

# Managed AWS Services

## Part 3: Managed AWS Services



*This white paper is the last in a series of three that examines managed AWS services and hybrid cloud environments. We've already identified and debunked 7 myths about AWS services ([Part 1](#)), and presented a roadmap of how MSPs can get started offering AWS services ([Part 2](#)). Now we take on the key challenges MSPs face when offering AWS services to their customers, as well as the challenges faced by enterprises, the users of AWS services.*

Whether adding managed AWS services to your MSP service catalog or migrating some of your enterprise workload to AWS services, the new hybrid environment will present numerous challenges to you and your business. This paper seeks to explore some of the most common ones, and to present actual solutions to overcoming the impediments, bumps, and roadblocks involved in managed AWS services.

Favoring practice over theory, we base the content of this white paper on real-life examples drawn from actual AWS operations. We look at specific challenges from both the MSP and the enterprise perspectives, including:

- the need for visibility
- the cost of deployment
- the right payment model
- using performance metrics to save money
- the costly “Lone Wolf” syndrome
- virtual private cloud security
- load balancer issues
- avoiding outages

We conclude with solutions to achieving end-to-end visibility of hybrid environments and the applications that run in them, and we offer some suggestions for both MSPs and enterprises on how to get the most out of their managed AWS services.

## The Need for Visibility

An often overlooked challenge that crops up immediately after migrating to the AWS cloud — or any cloud for that matter — is the degree of transparency and control required for hybrid cloud environments. Once you migrate a workload to a cloud provider to improve efficiency and agility, you have de facto split your resources between on-premises and off-premises. That dichotomy makes viewing end-to-end the health and availability of the business service that the application delivers very difficult. Suddenly, you can no longer see how your business service is doing.

You could use local cloud tools to look at deployed instances in a third party cloud, but you wouldn't be able to see a correlation between the applications running on top of those instances and the ones running on-premises. Even harder to see is how storage relates to your compute cycles in the new environment. Blindness and ignorance in the hybrid world make it difficult to know when there are outage or performance problems, and they make it just about impossible to do root cause analysis.

You need someone to provide the visibility and knowledge you must have in order to understand the hybrid world you now inhabit.

## The Cost of Deployment

If you migrate large and IO-intensive workloads into the cloud, you are likely to encounter sticker shock, especially if you initiate IaaS deployments without any reservations or contracts in place. A



lack of control over the total volume of autoscaling permissible for ones instances will only make matters worse. To avoid sticker shock and to execute more strategic cloud deployments, you should balance the seemingly high cost of an IaaS deployment with your historical operations, MTTR, licensing, human resources, networking, storage, and hardware maintenance and operational costs.

Very defensible and reliable calculators exist to help you with this pre-budgetary planning. These calculators can eliminate the fear of an insufficient budget as an impediment to moving to the cloud. They incorporate all the above-mentioned sunk costs into the equation, and showcase the long term TCO reduction available with cloud when you include all the variables specific to your operation.

Those calculators are available from the AWS website itself, and also from ScienceLogic. By contrast, our platform has billing and chargeback capabilities built in and included at no additional fee.

## The Right Payment Model

Many IT administrators fail to determine the right payment model for their cloud deployment before shifting workloads over to AWS. What is the best



payment model for your company — spot pricing, on-demand pricing, or reserved instances? Which services and what size should you choose? You need to know the answers to these questions ahead of time.

While AWS has dropped its prices 42 times since 2006, the pricing paradigm you choose will determine your ultimate AWS costs. Choosing ineptly can cost you plenty. You need to spend time determining, based on your specific cloud needs, when to use:

- **spot prices** (instances whose prices are bid on and used until a higher bid comes through)
- **on-demand pricing** (by the hour), and
- **reserved instances** (for dedicated or committed resources)

Additionally, you need to assess which among the approximately 40 AWS services (and innumerable iterations within each) best suits your requirements. Then you'll have to decide how best to aggregate, plan and limit the costs of those services, including which services to reserve, and by what size.

These variables seem to make cloud deployments a lot of guesswork for first time users. So before migrating to AWS, customers like you are in need of an apples-to-apples comparison between what the cloud offers and your internal KPIs and costs. Comparison tools exist. Here are some examples.

