



VMware vSphere DRS Dump Insight User Guide

June 4, 2021



VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com

Copyright © 2021 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at <http://www.vmware.com/go/patents>. VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

Table of Contents

Introduction	3
Collecting and Uploading drmdump Files	4
About the drmdump files.....	4
DRS 1.0: vSphere versions 6.0–6.7.....	4
DRS 2.0: vSphere version 7.0 and later.....	4
Locating drmdump files and uploading them.....	4
Working with the vMotion Timeline in DRS 1.0 for vSphere 6.0-6.7	5
Understanding the Summary.....	9
Summary.....	9
Audit.....	9
List of DRS recommended moves.....	10
Cluster resource distribution/host usage	10
Host load standard deviation.....	10
Hosts resource usage.....	10
What-if analysis.....	11
Exporting results of a DRS 1.0 dump	13
Working with the vMotion Timeline for DRS 2.0 in vSphere 7.0 and 7.0 U1	14
Some things to know about DRS 1.0 vs 2.0.....	14
The vMotion Timeline of DRS Dump Insight for DRS 2.0.....	15
Understanding the dots.....	16
What are the green dots?	16
What are the blue dots?.....	16
What are the yellow/red dots?	18
Summaries of partial and full dump analyses.....	18
Partial dump analysis summary	19
Full dump analysis summary	19
Exporting results of a DRS 2.0 dump.....	20
Conclusion	20

Introduction

The VMware vSphere® Distributed Resource Scheduler™ (DRS) dynamically allocates and balances computing capacity across a collection of hardware resources aggregated into logical resource pools. DRS continuously monitors use across resource pools and allocates available resources among the virtual machines based on predefined rules that reflect business needs and changing priorities. When a virtual machine experiences an increased load, VMware DRS automatically allocates additional resources by redistributing virtual machines among the physical servers in the resource pool.

NOTE: This user guide assumes you have already [set up a DRS cluster](#) and want to use the [DRS Dump Insight web app](#) fling.

The DRS Dump Insight web app takes existing *drmdump* log files (each *drmdump* log file contains the actions proposed and taken by DRS within a certain time or file size) and provides a summary of the DRS run, along with changes in host resource consumption in a plain text format.

With this tool, you can answer questions like:

- What recommendations did DRS drop due to cost/benefit analysis?
- Why did DRS make a certain recommendation?
- Why is DRS not making any recommendations to balance my cluster?
- How is my custom affinity/anti-affinity rule affecting load balancing in my cluster?
- Can I get all the recommendations made by DRS?

How does this tool work? DRS Dump Insight:

1. Automates dump replay (with pre-built custom DRS replayers)
2. Uses custom replayers to give additional info not available in regular replayer logs (like recommendations dropped due to rule conflicts)
3. Parses and analyzes a custom replayed log to get insight into DRS decisions
4. Generates a final summary in a plain text format

Where is the web app for this tool?

The web app is at <https://www.drsdumpinsight.vmware.com/>

This is a VMware fling, which was announced at <https://flings.vmware.com/drs-dump-insight>.

Collecting and Uploading drmdump Files

About the drmdump files

DRS 1.0: vSphere versions 6.0–6.7

Each drmdump file contains the actions proposed and taken by DRS for a given 5-minute time interval. (DRS runs once every 5 minutes.)

These files are named with the suffix of `*.proposeNonResAction.dump`. An example of a full drsdump filename is `6928191775-proposeNonResActions.dump`.

DRS 2.0: vSphere version 7.0 and later

Each drmdump file contains actions proposed and taken by DRS for a variable time interval limited on a file size of 21 MB.

These files are named with the prefix and suffix of `P1mt-*.log`. An example of a full drmdump filename is `P1mt-67.log`.

Locating drmdump files and uploading them

To find and upload the drmdump files:

1. Find drmdump files in the vCenter support bundle. To learn how to generate and download the vCenter support bundle, see <https://kb.vmware.com/kb/1011641>.
2. Look for the drmdump files in the bundle. Alternatively, you can get the drmdump files directly from your vCenter Server, see <https://kb.vmware.com/kb/1021804>.
3. Unzip the drmdump files for the cluster of interest. Unzipped drmdump files to be analyzed are collected into a directory.

Important: The drmdump files collected should be from the same cluster (domain).

4. Upload the folder in the web portal and select the corresponding vCenter version.

Important: Do not rename the drmdump files or modify the file content; doing so might result in errors.

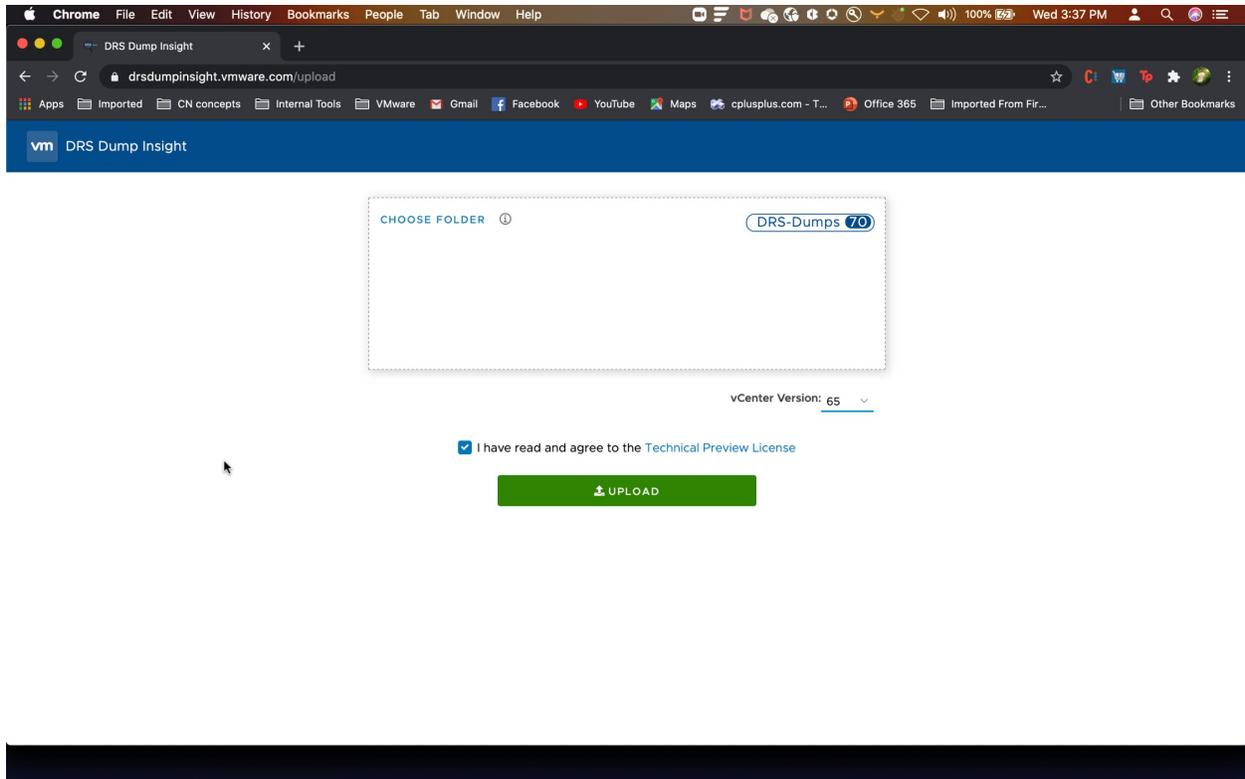


Figure 1. When you go to the DRS Dump Insight web portal, you'll see an **Upload** button. To make it active, choose the vCenter version you're using, then choose the folder where you extracted your DRS dump files from steps 1–3, above.

