater savings. While in active use, the IOPS/W is the most important factor.

Since the active test is based on 100% load, and the product will generally operate significantly below this, the active use is also affected by the idle efficiency. The weighting is therefore a function of both the time spent in active and idle modes and the intensity of use while in active mode.

A valid metric also needs to apply to a wide range of situations and transactional loads. A relatively limited amount of research is available that shows the utilisation levels of storage products under different workloads, and is discussed in this section.

2.3.1 Utilisation data analysis

The storage utilisation is closely related to the duty profile of the storage device, as discused in Task 4. However, rather than providing the time in a particular state such as active and idle, it tries to ascertain what proportion of the total performance is being used over time. This is achieved through various methods, including the number of concurrent data requests, the usage of the individual storage devices within the storage product, and the disk I/O rate. Utilisation is used, since it is better aligned to the mathematical metric approach.

Figure 12 shows a typical duty profile and presents the activity levels for a variety of workloads as well as how it varies across the day and week. It is taken over a two month period, from two large scale corporate enterprise file servers covering 1000 employees. The first three activities closely follow the working week, with little to no activity in evenings and weekends. The backups and updates, however, have a steady activity over the day and the backups peak at weekends. A simple interpretation of this information based on activity during 8am-6pm work hours gives an absolute minimum utilisation level of 50hr/168hr, or just 30%.

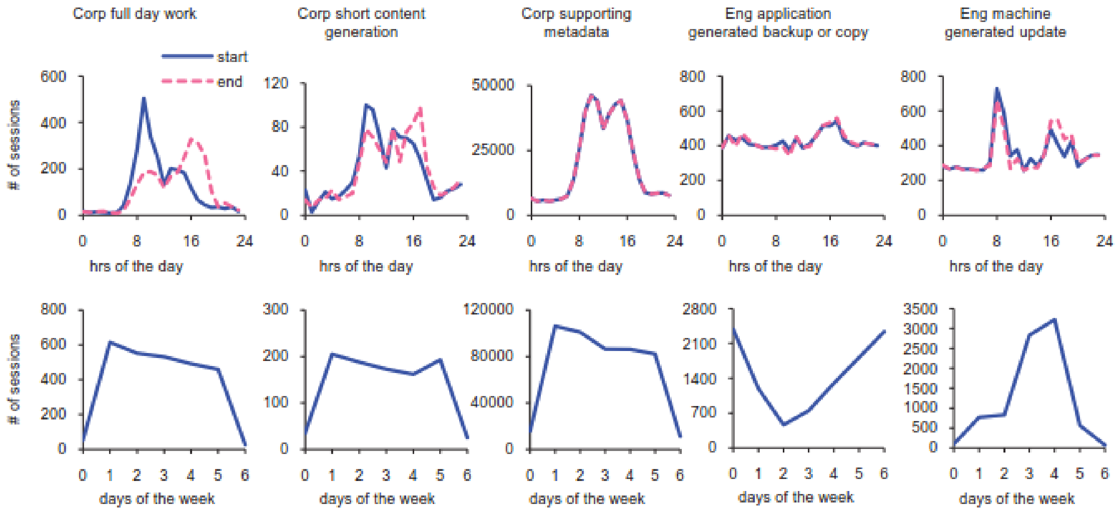

Figure 12 DC enterprise storage activity profiles. Source: Chen, 2011[1]

Figure 13 shows the activity of the individual drives that are active in an enterprise storage system and is representative of a large number of small to medium size enterprise

[1] Chen, Y et al (2011) *Design implications for enterprise storage systems via multi-dimensional trace analysis*, SOSP '11 Proceedings of the Twenty-Third ACM Symposium on Operating Systems Principles

data centres. This shows that the mean is 60% including read and write. The median value is also very close to the mean and may also be a useful indicator of utilisation. The purpose of this research is to demonstrate that the writes are responsible for a significant part of the utilisation and that methods can be developed to reduce the write utilisation and save energy. Therefore, the 60% value is interpreted as an upper bound for the active utilisation.

Table II. Number of Concurrently Active Volumes

	Mean	Median	99th pctile	Max
Read/ write	21.7 (60%)	22 (61%)	27 (75%)	31 (86%)
Read-only	7.6 (21%)	7 (19%)	15 (42%)	22 (61%)

Numbers in parentheses show the number of active volumes as a percentage of the total number of volumes (36).

Figure 13 Disk activity in storage system. Source: Narayan, 2008[2]

Figure 14 and Figure 15 show the utilisation level for Microsoft servers running Hotmail and Messenger over a week relative to peak utilisation. While these are old services that are no longer available, they are examples of heavily-used services running across tens of thousands of servers that may be more indicative of public cloud-type applications. Figure 14 shows the disk I/O performance follows a diurnal pattern and spends the majority of time between 40-70% utilisation. Figure 15 is a cumulative distribution function graph which shows the proportion of time spent at each utilisation level. The fraction of time accumulates to 100% as the time spent at each utilisation level is summed together. A steeper gradient at a particular utilisation level shows that the amount of time spent is large.

It shows that utilisation is never 0% and has a minimum of approximately 30%. From the report, the average value is calculated at 42% (Hotmail) and 60% (Messenger). However, this utilisation is relative to the peak hour and may be less than the absolute maximum utilisation of the storage product.

[2] Narayanan, D., Donnelly, A., and Rowstron, A. (2008) *Write off-loading: Practical power management for enterprise storage*. ACM Trans. Storage 4, 3, Article 10 (November 2008)

Based on this data, the weighting upper bound is around 60% and the mean between 40-60% for the active performance test.

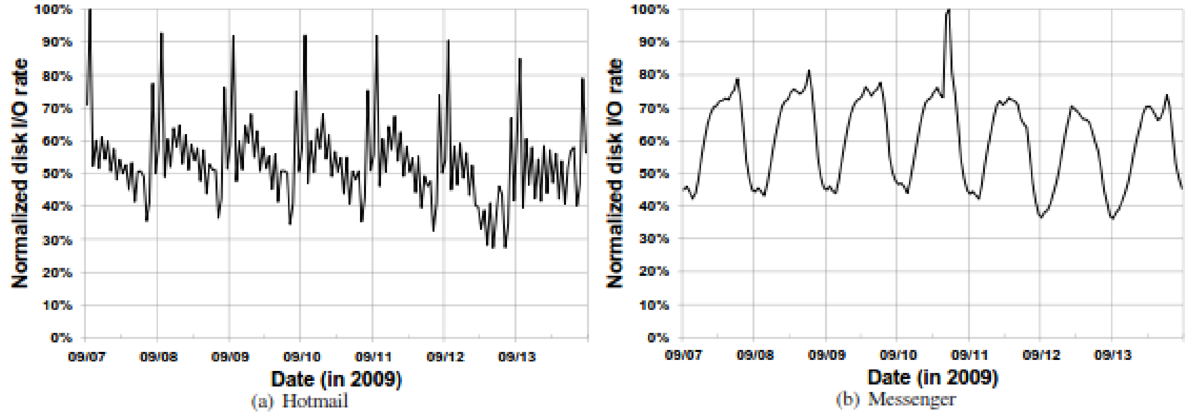

Figure 14 Disk I/O rate relative to peak utilisation for large services. Source: Thereska, 2011[3]

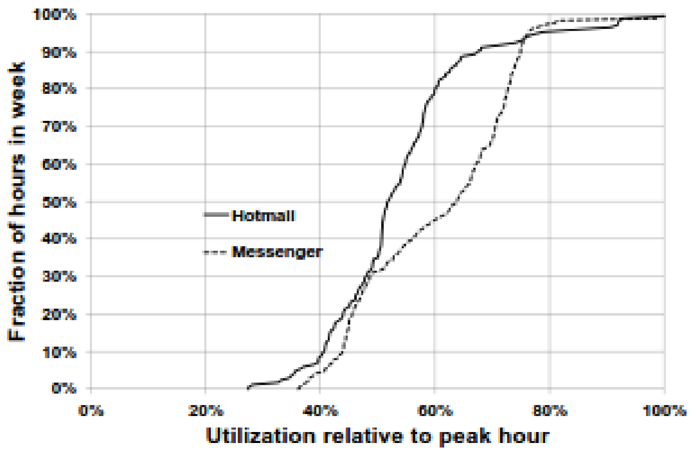

Figure 15 CDF utilisation for large services. Source: Thereska, 2011[4]

2.3.2 Proposed weightings

There is limited data available but there is relatively strong agreement in the utilisation level between 50-65% and an absolute lower bound of 30%. The proposed weighting is therefore 50% active:idle.

