

Cisco Workload Optimization Manager 2.2 Released!

The Turbonomic and Cisco teams have released our next Cisco Workload Optimization Manager platform with the most recent update to version 2.2, packed with much more cloudy goodness and also with the addition of new targets and more features in the cloud for both planning and real-time optimization.

One of my favorite parts of building out what I love to call the Cisco Stack is that the integration from the application (Cisco AppDynamics), down to the containers (Cisco Container Platform), into the virtualization layer (VMware, Hyper-V, OpenStack), and down to the metal (Cisco UCS, Cisco HyperFlex) including the network (Cisco Nexus, Cisco ACI, Cisco Tetration).

What's Inside CWOM 2.2?

Big updates in this one also include cloud pricing enhancements, custom pricing (any cloud), Azure CSP rate cards, and the ability to do future planning for reserved capacity purchases on AWS which is not available in any platform to-date.

The release aligns with Turbonomic version 6.3 for the feature set so you can get a quick view of what's inside the latest in my sizzle reel here:



You can visit the main solution page here:

<https://www.cisco.com/c/en/us/products/servers-unified-computing/workload-optimization-manager/index.html>

Full User Guide for CWOM 2.2 here:

https://www.cisco.com/c/dam/en/us/td/docs/unified_computing/ucs/ucs-workload-optimization-mgr/user-guide/2-2/cisco-ucs-wom-user-guide-2-2.pdf

Download the latest edition here:

<https://software.cisco.com/download/home/286321574/type/286317011/release/2.2.0>

Full list of Turbonomic and Cisco partner resources here: <https://resources.turbonomic.com/cisco>

[Top vBlog 2018 Results Show - Live Webinar](#) [3/21 @ 11:30 AM ET](#)

It's that time of year again! Big thanks to Eric Siebert, Angelo Luciani, John Troyer and yours truly for the unveiling of the top 25 Virtualization Bloggers of 2018. We are super excited to see who will be on the list and most importantly, how many new faces and blogs are showing up among the overall top 100.

I'm happy to be part of the Turbonomic team sponsoring this year's show and I'll be providing some fun commentary along the way during the event.

Jump on in to the webinar registration which is here:

https://turbonomic.zoom.us/webinar/register/WN_RxOEeDKVTOa3kBy8v1GGZA

The event will be held LIVE on Thursday March 21st at 11:30 AM Eastern Time. We are hosting it on Zoom and everyone is welcome!



See you on the live show!

[Quick Demo of Tags and Grouping for Automation Policies in Turbonomic 6.1](#)

Tagging is a phenomenal way to identify your workloads. This means VMs, containers, cloud instances, and pretty much every single layer of the application and virtualization/cloud stack. I'm often asked how tagging comes into play with Turbonomic, so here is a really quick demo of how tagging is used to be able to do things like creating dynamic groups, policies for automation, and many other important use-cases.

The ability to pull in tags automatically and dynamically without having the guess work makes it so

much easier than the way many products handle tagging. Especially when we think of the full-stack coverage which you're going to need to use it for when it comes to the Turbonomic platform.

More tagging is always good!

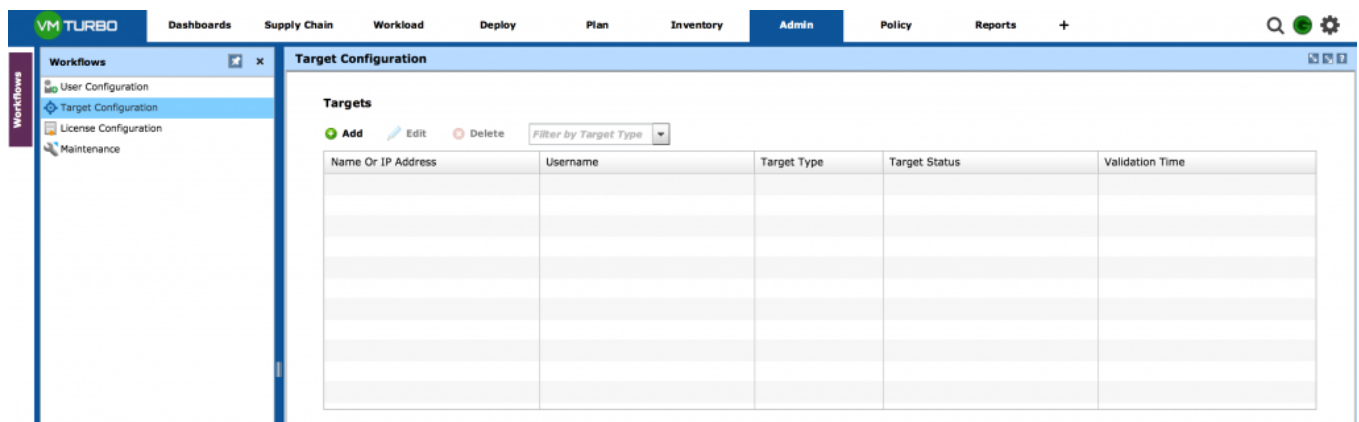
[Attaching Turbonomic to your AWS Environment](#)