Working with the vMotion Timeline in DRS 1.0 for vSphere 6.0-6.7

The *vMotion Timeline* is where you view DRS-recommended VM migrations. Each point on the timeline corresponds to a single drmdump analysis (DRS run analysis). The timeline appears after you have uploaded and analyzed your DRS dump files.

NOTE: If you're using vSphere 7.0 or later, go to the instructions at "[Working with the vMotion Timeline for DRS 2.0 in vSphere 7.0 and 7.0 U1.](#)"

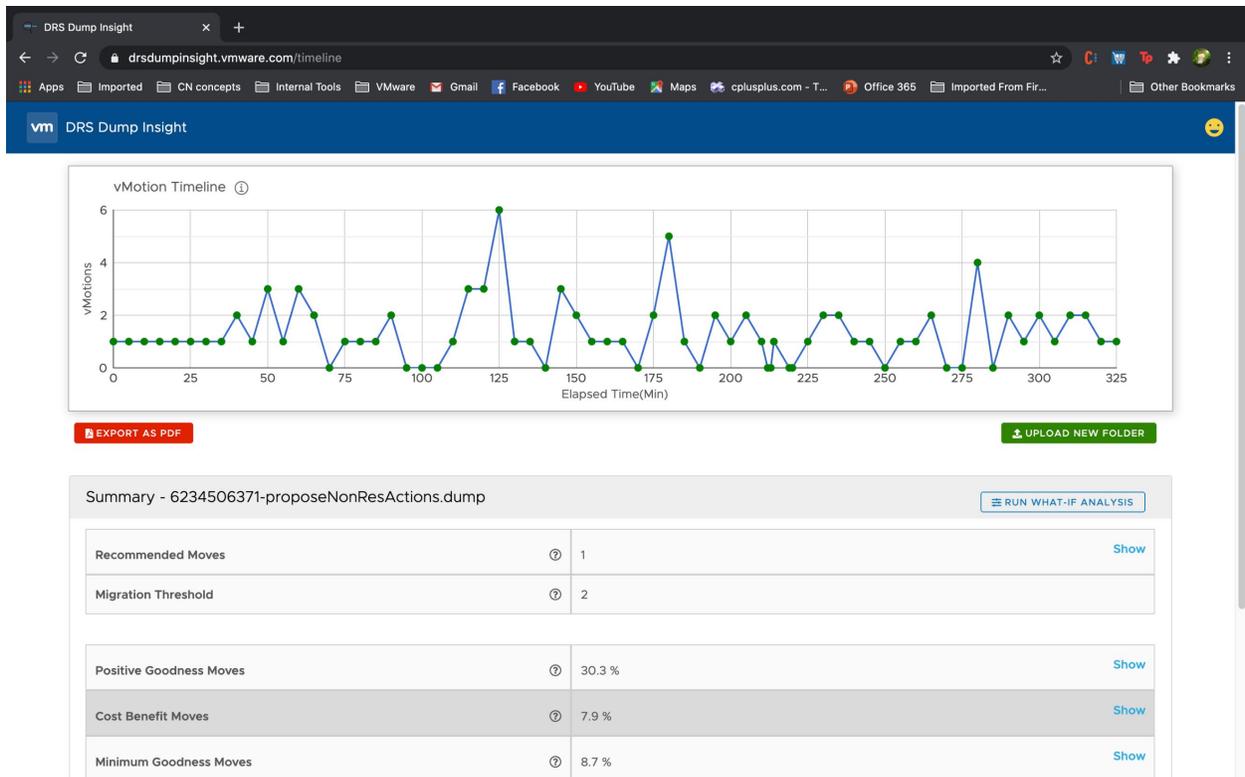


Figure 2. The vMotion Timeline shows you when and how many virtual machine migrations took place.

To display information about the migration, hover over the points as shown in figure 3.

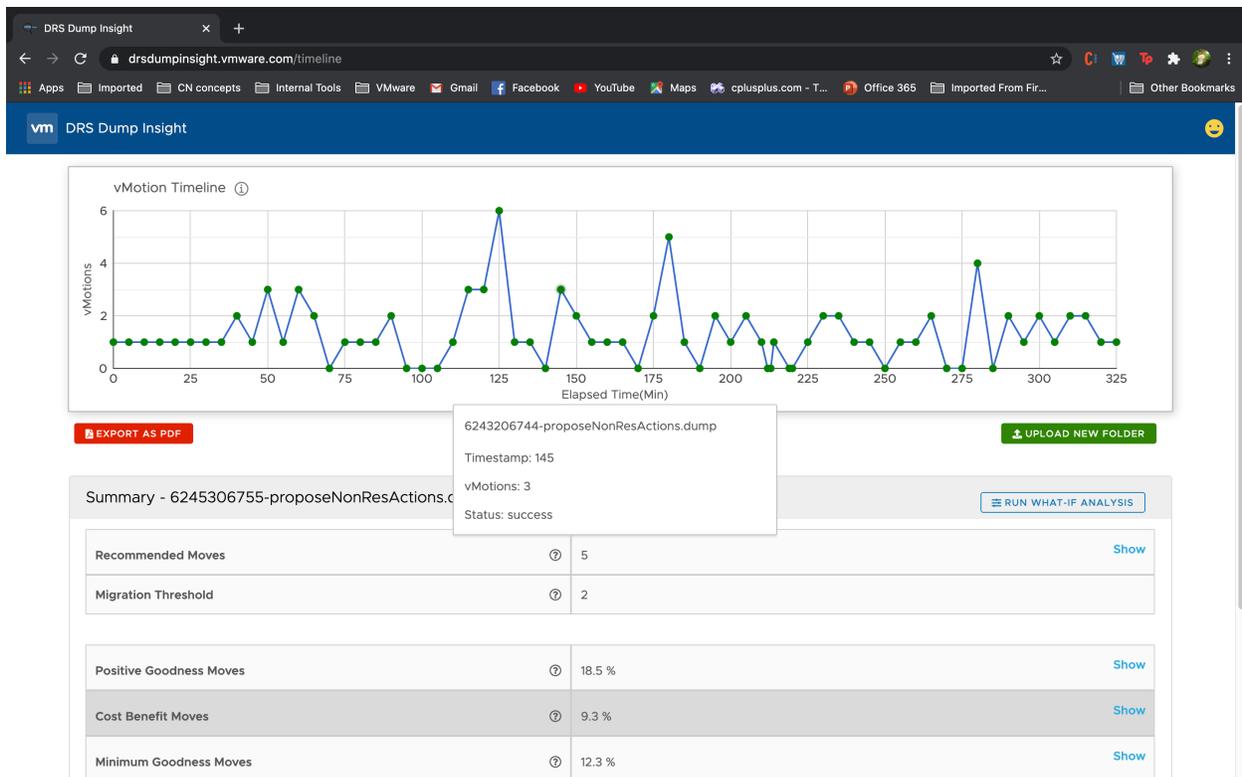


Figure 3. Hovering over a point displays some information about the migration including the dump file, timestamp, number of vMotions, and status.