2.3.2.1 Transactional metric

$$transactional\ metric = hot\ band\ perf^{0.5} \times ready\ idle\ perf^{0.5}$$

2.3.2.2 Streaming metric

$$streaming\ metric = write\ perf^{0.17} \times read\ perf^{0.33} \times ready\ idle\ perf^{0.5}$$

The relative weighting of the write and read is based on the 2:1 ratio, which is similar to the generally accepted 70:30 ratio, as used in the SNIA Emerald[5]. Due to the good cor-

[3] Thereska, E., Donnelly, A., and Narayanan, D. (2011) *Sierra: practical power-proportionality for data center storage.* ACM 978-1-4503-0634-8/11/04

[4] Thereska, E., Donnelly, A., and Narayanan, D. (2011) *Sierra: practical power-proportionality for data center storage.* ACM 978-1-4503-0634-8/11/04

relation of the read:write energy performance (Figure 10), knowing the precise ratio will only have a small impact on the overall metric.

2.3.2.3 Capacity metric

$$capacity\ metric = ready\ idle\ performance$$

2.4 Sensitivity analysis

This sensitivity analysis indicates how sensitive the metric is to the utilisation ratio and what impact the ratio has on the overall metric. A high sensitivity would mean that the utilisation ratio level strongly affects the results and will only be representative of a narrow band of real life situations which match the utilisation ratio. Additional research may then be necessary to further refine the utilisation level. Conversely a low sensitivity means that the metric is not very sensitive to the exact utilisation ratio, and thus the metric can be applied to a range of situations.

2.4.1 Transactional

The efficiency metric is calculated for Online 3 storage products with 24 or fewer drives (Table 2). These devices were selected due to the relatively large data set of similar configurations, which would highlight any sensitivity. The first three columns describe the configurations of the devices and then the active and idle performance tests. The following five columns give the metric values at different utilisation ratios. The columns are then shaded based on the metric value, with dark green being the highest and dark red the lowest. The table is ordered from highest to lowest for the 50% active:idle metric calculation.

From the shading, the table gives a visual indication of the ranking, which are generally similar across all the metric utilisation ratios. The differences become most apparent at 70% active at which point the average difference in ranking is 4 places out of 29. At 30% active there are also a few noticeable differences and the average difference is 3 places. This suggests that the performance is not very sensitive to the utilisation ratio and a ratio of 50% will give a meaningful result that is applicable across utilisation ratios ranging from 60-40%.

[5] www.snia.org/sites/default/files/UserGuideEmeraldMeasurementSpecV1_0.pdf

Table 2 Calculated efficiency score of OL3 products for different utilisation weightings

Trans Device Rated Speed (RPM)	Trans Total Num Installed Storage Devices Optimal Point	Trans total capacity	Trans Optimal Point Hot Band Workload Test (IOPS/W)	Trans Optimal Point Ready Idle Workload Test (GB/W)	Trans metric 70% active, 45% idle	Trans metric 60% active, 40% idle	Trans metric 50% active, 50% idle	Trans metric 40% active, 60% idle	Trans metric 30% active, 70% idle
10000	24	14400	34.41	46.74	37.72	38.89	40.10	41.35	42.64
10000	24	14400	34.42	43.3	36.87	37.73	38.61	39.50	40.42
7200	12	24000	14.42	76.38	23.78	28.09	33.19	39.21	46.32
7200	12	24000	12.91	82.61	22.53	27.12	32.66	39.32	47.34
10000	24	28800	28.4	36.6	30.65	31.43	32.24	33.07	33.92
15000	24	7200	35.6	18.3	29.16	27.28	25.52	23.88	22.34
7200	12	48000	5	122.3	13.05	17.96	24.73	34.04	46.87
7200	12	12000	11.4	46.9	17.43	20.07	23.12	26.64	30.68
10000	24	14400	35.7	12.7	26.18	23.61	21.29	19.20	17.32
15000	24	3504	34.6	12	25.18	22.65	20.38	18.33	16.49
10000	24	14400	30.4	13	23.56	21.64	19.88	18.26	16.77
10000	24	14400	27.3	12.5	21.60	19.97	18.47	17.08	15.80
10000	12	7200	21.7	15.5	19.62	18.97	18.34	17.73	17.15
15000	12	7200	14.1	20.3	15.73	16.31	16.92	17.55	18.20
7200	24	24000	8.8	29.1	12.60	14.20	16.00	18.04	20.33
7200	12	12000	8.7	28.1	12.37	13.91	15.64	17.58	19.77
7200		4608	4.88	49.9	9.80	12.37	15.60	19.69	24.84
15000	24	14400	13	17.5	14.21	14.64	15.08	15.54	16.01
10000		256	4.6	36.49	8.56	10.53	12.96	15.94	19.60
7200	12	12000	6.4	23.3	9.43	10.73	12.21	13.90	15.81
15000	24	14400	19.3	6.8	14.11	12.72	11.46	10.32	9.30
7200	12	12000	5.5	19.4	8.03	9.11	10.33	11.72	13.29
15000	24	3504	33.1	2.7	15.61	12.15	9.45	7.36	5.73
15000	24	3504	29.1	2.8	14.42	11.41	9.03	7.14	5.65
15000	12	7200	14.2	5.3	10.57	9.57	8.68	7.86	7.12
15000	24	14400	10.3	7.3	9.29	8.97	8.67	8.38	8.09
15000	24	14400	25.7	2.7	13.07	10.44	8.33	6.65	5.31
15000	24	3504	20.6	3.1	11.67	9.66	7.99	6.61	5.47
15000	24	7200	4.2	11.3	5.65	6.24	6.89	7.61	8.40

2.4.2 Rank correlation

A more rigorous analysis can be made using rank correlation. The rank correlation gives a ranking for each product in the entire dataset, with the most efficient ranked 1 and the second ranked 2 etc. The similarity in ranking can then be assessed using the rank correlation formula. Table 3 shows the Spearman rank correlation for both streaming and transactional applications compared to the 50:50 weighting. Based on this analysis weightings of 60:40 active:idle and 40:60 active:idle, shown in blue, have a correlation of over 90% and therefore the 50:50 metric provides a useful indicator over the average range of utilisation identified in section 2.3.1.

Table 3 Metric Spearman rank correlation compared to 50:50 weighting

Active weighting	*0.7*	*0.6*	*0.5*	*0.4*	*0.3*
Idle weighing	*0.3*	*0.4*	*0.5*	*0.6*	*0.7*
	Correlation against 50:50 weighting				
Capacity	82%	93%	100%	98%	92%
Streaming	93%	96%	100%	93%	85%

2.5 Capacity optimisation methods weighting

Capacity optimisation methods (COMs) increase the utilisation of the raw data capacity by allowing more useful data to be stored. This can be considered to be the 'usable capacity' and gives a new metric for ready idle in terms of usable capacity/W.

2.5.1 COM availablility

The number of type of COMs varies between products and their application. Figure 16 shows that delta snapshots and thin provisioning are by far the most common COMs available and that they are most commonly available on transaction optimised products which tend to have the greatest number of COMs. Streaming and deduplication is more likely to be available on streaming products.

