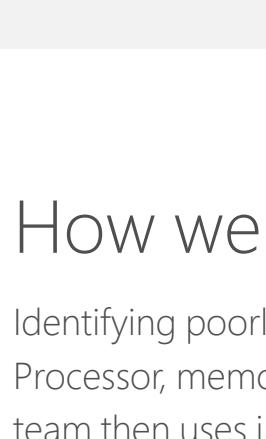




## Optimizing resource efficiency in Microsoft Azure

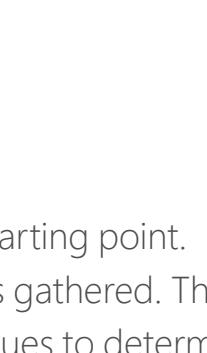
Microsoft IT is moving the bulk of its computing resources to Azure. The team first focused on identifying on-premise resources that could be completely eliminated. Migration of remaining resources to Azure followed. But once migrated, use of the Azure environment requires evaluation and optimization. To that end, Microsoft IT has set standards to manage resources, identify underutilized and noncompliant workloads, and to retire those that are unused.

A focus on efficiency is providing a significant reduction in cloud spending, and an increase in CPU utilization efficiency across our infrastructure in both IaaS and PaaS, while still aggressively reducing our on-premises footprint.



90%

Computing resources hosted in Azure by July 2017



400%

CPU efficiency increase across Infrastructure as a Service (IaaS)



-38%

Reduction in cloud spend

## Investing in Azure at Microsoft

Microsoft computing resources hosted on Azure IaaS (not including PaaS)

20,000

Azure virtual machines

110,000

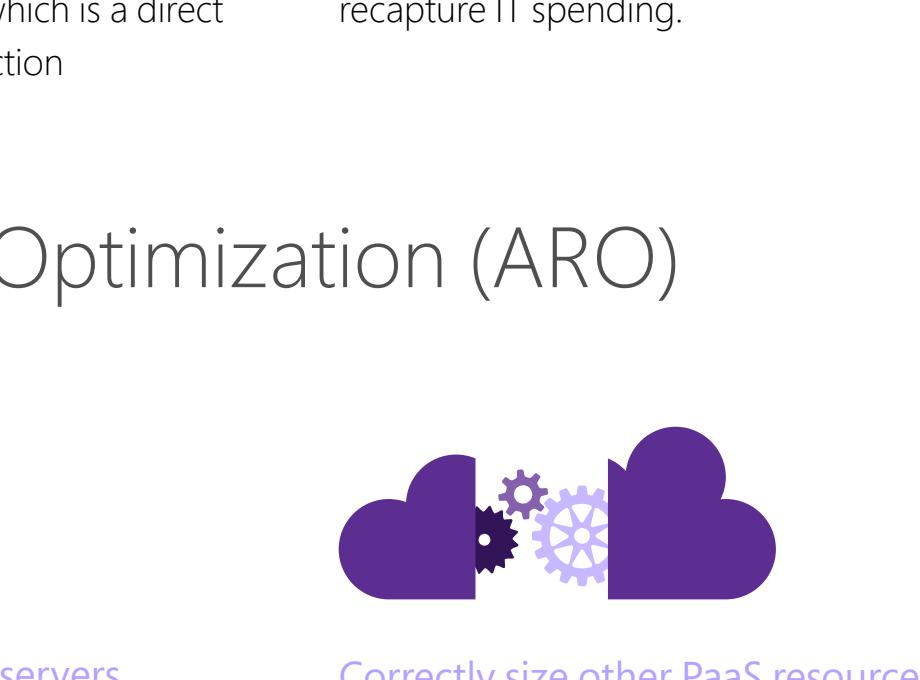
cores utilized

1,500

Azure-based applications

## How we optimize

Identifying poorly utilized servers is the starting point. Processor, memory, and hard drive data is gathered. The team then uses industry-standard P95 values to determine whether specific assets are underutilized. We categorize all servers into five performance categories: frozen, cold, warm, hot, and on fire. Underutilized servers that are on-premises are prime candidates to move to the cloud.



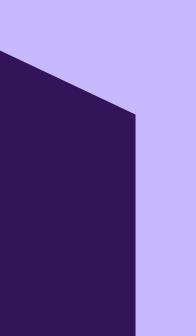
## Optimizing resource efficiency in Microsoft Azure

Our toolset involves both graphical and automated solutions that interact with our Azure environment:



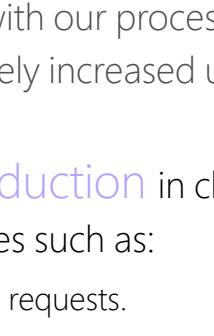
### Measure

Our internally developed dashboard is where it all starts. Aggregating a simple cost view with the recommendations across IaaS and PaaS resources shows how much of current spending could be removed without business impact.