To access timeline navigation options and the legend, click the info symbol:  as shown in figure 4.

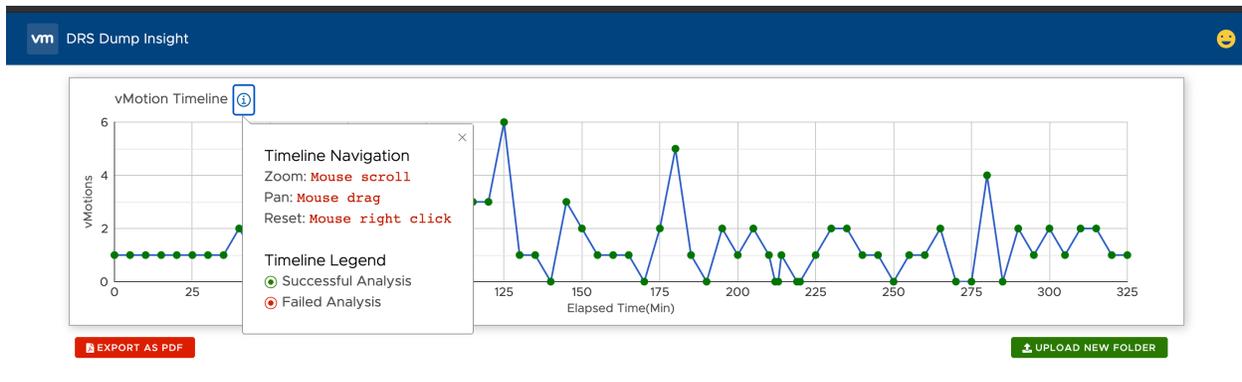


Figure 4. View navigation hints and the timeline legend by clicking the info button next to the vMotion Timeline label.

To see the summary of a DRS run with links to more detailed information, click a point on the vMotion Timeline. You'll see a table beneath the vMotion Timeline with data, which includes:

- The name of the DRS dump file next to the Summary label.
- The number of recommended moves.
- The migration threshold, which specifies how aggressively DRS recommends vMotions. It can be varied between 1-5 with 1 being less frequent vMotions and 5 being very frequent vMotions.
- The positive “goodness” of the moves (%). For a description of goodness, see <https://blogs.vmware.com/vsphere/2020/05/vsphere-7-a-closer-look-at-the-vm-drs-score.html>.
- The cost benefit of the moves (%). See the following section “[Audit](#)” for a description of cost-benefit analysis.
- The minimum “goodness” of the moves (%).
- The negative “goodness” of the moves (%).
- The standard deviation of the host’s resource usage broken down by CPU and memory, before and after the migration.

Click the **Show** label (located to the far right of the summary table) to see more details about each item in the table. (See figure 5.)

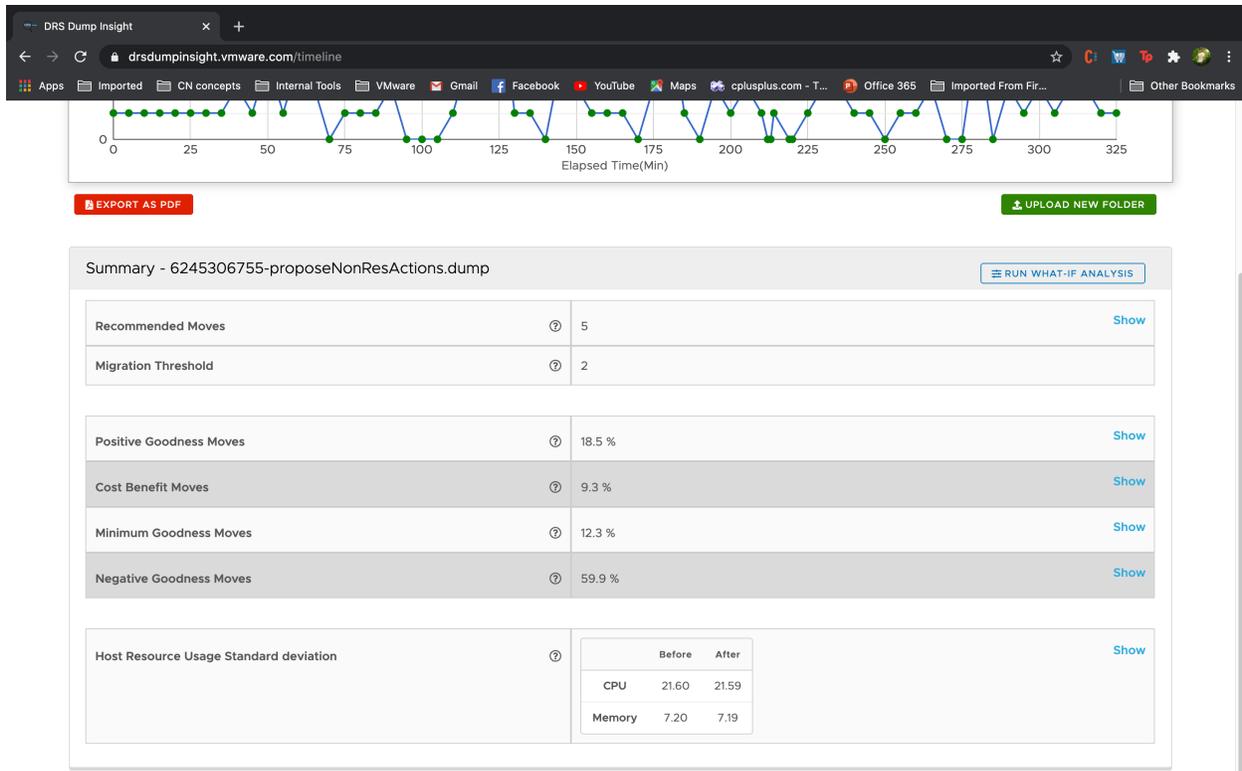


Figure 5. Click on a point in the vMotion Timeline to see a summary of that DRS dump file.