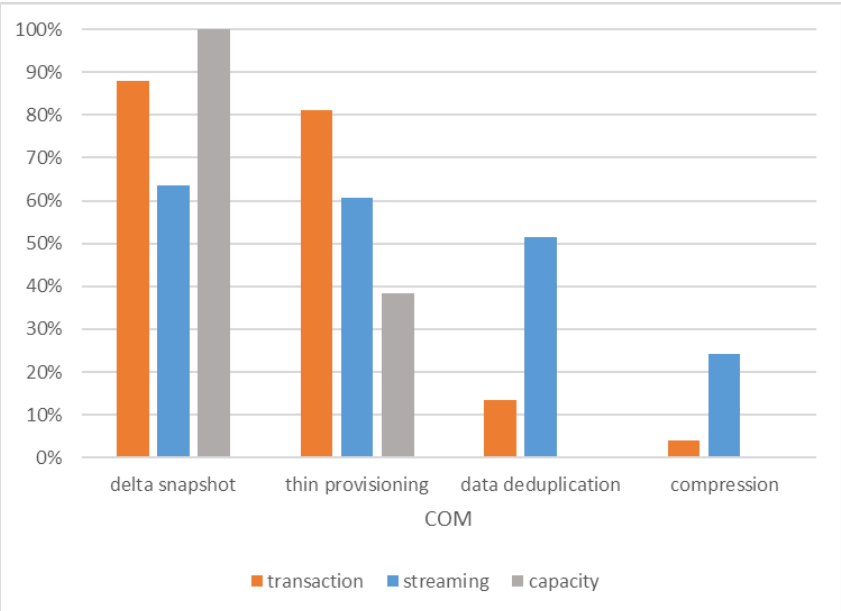

Figure 16 Availability of different COMs for different storage product types

The majority of transactional optimised products have 2 different COMs available, while streaming applications have a more even distribution from 0 to 4 (Figure 17). Capacity optimised products tend to have 1 or 2 COMs. This is important because it can influence the metric's ability to reward and encourage a greater number of COMs, assuming they provide real world benefits.

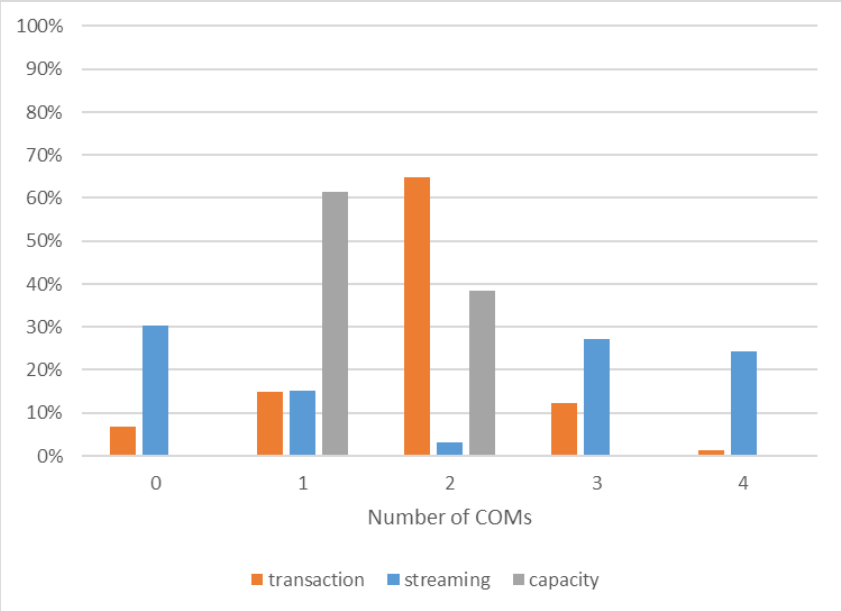

Figure 17 Number of COMs available for different storage product types

2.5.2 COM benefits

There is limited information about the effectiveness of each COM. The figure below shows the physical capacity savings and power savings from one test conducted for The Green Grid. This was carried out assuming that the amount of useful storage required is fixed, and that the configuration can be changed. Therefore, the number of physical drives is reduced (when using COMs) to store the same amount of data and, because there are fewer drives which can be simultaneously accessed and provide data, the performance also falls. This performance drop can be significantly greater than the capacity increases, for example compression can half the number of drives but causes performance to fall 75%.

However, the testing is carried out on a specific configuration with a fixed physical storage capacity for which the metric is valid. Applying the COM on a fixed configuration would therefore result in an increase in useful capacity with no change in physical capacity resulting in power savings relative to the useful capacity[6].

[6] This is linked to the fact that the performance of a data storage device with a COM is compared to the performance of the equivalent number of physical devices without a COM (e.g. 1 physical HDD+COM = 2 physical HDD)

COM Feature	Physical capacity savings	Idle / active power savings	IOPs (70/30 R-W Workload)
Parity RAID	33%	34% / 30%	0.47 X
Compression	50%	38% / 35%	0.25 X
Thin Provision	50%	38% / 35%	0.50 X
Thin Provision + Deduplication	75%	57% / 68%	0.20X
Snapshot	not tested		

Note: capacity savings = fewer drives

Figure 18 Summary of COM benefits. Source: The Green Grid (2016)[7]

Assuming the active performance is proportional to the number of drives, the capacity and performance can be recalculated assuming the number of drives do not change (Table 4). A COM metric adjustment factor can then be calculated, based on the weighting of active and idle established in the previous sections.

Table 4 COM metric adjustment factor

	utilisation	IOPS	COM metric factor
parity raid	149%	70%	102%
compression	200%	50%	100%
thin provision	200%	100%	141%
thin provision + de-duplication	400%	80%	179%

The metric factor varies depending on the specific COM, and would also vary depending on the exact software, usage and other factors. Perhaps surprisingly, compression and parity raid do not give any overall benefit in this particular situation.

Given the limited information, a COM metric factor of 150% is used for transactional applications and streaming applications. This assumes that at least 2 appropriate COMs are applied in each situation.

For capacity applications, the performance is not a factor and therefore the utilisation is the only consideration rather than the calculated COM metric factor. If the COMs all work maximally when combined, the utilisation could be improved by as much as 800%, by applying three COMs. However, capacity applications frequently involve data that has already been reduced as much as possible through compression and other techniques before storing them (e.g. compressing through use of ZIP and JPG), which means the

[7] The Green Grid (2015) *Preliminary Assessment of Emerald Data* [Accessed online: http://www.snia.org/sites/default/files/emerald/EPA_Storage_Stakeholders_Nov-2015/TGG_Emerald_Analysis_Discussion_v9_110615.pptx]

COMs will not achieve this. The extent of this is not known, and therefore a conservative factor of 300% is applied (i.e., the $COMfactor_{cap}$ is set to 3, as shown in Section 2.5.2.3 below). It should be noted that the accuracy of this COM factor should be improved e.g. via a dedicated testing/simulation activity.

The weighting is applied for three or more COMs to improve the chances that the COMs have a real impact and encourage more availability,. This balances the high utilisation achieved in The Green Grid testing, together with the more limited benefit based on the type of data being streamed.

The lack of data means that confidence in the COM factor is relatively weak. Therefore, the proposed factors are conservative and may underestimate the savings. In addition, the required number of COMs to be eligible for the COM factor is relatively low given that many transactional products already have 2 COMs available (delta snapshots and thin provisioning). This may not reward, or encourage manufacturers to include more COMs in their products, particularly data deduplication, which is effective but relatively complex to implement effectively.

2.5.2.1 Transactional metric

$$transactional\ metric = hot\ band\ perf^{0.5} \times ready\ idle\ perf^{0.5} \times COMfactor_{trans}$$

2.5.2.2 Where $COMfactor_{trans}$= 1.5 for two of more available COMs. Streaming metric

$$streaming\ metric = write\ perf^{0.17} \times read\ perf^{0.33} \times ready\ idle\ perf^{0.5} \times COMfactor_{stream}$$

Where $COMfactor_{stream}$=1.5 for two or more available COMs.

2.5.2.3 Capacity metric

$$capacity\ metric = ready\ idle\ performance \times COMfactor_{cap}$$

Where $COMfactor_{cap}$= 3 for three or more available COMs.

3. **Task 5B Summary of calculations**

This table summarises the steps needed to calculate the energy budget for an individual model/configuration, for example, by the manufacturer, or for market surveillance. It covers the information required that must be available to the market surveillance authority and the formulae applied at each step.