While it's a seemingly simple task, I wanted to document it quickly and explain one of the very cool things about attaching your Turbonomic instance to AWS. For the latest release of the [TurboStack Cloud Lab giveaway](#), we wanted to move further up the stack to include the AWS cloud as a target.

Even without the TurboStack goodies, you can already attach to AWS and get back great value in a few ways. Let's see the very simple steps to attach to the cloud with your control instance.

Attaching to an AWS Target

First, log in to your server UI with administrative credentials that will allow you to add a target. Go to the **Admin** view and select **Target Configuration** from the workflows panel:



Click on the **Add** button and enter the following details:

- **Address:** *aws.amazon.com*
- **Username:** Your Access Key from AWS
- **Password:** Your Access Key Secret from AWS

Target Configuration

Target Type: vCloudDirector CloudStack OpenStack VMM AWS Azure SoftLayer

Hypervisor

Cloud Management

PaaS

Load Balancer

Storage

Fabric

Network

Operations Manager Appliance

Application Server

Database

Windows Application

Address:

Username:

Password:

Address: Leave blank, or enter aws.amazon.com
 Username: Enter the Access Key ID for this account
 Password: Enter the Access Key Secret for this account

Pending targets

[Add](#) [Edit](#) [Delete](#)

Name Or IP Address	Username

Next click the **Add** button in the Pending Targets area below the form, then press the **Apply** button at the very bottom. That will take you to the next step which validates your configuration:

Targets

[Add](#) [Edit](#) [Delete](#)

Name Or IP Address	Username	Target Type	Target Status	Validation Time
aws.amazon.com	AKIAJ2J AKIAKIAJ2J AKI	AWS	Validated	Tue Jul 5 16:54:32 GMT-0400 2016

Now that you are validated, you will begin to see data populating as it is discovered from AWS. The discovery cycle runs by default every 10 minutes, and as each entity is discovered, it is polled asynchronously from there for continuous gathering of instrumentation.

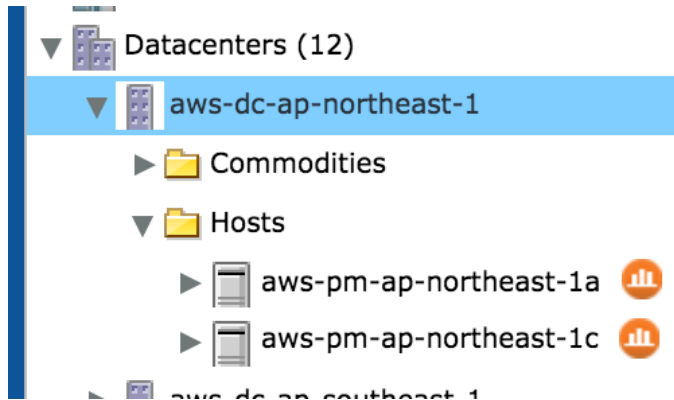
In your **Inventory** view, you will see the addition of AWS entities under the **Datacenters**, **Physical Machines**, **Virtual Machines**, and **Applications** sections:

The screenshot shows the VM Turbo Inventory view with the following components:

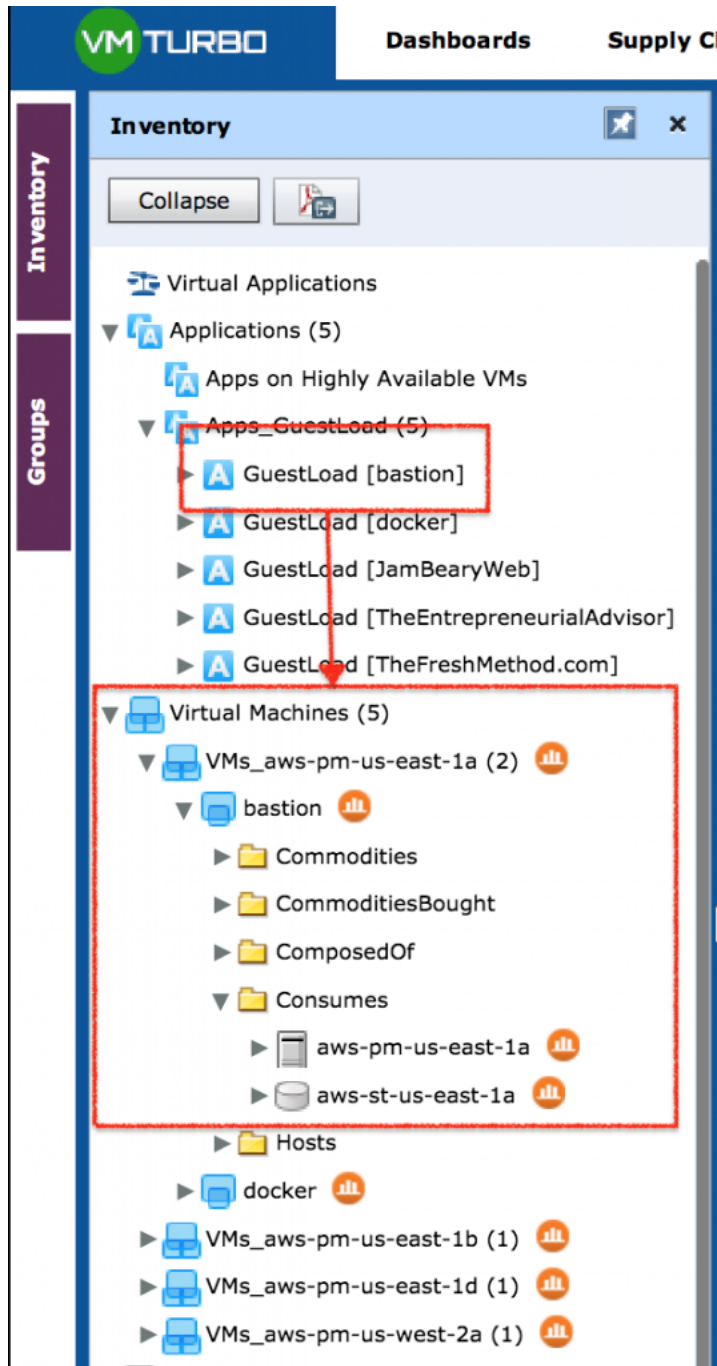
- Left Sidebar:** A tree view showing the hierarchy: Inventory > Groups > Datacenters (10). Under Datacenters, several AWS regions are listed, such as 'aws-dc-ap-northeast-1', 'aws-dc-ap-southeast-1', 'aws-dc-cn-north-1', 'aws-dc-eu-central-1', 'aws-dc-eu-west-1', 'aws-dc-sa-east-1', 'aws-dc-us-east-1', 'aws-dc-us-west-1', and 'aws-dc-us-west-2'.
- Top Panel:** 'Utilization - Physical Machines By UI' bar chart showing metrics for UI, Mem, CPU, IO, and Net over time.
- Right Panel:** 'Trend - Physical Machines' line chart showing % utilization over time.
- Bottom Panel:** 'Usage - Physical Machines' table with columns: Name, UI, Me..., CPU..., IO %, Net %, Buff..., Swa..., Ball..., ICP..., 2CP..., 4CP..., #VMs. The table lists various physical machines across different AWS regions.

If you expand one of the **Datacenters**, you will see that it is defined by **Regions** (example: ap-

northeast-1) and then underneath that, you can expand to see the **Availability Zones** represented as **Hosts**:




Let's expand our Applications, and Virtual Machines, where you see the stitching of the entities across all of the different entity types:



You can see that we have a Virtual Machine (EC2 instance) named **bastion**, which also has an **Application** entity, and you can see that it consumes resources from the us-east-1a AZ, with an EBS volume in the same AZ.

You can also see the cumulative totals under the **Virtual Machines** list to get a sense of how many instances you have running across the entire AWS environment. The running instances are counted in brackets at the end of each region listing. How cool is that?! As someone who forgets instances all the time running across numerous regions from testing, this has been a saving grace for me.

You can also use the Physical Machines section to view each region. When you drill down into the PMs_ section, you will see the AZ listings underneath there.



That's all it takes to get your AWS environment attached, and we will dive into some other use cases on what you can do for application level discovery, and much more on the AWS public cloud.