Understanding the Summary

The typical summary output from DRS dump insight has the following main sections.

Summary

The Summary section shows the overall summary with the total vMotions (or moves) DRS recommended for the given run. When DRS has recommended no moves for a given run, DRS Dump Insight displays possible reasons for this.

Audit

During each run, the DRS algorithm runs several iterations, going over the combination of hosts and VMs in the cluster, to come up with the best moves that can improve the load balance in the cluster. This section lists the total moves that DRS considered to provide the final list of recommendations.

The section also lists the number of moves out of the total considered that can improve the load balance and those that cannot improve the load balance.

Some moves can improve the load balance slightly, but the cost of vMotion is much higher, owing to the host and network resources needed for the vMotion (for example, if it's a VM with an active, high memory footprint, and the resulting balance improvement upon vMotion is minimal). DRS drops such moves in what is known as *cost-benefit analysis*. This section lists the number of moves that were dropped due to DRS' cost-benefit analysis.

DRS also drops moves that have a minimal impact on the load balance, to ensure that VMs do not suffer due to unnecessary vMotions. This section lists the number of such moves dropped during the given run of DRS.

List of DRS recommended moves

In this section, the final list of DRS recommended moves is listed with details for each move, which includes information like the target VM, the source and destination hosts, and the reason for the move.

Cluster resource distribution/host usage

DRS dump insight shows the following information with respect to the CPU and memory utilization in the cluster.

Host load standard deviation

This is the standard deviation of CPU and memory resource utilization values across all the hosts in the cluster. The lower this value, the better is the load distribution in the cluster. DRS Dump Insight shows the value for the host load standard deviation before and after the given DRS run. This gives an idea of how the cluster load balance has changed for the given DRS run.

Hosts resource usage

For every host in the cluster, resource utilization is shown before and after the given DRS run. CPU is shown as the percentage used across all cores, and memory is shown as the overall percentage used. This gives a clear picture of how the resource utilization has changed per host during the given DRS run.

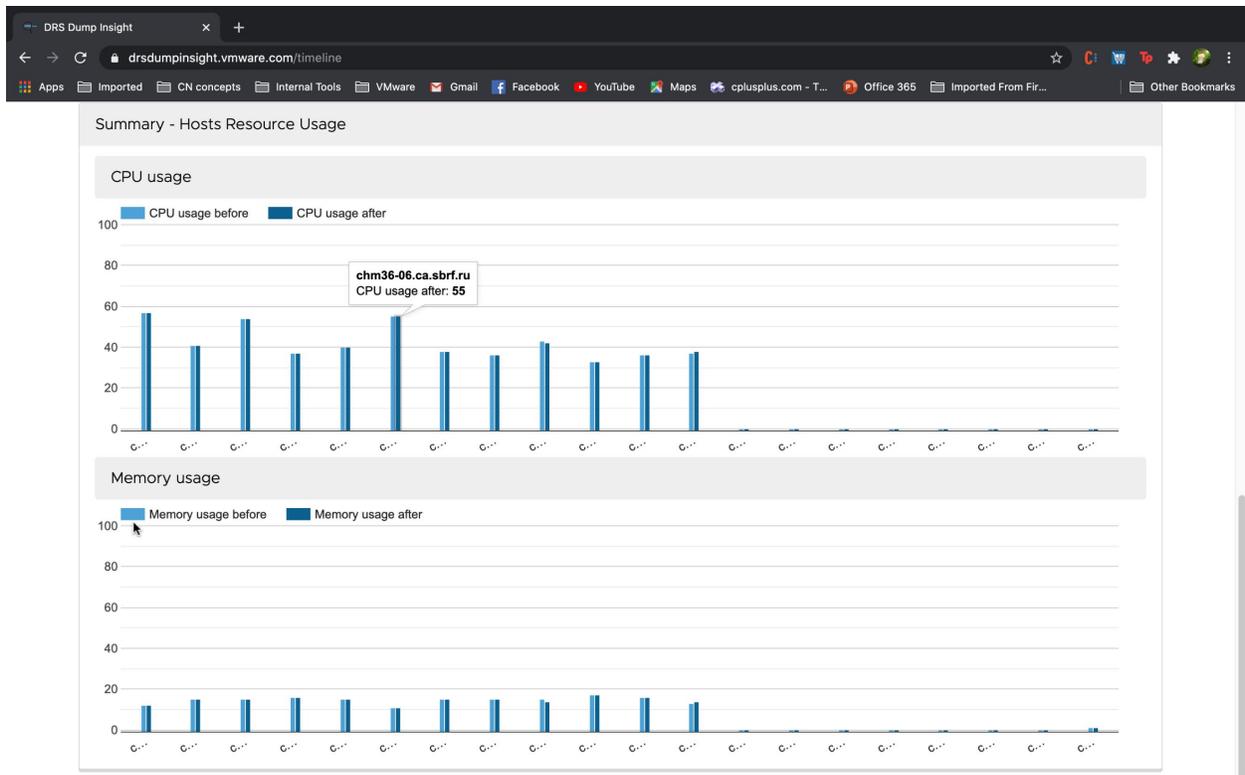


Figure 6. For every host in the cluster, resource usage is shown before and after the given DRS run.

What-if analysis

DRS Dump Insight lets you perform a *what-if* analysis on the DRS run. You can select from a given set of options to see how DRS would have come up with a different set of recommendations with those options set.