Table 5 Steps needed to calculate the energy budget

<table>
<tr><th>Step</th><th>Description</th><th>Information required</th><th>Calculation</th></tr>
<tr>
<td>I</td>
<td>Product performance based on SNIA Emerald v2.1.1 test results</td>
<td>
The principal application of the product and configuration.
SNIA Emerald test results relevant for that application (capacity, transactional or streaming) based on table below:
<table>
<tr><th></th><th>Hot band test IOPS/W</th><th>Sequential test MiBPS/W</th><th>Ready idle capacity GB/W</th></tr>
<tr><td>Capacity applications</td><td>N</td><td>N</td><td>Y</td></tr>
<tr><td>Transactional applications</td><td>Y</td><td>N</td><td>Y</td></tr>
<tr><td>Streaming applications</td><td>N</td><td>Y</td><td>Y</td></tr>
</table>
Result of COMS tests from SNIA Emerald[8]
</td>
<td>
Product performance is the weighted geometric mean of the test results weighted according to the table below.
<table>
<tr><th></th><th>Hot band workload test
IOPS/W</th><th>Seq read workload test
MiBPS/W</th><th>Seq write workload test
MiBPS/W</th><th>Ready Idle workload test
GB/W</th></tr>
<tr><td>Transactional</td><td>50%</td><td>0%</td><td>0%</td><td>50%</td></tr>
<tr><td>Streaming</td><td>0%</td><td>13%</td><td>37%</td><td>50%</td></tr>
<tr><td>Capacity</td><td>0%</td><td>0%</td><td>0%</td><td>100%</td></tr>
</table>
If COMs are available, multiply result by the COM metric factor in the table below:
<table>
<tr><th></th><th>Minimum Number of COMS</th><th>COM metric
Factor</th></tr>
<tr><td>Transactional</td><td>2</td><td>1.5</td></tr>
<tr><td>Streaming</td><td>2</td><td>1.5</td></tr>
<tr><td>Capacity</td><td>3</td><td>3</td></tr>
</table>
</td>
</tr>
</table>

[8] For market surveillance authorities, SNIA operates a 'recognised tester program' for those who are able to carry out testing .

Step	Description	Information required	Calculation
II	Product energy budget	Product performance from previous step	Inverse of the product performance, Energy budget = (product performance)$^{-1}$
III	*Total energy budget including power and cooling*		*Total energy = PUE x energy budget* *Where PUE = 5/3*
Optional Steps IV and V			
IV	*ASHRAE adjusted total energy budget*	*ASHRAE operating condition. This is provided by the manufacturer.*	*Total energy$_{ASHRAE}$ = ASHRAE factor x total energy* *Where ASHRAE factor =* *1 for ASHRAE 1 or 2* *0.96 for ASHRAE 3 or 4*

Step	*Description*	*Information required*	*Calculation*
V	*Guidance total energy budget*	*Is good guidance provided by the manufacturer with the product or online?* *The case study does not specify how to evaluate the quality of the guidance.*	*If good guidance is provided:* *total energy$_{guidance}$ = Step III total energy x 0.97* *or if option IV is used:* *total energy$_{guidance+ASHRAE}$ = Step IV total energy$_{ASHRAE}$ x 0.97*
Required Final Step VI			
VI	EEI and points	Reference case total energy budget from regulation	EEI = total energy / reference case total energy (where “total energy” can be *total energy$_{ASHRAE}$, total energy$_{guidance}$ or total energy$_{guidance+ASHRAE}$ depending on use of optional steps)* Points = 1- EEI

4. **Task 5C - Scalability analysis**

The scalability analysis looks at how the metric changes across the entire dataset and what factors influence the efficiency based on the proposed metric. The analysis covers:

- Number of drives
- Rotational drive speed
- Capacity
- Classification

4.1.1 Transactional Metric scalability analysis

Analysing firstly the online 3 products gives a better picture of the metric behaviour within a classification. Users are more likely to be comparing devices within a classification at purchase and so this is important.

Table 6 shows that 15k rpm drives have the worst efficiency, suggesting that the increase in IO does not offset the increase in power consumption, and the loss in ready idle performance. This may create a problem because 15k drives provide a useful niche in the market for users that have relatively small capacity requirements but very high performance needs and the metric may negatively impact this niche. However, now that the cost of SSD are comparable in cost to 15k HDD, this is less important and the metric may positively influence the transition to SSD by providing a clear differentiation in efficiency. Lower capacity products also tend to have a lower efficiency but this is more related more to the smaller capacity of the 15k drives. 7.2K and 10k HDD show no particular trends and the efficiency is related to the product design itself which means the metric can be applied and compared across different product types without negatively influencing choice.

Table 6 Efficiency score for OL3 products

Trans Device Rated Speed (RPM)	Trans Total Num Installed Storage Devices Optimal Point	Trans total capacity	Trans Optimal Point Hot Band Workload Test (IOPS/W)	Trans Optimal Point Ready Idle Workload Test (GB/W)	Trans metric 50% active, 50% idle
10000	24	14400	34.41	46.74	40.10
10000	24	14400	34.42	43.3	38.61
7200	12	24000	14.42	76.38	33.19
7200	12	24000	12.91	82.61	32.66
10000	24	28800	28.4	36.6	32.24
15000	24	7200	35.6	18.3	25.52
7200	12	48000	5	122.3	24.73
7200	12	12000	11.4	46.9	23.12
10000	24	14400	35.7	12.7	21.29
15000	24	3504	34.6	12	20.38
10000	24	14400	30.4	13	19.88
10000	24	14400	27.3	12.5	18.47
10000	12	7200	21.7	15.5	18.34
15000	12	7200	14.1	20.3	16.92
7200	24	24000	8.8	29.1	16.00
7200	12	12000	8.7	28.1	15.64
7200		4608	4.88	49.9	15.60
15000	24	14400	13	17.5	15.08
10000		256	4.6	36.49	12.96
7200	12	12000	6.4	23.3	12.21
15000	24	14400	19.3	6.8	11.46
7200	12	12000	5.5	19.4	10.33
15000	24	3504	33.1	2.7	9.45
15000	24	3504	29.1	2.8	9.03
15000	12	7200	14.2	5.3	8.68
15000	24	14400	10.3	7.3	8.67
15000	24	14400	25.7	2.7	8.33
15000	24	3504	20.6	3.1	7.99
15000	24	7200	4.2	11.3	6.89

Figure 19 and Figure 20 show the efficiency plotted against the drive count and capacity for the whole data set. This includes all products, including the top performing 10%. The following observations can be made:

- the metric does not favour any single classification of product which means users should not be pushed towards an unsuitable product
- products within a classification are distributed across the efficiency range and therefore the metric is successful in differentiating products by efficiency
- 15K are consistently the least efficient. These drives serve a particular requirement but more efficient SSD are a suitable replacement
- higher drive count/capacity are generally more efficient, and have a higher minimum efficiency. However, this is unlikely to encourage sale and purchase of over specified capacity requirements because the trends suggest a large (and likely costly) increase

in capacity is required to produce a relatively small increase in efficiency. Other design factors are likely to be more effective ways to improve efficiency

- SSD are significantly more efficient, despite the lower drive count (figure OL2 SSD). This is due to the very high active IO performance. For very high end products, the performance difference is so large it overwhelms all the other factors and based on this metric SSD are twice as efficient as any other products. These high end products are currently extremely expensive and should be considered as BAT. However, as technology improves and costs fall there may be a step change in the apparent efficiency of products, as products provide IO performance beyond what is required. There is a risk that the product could be optimised for the test procedure using a small number of SSDs mixed with HDDs and could achieve this apparent efficiency improvement without creating any real savings. More testing of SSD products would provide a clearer understanding about how future products might behave under this efficiency metric.

