



DPACK: Workload Concentration (WC)

Tech Briefing on how to leverage the new Workload Concentration features

2016

Getting Data Insight

As DPACK progresses into its next generation offerings the goal will be to bring rapid visibility to areas of interest or concern within data captures.

DPACK itself is a Big Data and IoT project wrapped up in one as the collectors are sending back hundreds of thousands of data points to the Online Analytics Engine. Without the processing capabilities of this platform, these data points would be nearly impossible to surface.

Introducing Workload Concentration

WC is a representation of the facts about server data. This is displayed in a way that allows you to quickly see what drives [LUNs/Volumes] are creating the most activity in the environment as well as their ratio to overall data capacity.

Before going deeper into the data that WC shows, we need to understand a couple of key points so there is no misunderstanding as to what you are viewing.

Q: Does this predict tiering or sizing requirements?

A: No, this is simply a representation of the facts. It's unrelated to any platforms specific tiering or caching algorithms. However, the data does have meaningful correlation to any device that stores capacity.

Q: Is this the same as Heat Mapping?

A: No, heat mapping would indicate sub-LUN or sub-File level knowledge of how "hot or cold" is a portion of the data. WC only deals with the disks themselves and their relationship to all other disks captured.

The Alternate Views

WC can be viewed in two different methods: Variable Step Graph or Bubble Chart. Each show similar but slightly different information and depending on the environment captured one might work better than the other.

This document will walk through both variations and explain what is being shown, as well as what you should be interpreting.

Note: *These graphs will not currently print in any PDF.*

"Workload Concentration is a stack ranking of disk activity showing performance over capacity."

It is not a sub-LUN or sub-File heat mapping."

In This Briefing

- WC—New ways to see data
- What WC is not
- What should you be interpreting from these new graphical views

WC: Variable Width Step Graph

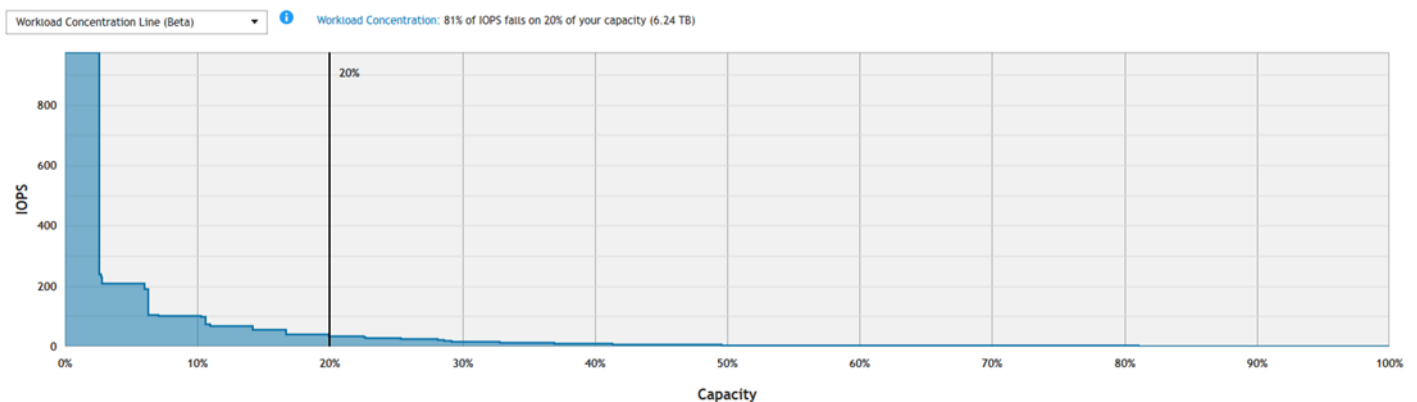
As DPACK collects the current environment it will understand the capacity and performance of each disk. In this document a disk will include any LUN, Volume, DataStore, or Internal Capacity that donates to the overall capacity figure determined by DPACK in the project summary.

WC is simply a plotting of the relationship of those disks' performance (IOPS) over capacity. With the Variable Width Step Graph, each disk is first sorted by its 95th percentile IOPS number from highest to lowest. Each of those disks is then plotted in a graph, much like the sample below, with the highest-performing disks on the far left and the lowest-performing disks on the far right.

Therefore, the height of each "Step" is representing the amount of IOPS in relation to all other steps.

Used Capacity of that disk is represented by the width of the step. The wider the step, the larger in capacity the disk.

The horizontal axis demonstrates all capacity from the environment. There is an arbitrary 20% capacity line drawn in the graph for the purpose of creating a calculation of what amount of IOPS fall into what amount of capacity.



A Workload Concentration Example that shows a very high percentage of IOPS coming from less than 20% of overall capacity

WC's Relationship to Tiering Value

Although WC is in no way a measure of how to size a tier in storage, WC is a great predictor of the environment's candidacy for storage tiering in general. The graph shows a propensity model that demonstrates how hot certain disks are in relation to the rest of the environment.

The higher the concentration of IOPS from a smaller portion of data capacity, the more suited the environment is to data tiering. Likewise, the flatter or more evenly-distributed this graph, the more suited the environment would be towards a monolithic or single-tier design, as performance is evenly distributed throughout capacity.

Since this graph has no knowledge of Sub-LUN level data tiering algorithms that would tier the LUN over multiple disk types, one could assume that a highly-concentrated environment would be even further optimized by this technology. One could also make the argument that even flat graphs could benefit from Sub-LUN level tiering or caching algorithms.