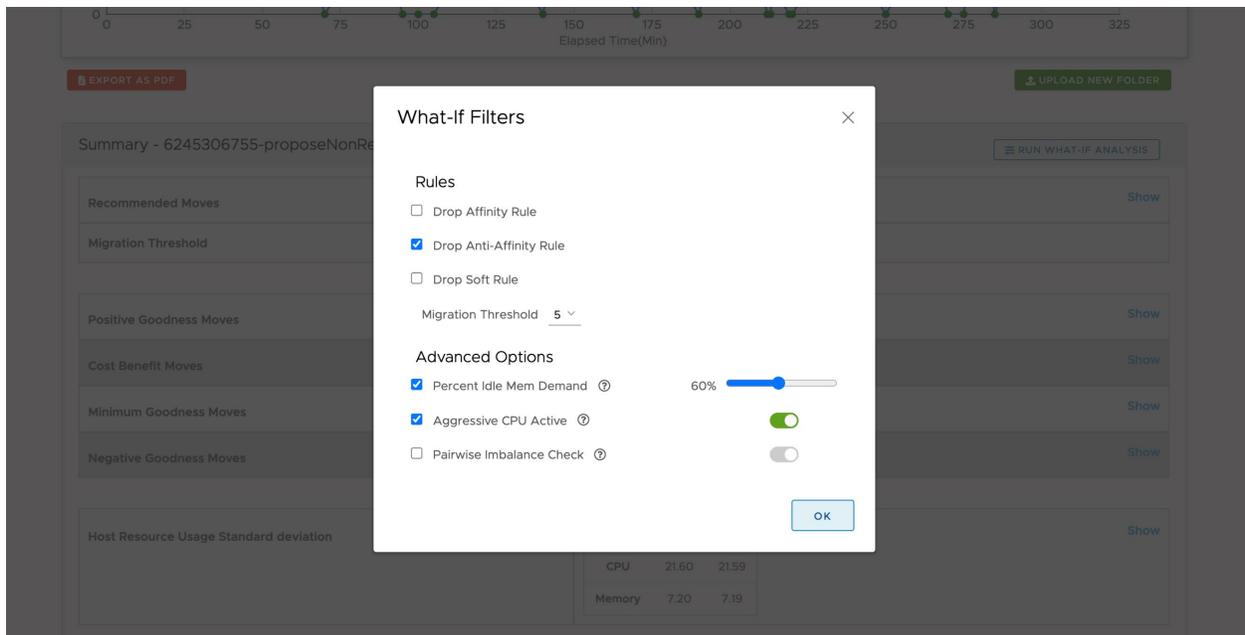


Figure 7. Run a what-if analysis to determine what might have happened when different rules and advanced options were set for your DRS run.

To start a what-if analysis, click the **Run what-if analysis** button on the summary table of a DRS run.

A what-if analysis lets you drop cluster rules and set advanced options:

- **Drop Soft Rule** drops VM-host soft rules.
- **Drop Affinity Rule** drops VM-VM affinity rules.
- **Drop Anti-Affinity Rule** drops VM-VM anti-affinity rules.

For more information about DRS affinity, see “Using DRS Affinity Rules” at <https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.resmgmt.doc/GUID-FF28F29C-8B67-4EFF-A2EF-63B3537E6934.html>.

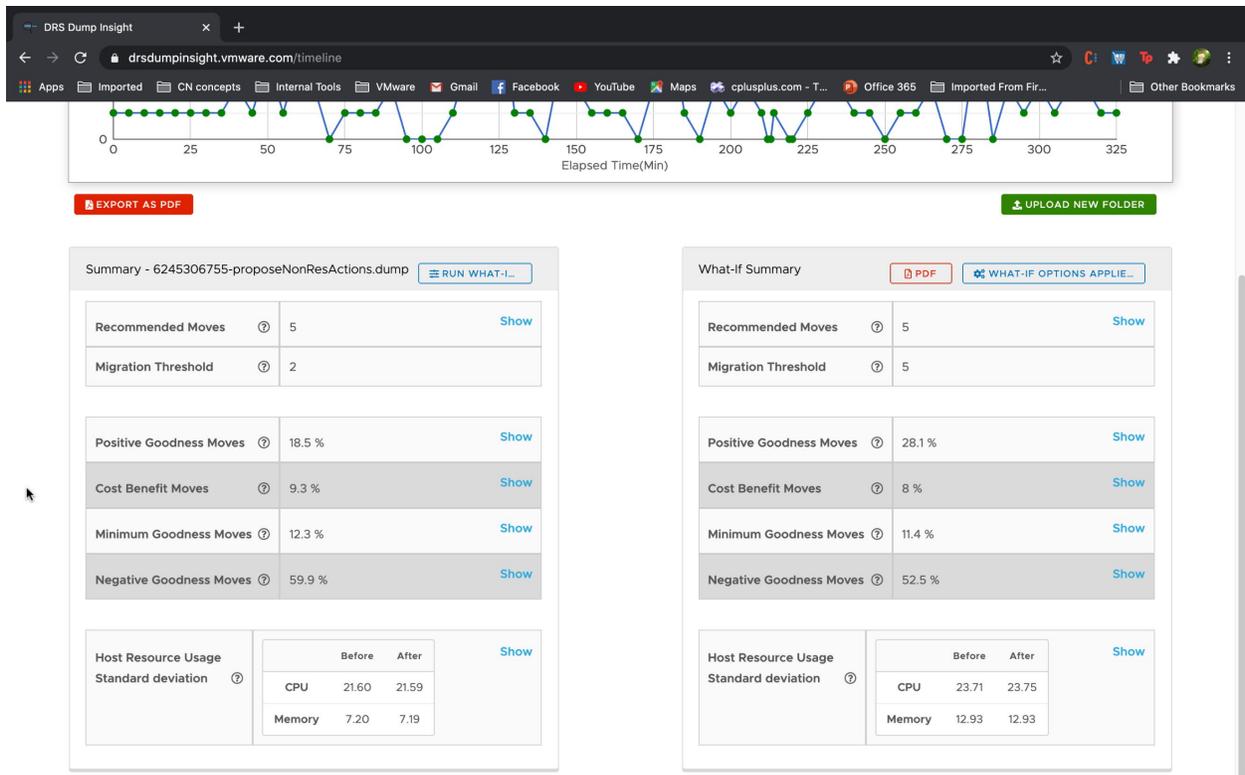


Figure 8. After the what-if analysis completes, you can see a side-by-side comparison of the details of the run with and without the applied what-if options.

Exporting results of a DRS 1.0 dump

DRS Dump Insight lets you export the DRS run analysis in PDF format. Click the **Export As PDF** button below the vMotion timeline to export the results.