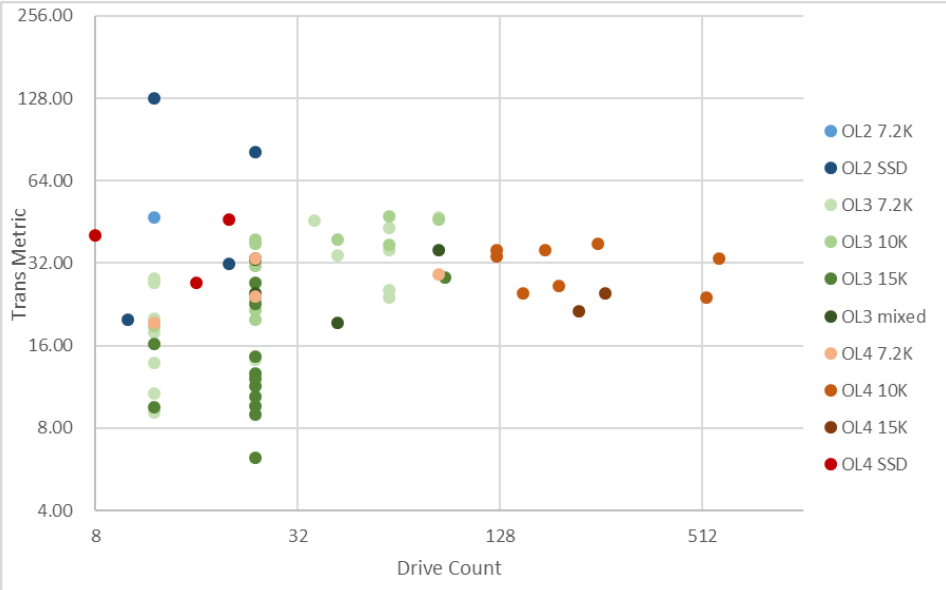

Figure 19 Transactional efficiency and drive count relationship

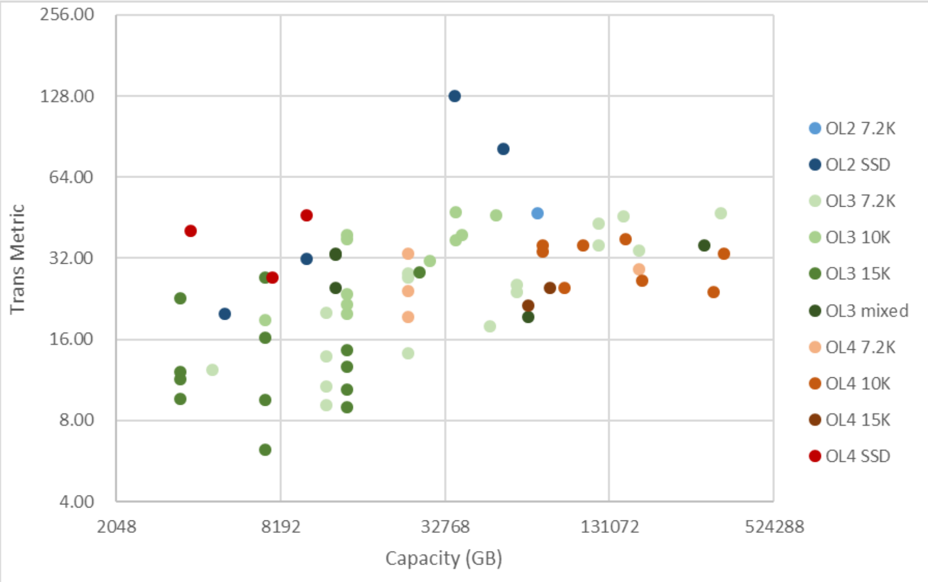

Figure 20 Transactional efficiency and storage capacity relationship

Comparing the metric against the date the product was first placed on the market (Figure 20), there is no general correlation. The only improvement is shown in the efficiency of 15k storage products and this may give a false impression of an appearance of overall improvement. In addition, the three OL2 SSD also provide a very clear line of improvement. Since the number of data points is so low, it is not clear whether this is a real trend, however, SSD technology was relatively immature during this timeframe and improvements are more likely.

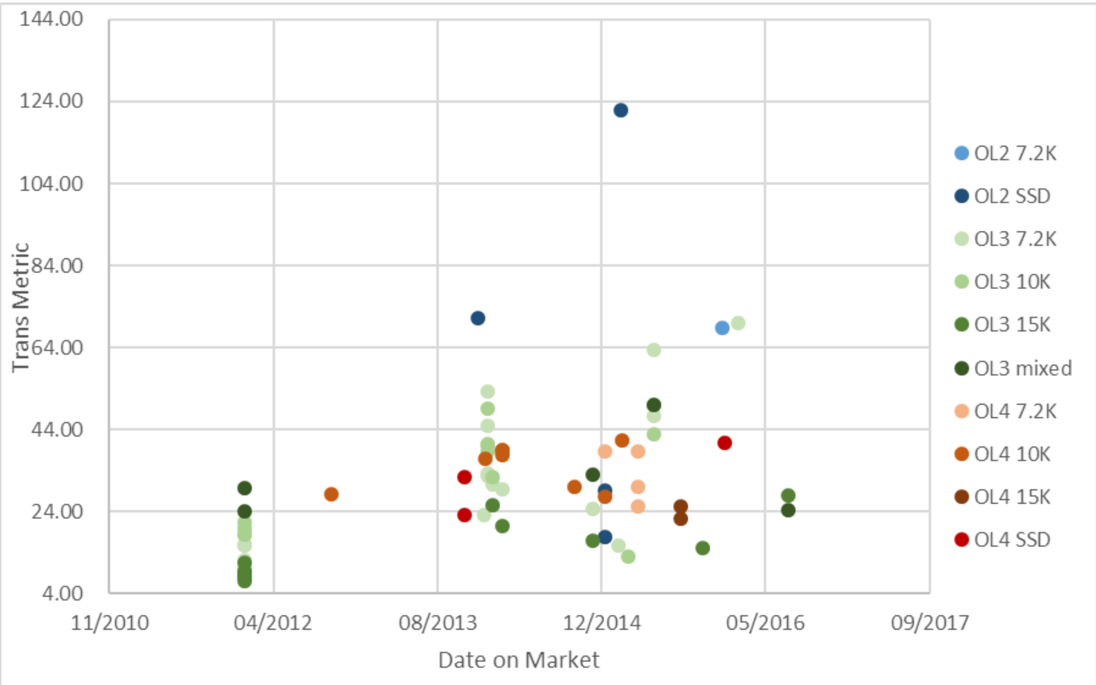

Figure 21 Transactional efficiency against date product available on the market

4.1.2 Transactional metric + COMs

The COMs factor improves the efficiency of products with two or more COMs by a factor of 1.5. This covers 84% of the products and as a result most of the products increase in efficiency. The number of COMs tends to correlate well with the classification of the product and therefore OL2 are the least likely to qualify for the COMs adder.

Figure 22 compared against Figure 20 shows the same overall pattern except with a higher efficiency since it is based on a multiplicative COM metric factor.