Why Does this Graph Work

Most applications will have disk layouts designed to accommodate the specific workload by the administrator. For example, many higher-IOPS disks are associated with a database or log files which might require much smaller capacity than other workloads that have much lower IO to capacity ratios, such as a File Server.

This view will give you a visual of how the application owners have provisioned storage in relationship to workloads.

WC: Bubble Graph

The WC bubble graph might be more appropriately named the “Workload Characterization” graph, due to its representation of more than just performance over capacity. The Bubble graph shares the same data represented in the Variable Width Step Graph, but leverages a Bubble chart which can show a 4th attribute. In this case that attribute is *Quality of Service* or simply the average latency of the disk’s IOPS.

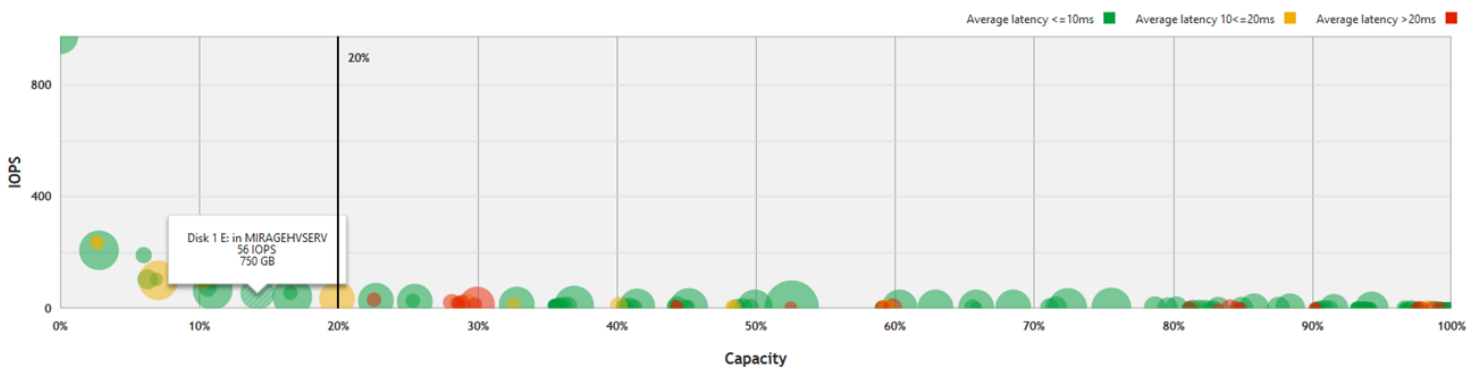
The WC Bubble Graph is simply a plotting of the relationship of all disks’ performance (IOPS) over capacity and their propensity to be a problem factor in the environment. Each disk is first sorted by its 95th percentile IOPS number from highest to lowest. Each of those disks is then plotted in a graph, much like the sample below, with the highest-performing disks in the far upper left and the lowest-performing disks in the far bottom right.

Therefore, the height of each bubble is representing the amount of IOPS in relation to all other bubbles.

Capacity of that disk is represented by the size of the bubble. The larger the bubble, the higher in capacity the disk.

The horizontal axis demonstrates all capacity from the environment.

The additional attribute of Quality of Service is introduced by the color indicator assigned to each bubble. Green is a healthy average latency, yellow represents a potential area of concern, and red indicates a disk that potentially needs to be inspected. By no means are these colors a predictor of a problem. They are just a visual indicator of the facts.



A Workload Concentration Bubble Example that shows many disks with varying average response times

Researching a Potential Problem

Both the Step Graph and the Bubble Chart will become more insightful when the project has a higher number of disks.

The WC Bubble chart can be particularly problematic if the environment has over a hundred disks or has little variation in the performance of each bubble. As seen in the example above, when there are many similar performing bubbles, the graph can produce a “snake” like effect which makes it hard to isolate a bubble. To gain visibility to dense regions of bubbles the chart can be zoomed with the graph controls below the chart.

Each bubble can be “hovered over” to see the disk name, performance, capacity, and the server that owns the disk. The bubbles are also interactive. If you click on any bubble the left hand tree navigation will automatically expand and locate the position of the disk in the tree menu. Once the disk of interest has been located you can use the other available graph data to investigate other performance attributes.

Bubbles marked with a gray color represent a disk that does less than 100 IOPS in either reads or writes. Low-IO disks can skew latency averages and these disks can be ignored. You can read more on the DPACK support site:

<https://dpacksupport.dell.com/entries/78846377>

Understanding more about DPACK

The DPACK team is happy to conduct team or territory level training with employees, partners, or End Users.

Our contact information can be found on the right in the blue column!

The DPACK site is located at <https://DPACK2.Dell.com>

The DPACK support site is located at <https://DPACKSupport.dell.com> or by emailing support@dpack.zendesk.com

The DPACK support site also has a vast library of other insightful attributes of DPACK and can be located here:

<https://dpacksupport.dell.com/forums>

Contact Us

The DPACK team can be contacted at the following address:

[DPACKsupport.dell.com](https://dpacksupport.dell.com)

General Manager:
Sam Kirchoff

NA:
Scott DesBles
Mike Bachman

EMEA:
Uwe Wiest

LatAM:
Alan Rabinovich

APJ:
Mike Bachman

DPACK: Know Your Load!