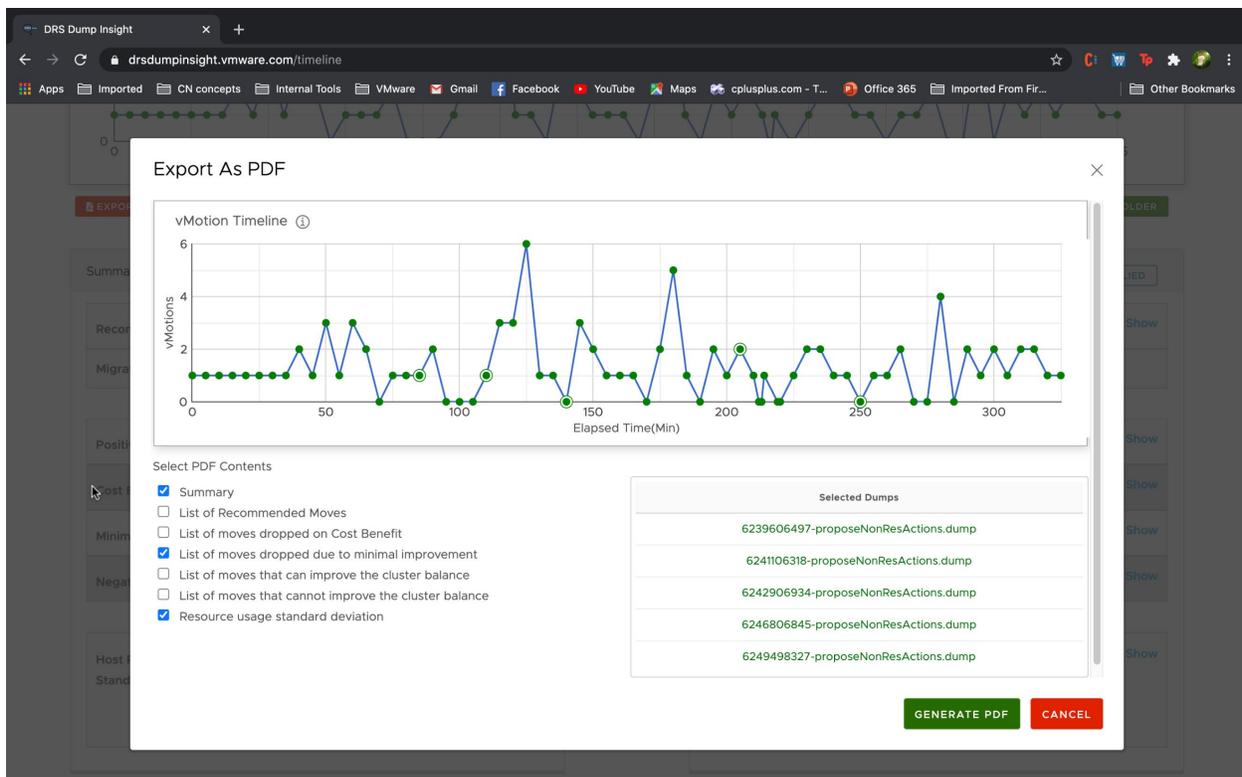


Figure 9. In the **Export As PDF** dialog box, you can choose one or multiple DRS run analyses to export. The results are exported in a PDF file for each DRS run.

Working with the vMotion Timeline for DRS 2.0 in vSphere 7.0 and 7.0 U1

The DRS in vSphere 6.0-6.7 is known as version 1.0. Likewise, the DRS in vSphere 7.0 and 7.0 U1 is known as DRS 2.0. This section describes DRS Dump Insight for DRS 2.0.

Some things to know about DRS 1.0 vs 2.0

With the introduction of DRS 2.0 in vSphere 7.0, the core DRS algorithm and drmdump format have changed significantly. This means there are significant differences while using DRS Dump Insight and understanding the drmdump analysis compared to the prior versions.

Some of the key differences between DRS 1.0 dumps and DRS 2.0 dumps include:

- Each DRS 1.0 dump corresponds to a single DRS run, no such correspondence exists for DRS 2.0 dumps.
- A DRS 2.0 dump contains two types of dump-sections:

- **Full dump (aka cluster dump):** Dumped periodically and the dump contains almost all the information required to reconstruct the cluster state.
- **Partial dump:** Dumped whenever a vMotion is initiated by DRS and there are power-on failures.

The vMotion Timeline of DRS Dump Insight for DRS 2.0

After DRS 2.0 dumps are uploaded for analysis, a vMotion timeline is generated.

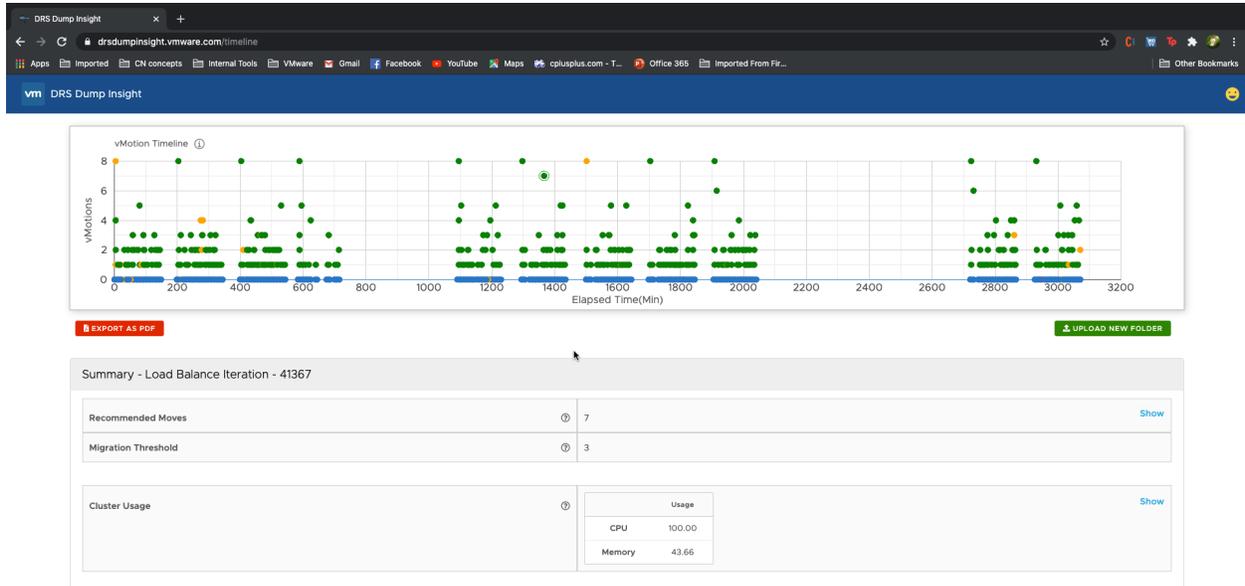


Figure 10. DRS Dump Insight for DRS 2.0 (vSphere 7.0 and 7.0 U1) displays the vMotion timeline when dump files are uploaded for analysis.

The immediate thing to notice about the vMotion timeline for DRS 2.0 dumps is the distribution of points on the timeline.

- The timeline chart is now a scatter chart in contrast to the line chart for DRS 1.0 dumps.
- The timeline navigation and legend are different.