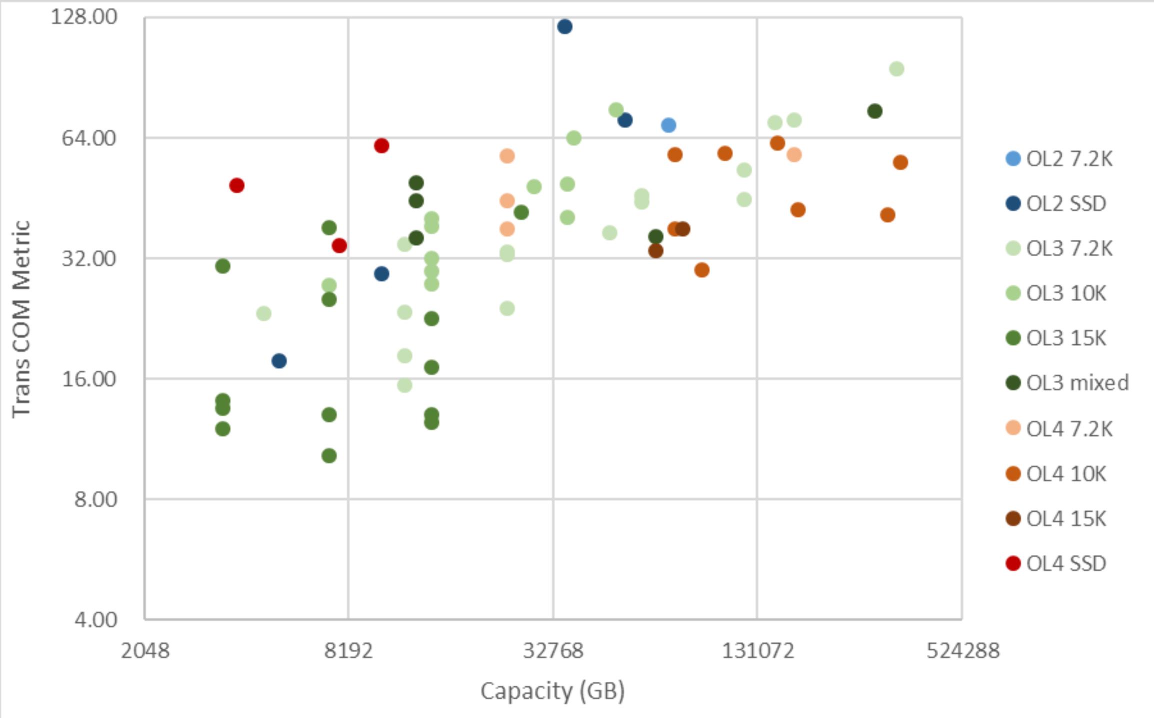

Figure 22 Transactional efficiency with COM factor and storage capacity relationship

Plotting the transactional metric against the transaction + COMS metric (Figure 23) shows that the majority of products have a higher efficiency and that the unqualified products are mostly OL2 and OL3 7.2K.

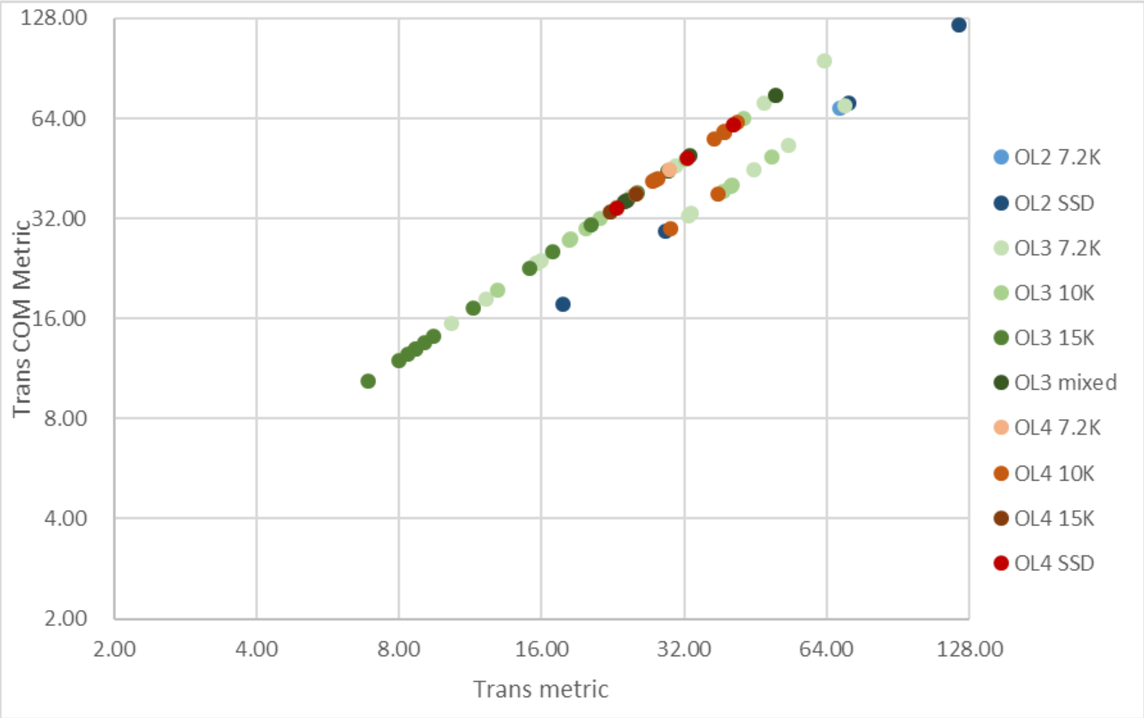

*Figure 23 Transactional efficiency **with** COM factor and transactional efficiency **without** COM factor relationship.*

Increasing the minimum number of qualifying COMs to three would reduce the number substantially and result in mainly OL4 products (in red) qualifying.

4.1.3 Streaming Metric

The metric shows a distribution of efficiencies across all product types and categories. However, there are few clear trends in the streaming metric and the limited dataset makes any trends hard to identify. Rotational speed appears to have no strong impact on the metric. There is also no clear pattern in the metric with respect to the drive count or capacity. Furthermore, the metric does not correlate with the date the product was first placed on the market, which is unexpected. It suggests the metric can be equally valid across all types of streaming products, but may also suggest that the analysis is insufficient to understand the products adequately. However, without additional test data and more detailed test results, such as IO performance this may not be possible. As a result, there is less confidence in the maturity of the metric.