### Snooze

Non-production cloud servers need to be online only when employees are actively working on them. We created a suite of on demand and scheduled tools which allows us—in some cases—to turn off, or de-allocated, over 70 percent of the time, which is a direct



### Resize

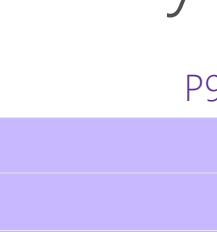
Azure provides the ability to quickly change the size of a server. With a simple reboot, dropping the size and cost of a virtual machine by more than half, at large scale, allows us to quickly use the flexibility of the cloud to recapture IT spending.

## Top tips for Azure Resource Optimization (ARO)



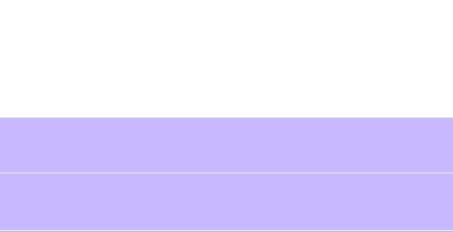
### Turn off unused servers

We see non-production servers can generally be turned off and save over 70% of the time and cost. It especially makes sense on nights and weekends.



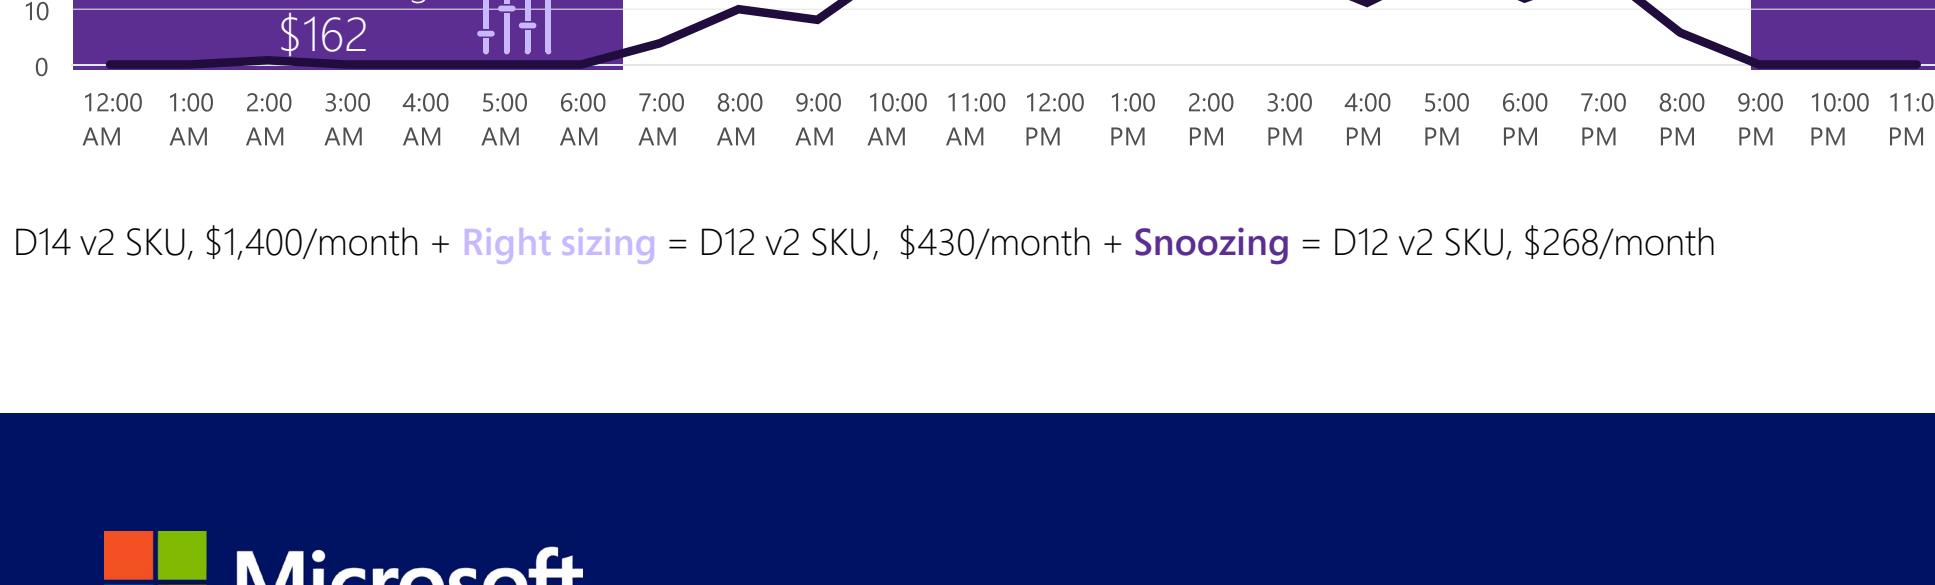
### Correctly size your servers

Look at your CPU and memory utilization, and select the right size that brings your workloads into the 40-80% CPU range.



### Correctly size other PaaS resources

For example, target your SQL PaaS databases to hit between 40-80% average DTU utilization.



## How do we save money?