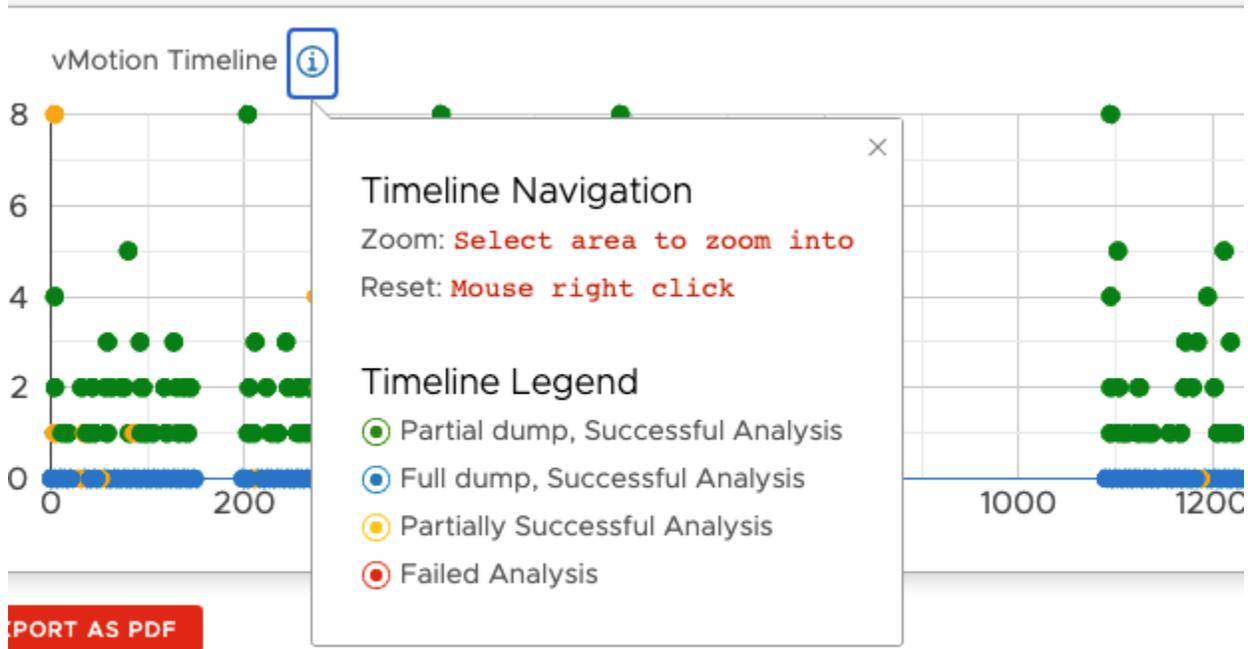


Figure 11. Clicking the circled i icon shows the timeline navigation and timeline legend. Just click any area of the timeline and right-click to zoom in.

Understanding the dots

What are the green dots?

A green dot/point indicates a successful partial dump analysis, which can be a breakdown of:

- A DRS load balance run
- DRS placement decisions while powering on VMs in a powered-off state (DRS initial placement)
- DRS placement decisions while evacuating VMs from a host that is being put under maintenance mode
- DRS placement decisions to resolve rule violations in the cluster (mandatory moves)

What are the blue dots?

A blue dot/point indicates a successful full dump analysis, and this dot always lies on the x-axis (that is, the number of vMotions is 0). The reasons for this follow.

- A full dump analysis only indicates what move (vMotion) recommendations DRS would have come up with if it was run at that point in time.
- Any actual move (vMotion) recommended and initiated by DRS in the cluster will have an associated partial dump.

If these moves from a cluster dump analysis are not actual moves observed in the cluster, why do you see them?

Even if the moves are not actual moves observed in the cluster, their analysis is useful to understand cases where DRS is not recommending any moves.

Consider a scenario where DRS did not recommend any vMotions—there will not be any associated partial dumps in the DRS 2.0 dump to analyse. However, full dumps are available in the DRS 2.0 dump at periodic intervals of 5 mins, which can be analysed to understand the cluster state.

By default, DRS Dump Insight does not analyse all the full dumps when a DRS 2.0 dump is uploaded. Only the full dumps near the partial dumps are analysed. This is to reduce the analysis time because there can be hundreds of full dumps in a single DRS 2.0 dump, which leads to a significant wait time for the analysis to complete.

If you need to analyse all the full dumps (in cases where no partial dumps are available due to DRS not initiating any vMotions), enable **All full dumps** when uploading the dump files.

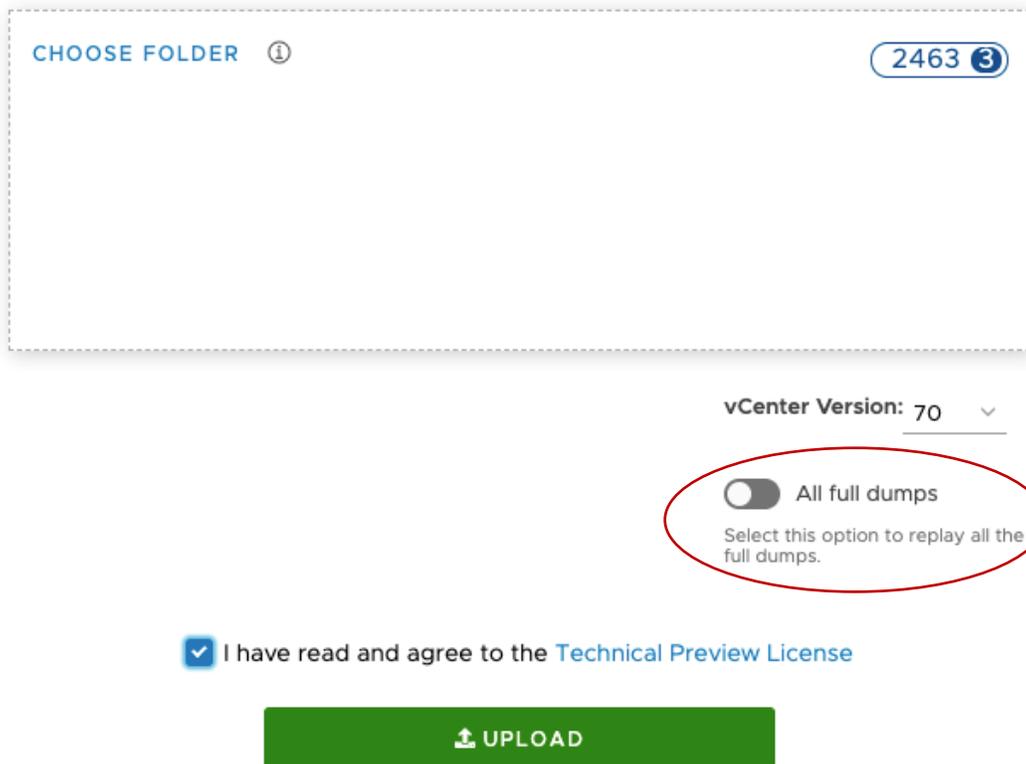


Figure 12. In this screenshot, you would need to enable **All full dumps** (as marked in red) before uploading your dump file. You would click the toggle to enable it.

What are the yellow/red dots?

A **yellow dot/point** indicates a dump that was partly successful. In this case, the DRS replayer might have failed, resulting in missing data. The *DRS replayer* is the tool that DRS Dump Insight invokes to replay the DRS dump file. *Replay* means to recreate the cluster from the dump and run DRS again to understand DRS behaviour when the dump was created.

A **red dot/point** indicates a failed partial or full dump analysis.

Summaries of partial and full dump analyses

Here, we give sample output of the previous scenarios.