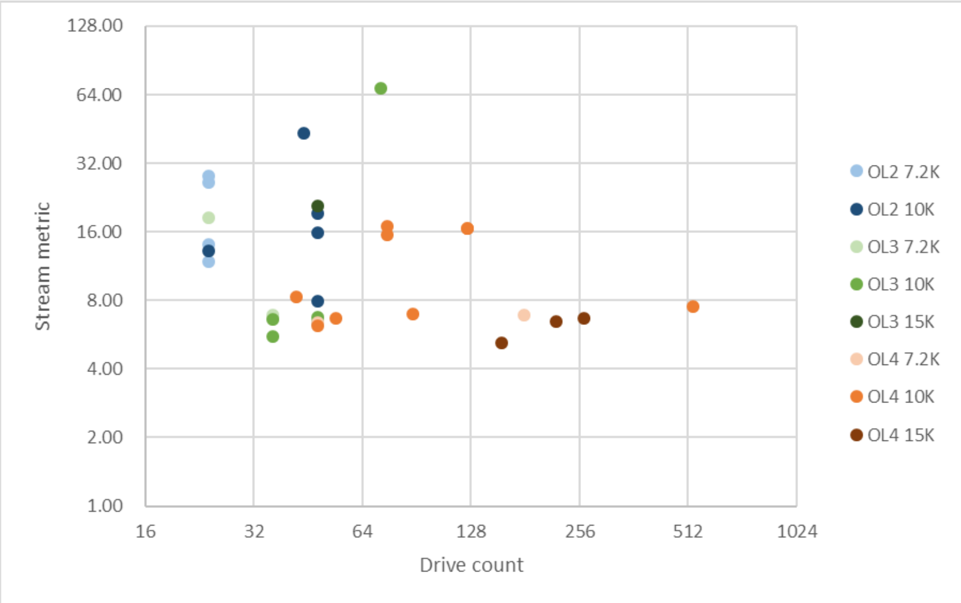

Figure 24 Streaming efficiency and drive count relationship

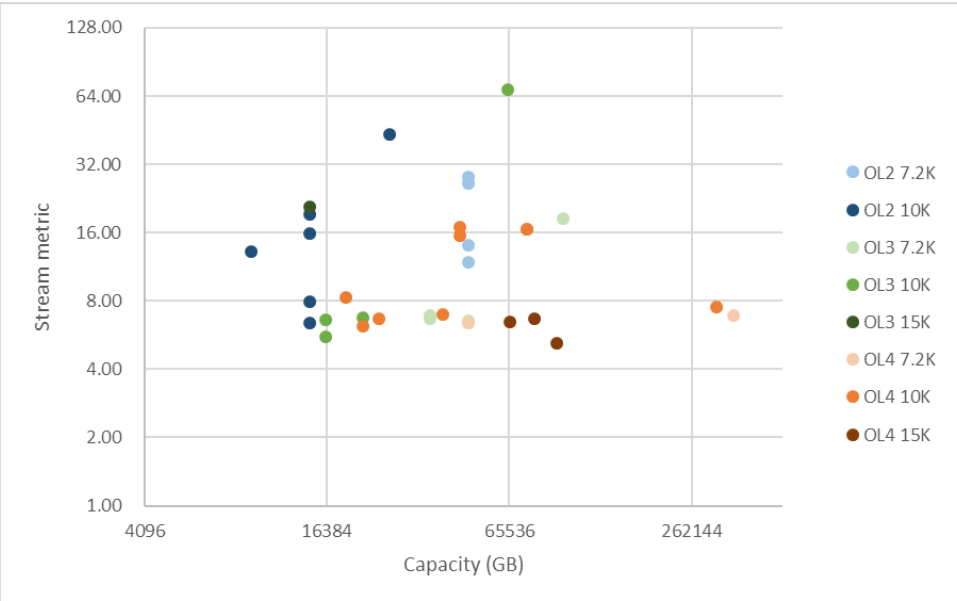

Figure 25 Streaming efficiency and storage capacity relationship

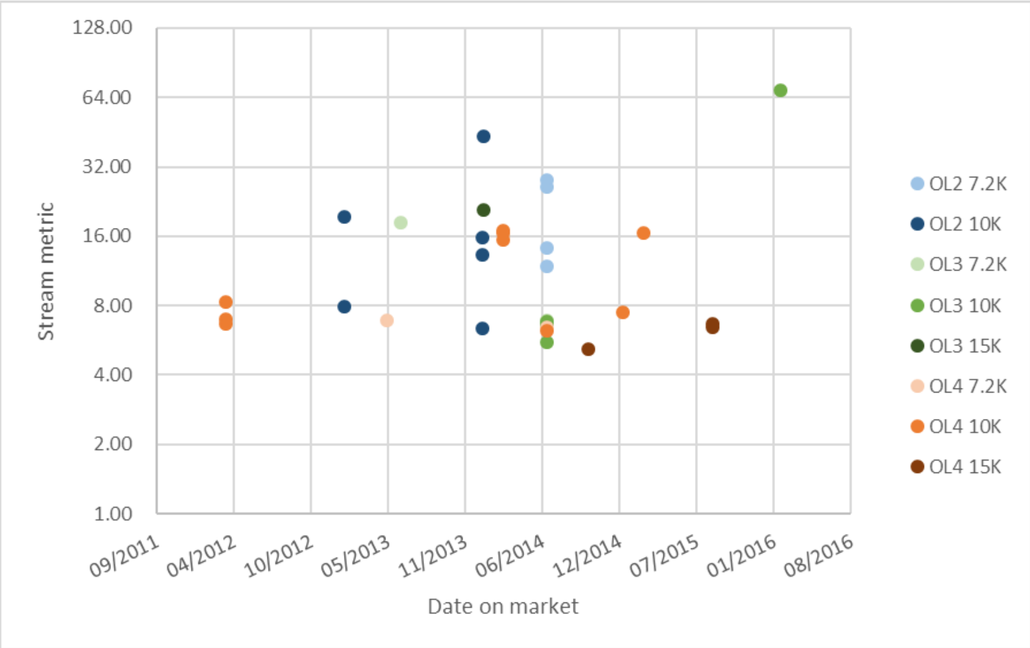

Figure 26 Streaming efficiency against date product placed on the market

The COMs factor improves the efficiency of products with two or more COMs by a factor of 1.5. This covers almost all the OL3 and OL4 products and none of the OL2 products.

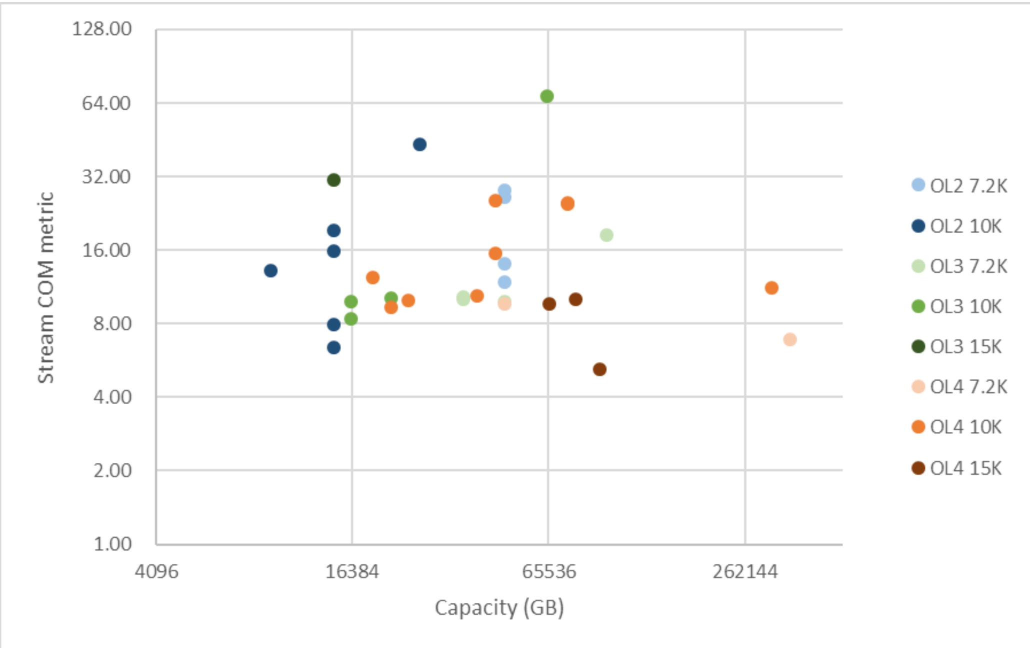

Figure 27 Streaming efficiency with COM factor and storage capacity relationship

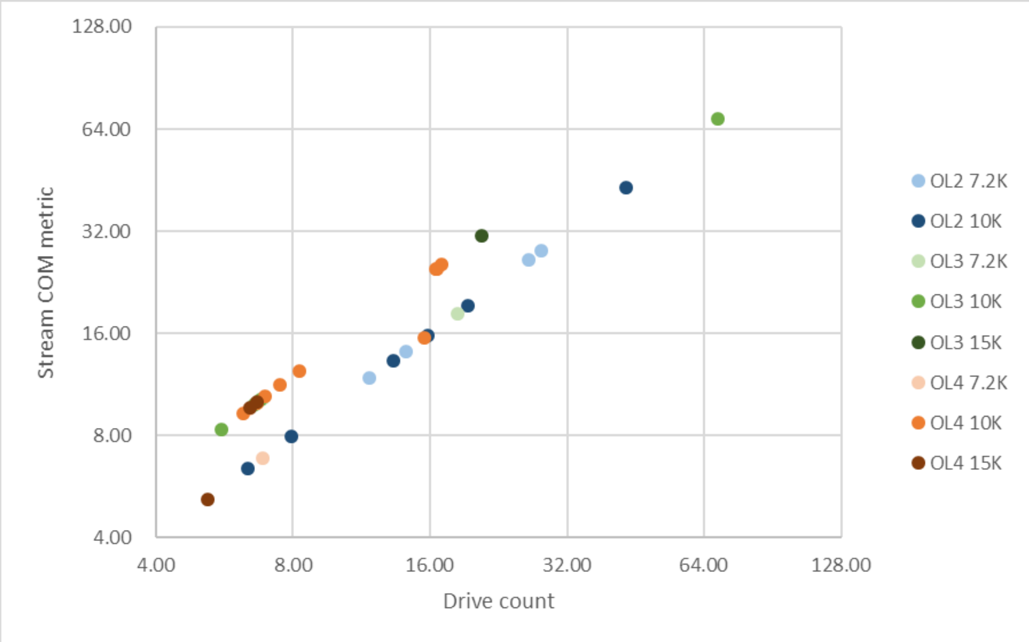

Figure 28 Streaming efficiency with COM factor and drive count relationship

4.1.5 Capacity metric

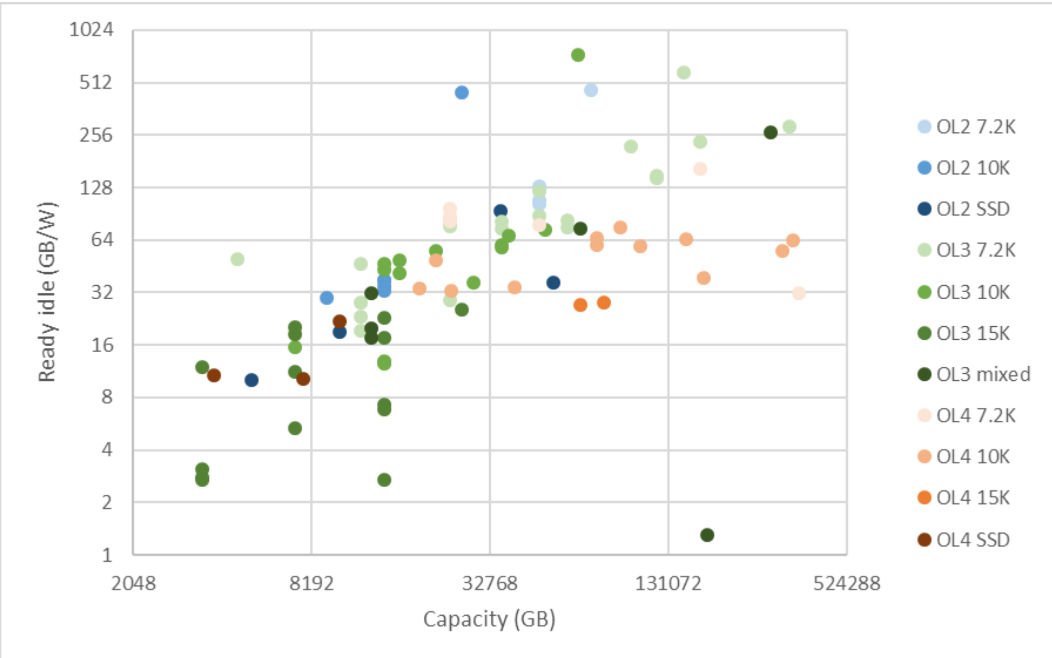

Figure 29 Ready idle efficiency and storage capacity relationship

Because the capacity metric is simply the ready idle, the results are very straightforward. The most efficient products have greater capacity, greater capacity per drive, and slower rotational speed (7.2K). SSDs do not show any significant advantage but this may be due to their relatively low capacity currently. SSD capacity will continue to increase, however, rotational hard drives are also steadily increasing in capacity and are likely to retain an advantage. In addition, the COMs weighting has a bigger impact in the overall metric.

4.1.6 Capacity metric + COMs

The COMs factor improves the efficiency of products with three or more COMs by a factor of 3. This covers 22% of the products, mostly OL4 products (Figure 30).

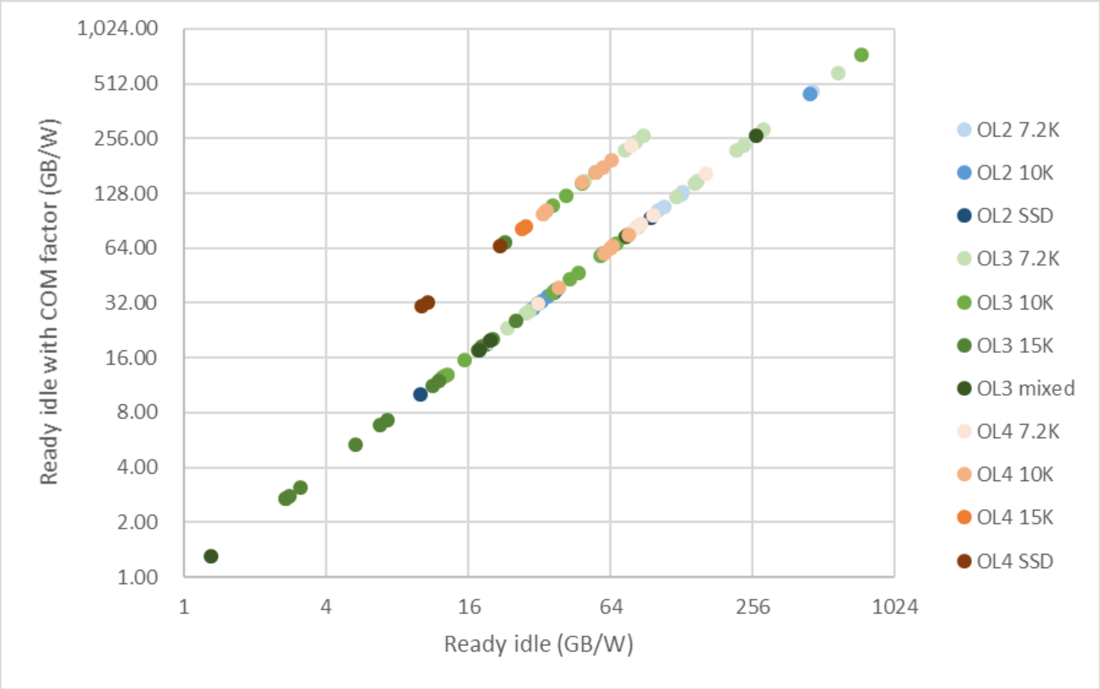

Figure 30 Ready idle efficiency with COM factor against ready idle efficiency without COM factor

5. Conclusions

The analysis shows that the final metrics can differentiate between different products, without favouring any particular class, and are applicable to a range of applications with varying levels of utilisation. However, the confidence levels of the COM weighting and the streaming metric are limited by the lack of data and the limited analysis that can be made as a result.

The even weighting of the active and idle energy performance tests was based on real utilisation data, and balances the sometimes competing requirements of low idle power and high performance. The sensitivity of the ratio is relatively low and suggests that it will also be applicable at other realistic utilisation levels.

The biggest impact on the transactional metric is likely to be for future SSD products, which have very high IO performance compared to current products at similar power consumption levels. The expected growth in SSD storage products means that active energy performance is expected to increase dramatically and the efficiency metric will reflect this. Further in the future, new storage technologies could push this even more. SSD and hard drive capacities are also projected to increase steadily and this will improve ready idle, but not at the rate that the IO can be improved.

COMs also have significant scope for improvement and more data could justify a more granular COM metric adjustment factor, for example, based on how many and which combination of COMs are available rather than a single value.

5.1.1.1 Transactional metric

$$transactional\ metric = hot\ band\ perf^{0.5} \times ready\ idle\ perf^{0.5} \times COMfactor_{trans}$$

Where $COMfactor_{trans}$= 1.5 for two available COMs. 1.8 for three or more COMs

5.1.1.2 Streaming metric

$$streaming\ metric = write\ perf^{0.17} \times read\ perf^{0.33} \times ready\ idle\ perf^{0.5} \times COMfactor_{stream}$$

Where $COMfactor_{stream}$=1.5 for two available COMs. 1.8 for three or more COMs

5.1.1.3 Capacity metric

$$capacity\ metric = ready\ idle\ performance \times COMfactor_{cap}$$

Where $COMfactor_{cap}$= 3 for three or more available COMs

HOW TO OBTAIN EU PUBLICATIONS

Free publications:

- one copy:
  via EU Bookshop (http://bookshop.europa.eu);
- more than one copy or posters/maps:
  from the European Union's representations (http://ec.europa.eu/represent_en.htm);
  from the delegations in non-EU countries
  (http://eeas.europa.eu/delegations/index_en.htm);
  by contacting the Europe Direct service (http://europa.eu/europedirect/index_en.htm)
  or calling 00 800 6 7 8 9 10 11 (freephone number from anywhere in the EU) (*).

  (*) The information given is free, as are most calls (though some operators, phone boxes or hotels may charge you).

Priced publications:

- via EU Bookshop (http://bookshop.europa.eu).

Priced subscriptions:

- via one of the sales agents of the Publications Office of the European Union
  (http://publications.europa.eu/others/agents/index_en.htm).

Catalogue number

Publications Office

doi:10.2873/211016