Partial dump analysis summary

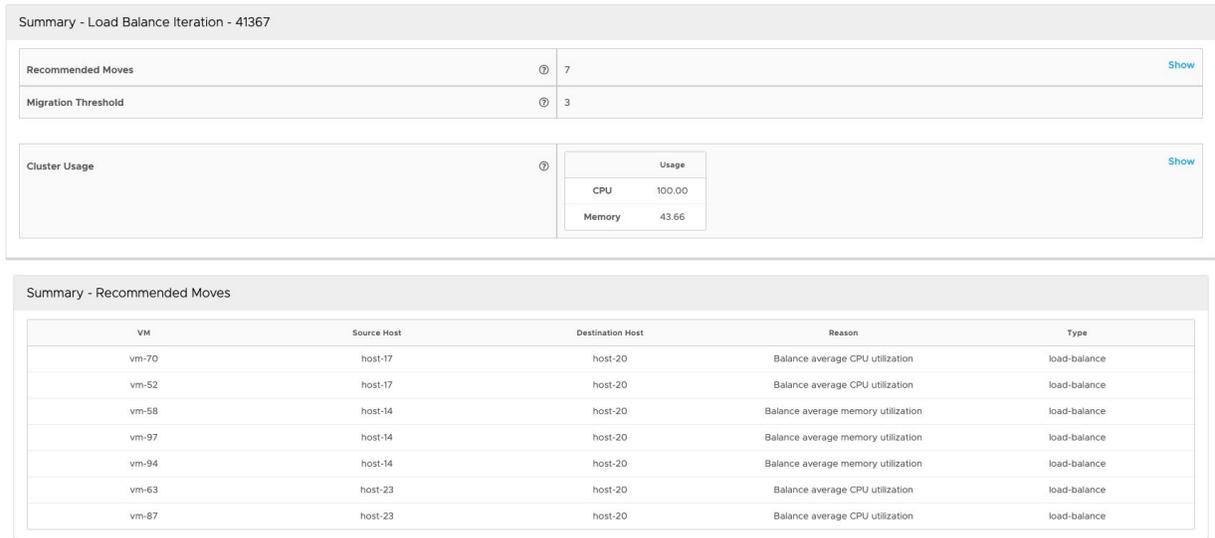


Figure 13. This is an example of a partial dump analysis summary.

Full dump analysis summary

This summary indicates DRS behavior if it was run at that point in time.

The full dump analysis summary lists the number of moves out of the total considered that can and cannot improve the load balance.

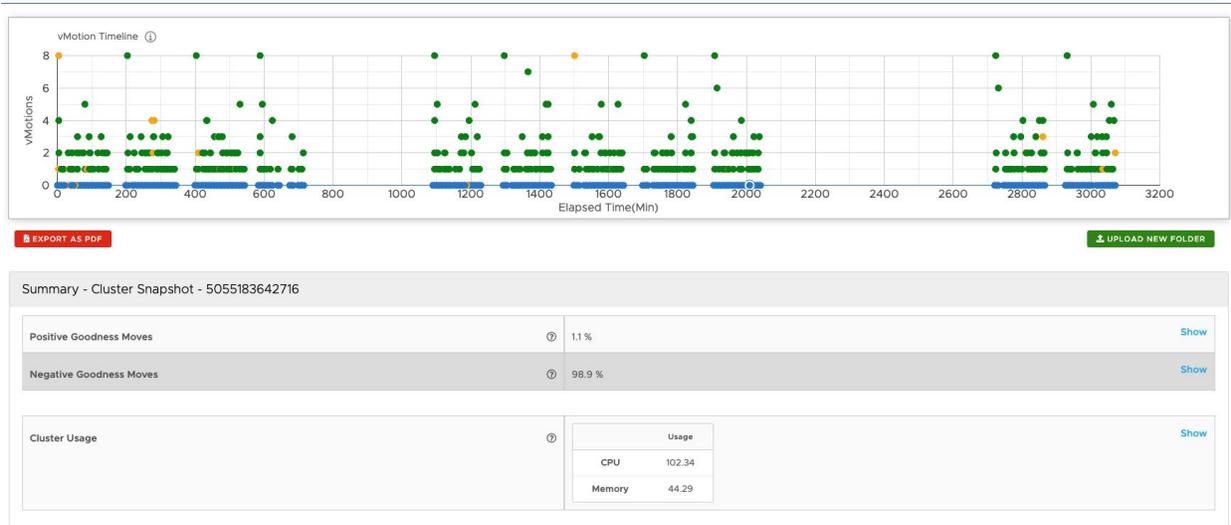


Figure 14. This is an example of a full dump analysis summary.

The full dump analysis summary does not contain a list of recommended moves due to the reasons indicated in the “[What are the blue dots?](#)” section.

Exporting results of a DRS 2.0 dump

You can export the results from your run to a PDF file.

From the vMotion Timeline:

1. Click **Export As PDF**.
2. On the dialog box that appears, click **Generate PDF**.

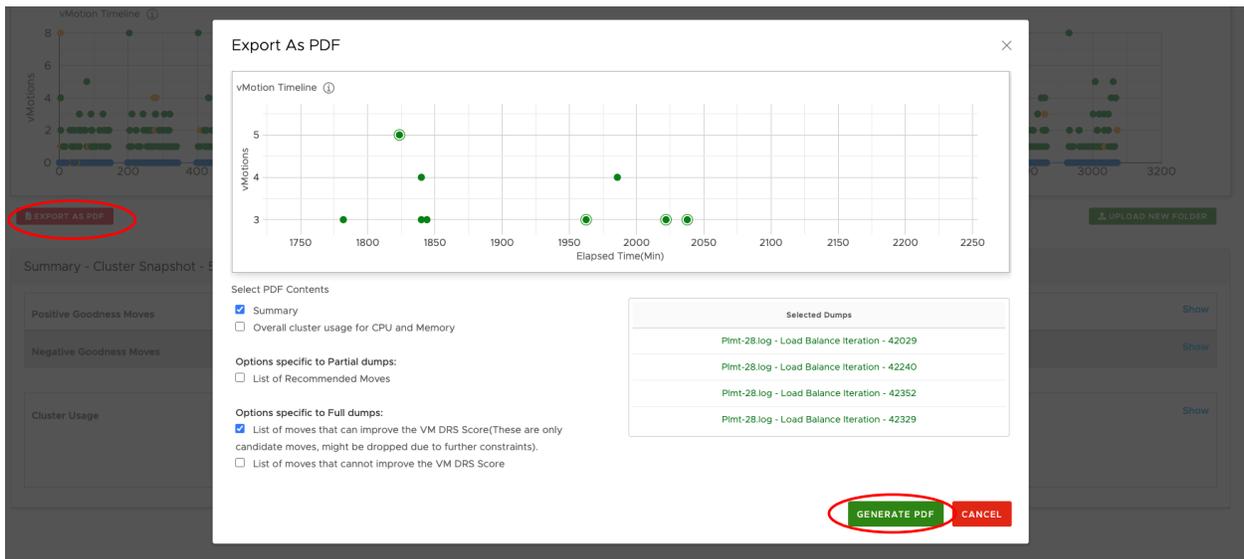


Figure 15. On the **Export As PDF** dialog box, click **Generate PDF**.

Conclusion

This guide has shown how to use the DRS Dump Insight web app.

- To log in to the web app, go to <https://www.drsdumpinsight.vmware.com>.