## Using Performance Metrics to Save Money

Like many cloud providers, AWS now makes historical performance metrics available, which extend beyond a matter of days but not on a per region basis. At ScienceLogic, we have found that material differences in performance can come at a significant cost. Inefficient performance can result



in overall costs that outweigh the lower prices on offer, especially for cloud platforms stretching across multiple regions.

ScienceLogic has a lot of parity with AWS's CloudWatch from a billing perspective. But when considering trends over time, favorable comparison leans strongly toward ScienceLogic because of our extensive and lengthy archives of performance metrics. Since CloudWatch does not keep metrics for long periods of time, tracking which availability zones have a greater likelihood of outages is very difficult. That makes assessing zone choice nearly impossible — and probably very costly. For example, if users pay the same price for deploying a particular Amazon service in two different availability zones, but Zone A has a much greater history of outages, it makes sense to deploy in Zone B — but only if you know the performance metrics for both zones over time.

Herein lies the MSP opportunity — without you, enterprise customers would not know about the disparity and could continue to spend money and deploy to zones that don't provide the best performance for their dollar. Your expertise can save them money — and gain you business and a solid reputation.

## The Costly “Lone Wolf” Syndrome

When firing up a few EC2 instances with other AWS services attached to those EC2 workloads, a user will almost surely employ a Business Analytics app — the number one application being outsourced to the cloud today and otherwise known as a Big Data app. In firing up the EC2 instances, the user first has the option to specify how many EMR (Elastic Map Reduce, which uses Hadoop to process large amounts of data) instances they want to spin up with the EC2 instances. When the process ends, the user will shut down the EC2 instances but, more often



than not, will forget about the EMR jobs that, having run, are still sitting in the background. This so-called “lone wolf” syndrome of several not-in-use EMR jobs — not running but just sitting out in the ether — costs the company money for resources not being used. Once the EC2 instance has been removed, there's no obvious way, without help, to detect that these unused, wasteful workloads are still sitting out in the ether.

But help for this costly syndrome is available. MSPs, enabled by ScienceLogic, can expose any EMR jobs tied to an EC2 instance directly to their customers, and can detect any “lone wolves” by making visible the complete topology of any virtual private cloud.

By resolving this common syndrome with ScienceLogic's help, MSPs can save their customers a discernible, and often significant, amount of money by detecting their unused EMRs and terminating them.

## Virtual Private Cloud Security

Virtual private clouds (VPCs) are everywhere these days, and their numbers are constantly multiplying. Organizations will often deploy multiple VPCs, each usually with its own set of security policies. MSPs may have hundreds of VPCs, and as many as one or two for every client. When deploying an instance in AWS, users can select from the list of available VPCs. With so many VPCs being set up, the probability that human error will come into play is increasingly likely.




In addition to enormous compliance issues, VPCs present a host of security challenges. While Amazon has gone to great lengths to provide additional, and impressive, security capabilities, some of those advantages can, by themselves, present a problem. One example is that when a VPC is created it is automatically given a random unique ID. While on the surface this is good, it can also result in a user deploying the wrong instance to the wrong VPC. In fact, once you have deployed all of your instances, you will be hard pressed to tell whether your instances are deployed in the correct VPCs.

Using ScienceLogic's dynamic device groups eliminates problems created by human error in VPC deployment. For example, with workloads deployed from multiple organizations, you can quickly and automatically detect when a deployed instance is aligned to the right organization but the wrong VPC. ScienceLogic alerts you to the problem immediately for ready remedy.

## Load Balancer Issues

You may choose to deploy a number of load balanced EC2 instances across multiple regions and multiple availability zones. But because the most common AWS monitoring tools have significant



**If a group of load balanced AWS instances spread across multiple regions has an issue, it is difficult to find the instance causing the problem.**

limitations, they will not allow you to see a load balanced group across regions. Instead, you will be able to see EC2 instances in only a single region at a time.

This restrictive view can be highly problematic for you. If your workloads encounter poor performance from a load balanced group, you would be extremely hard pressed to accurately detect which EC2 instances were the culprits.

With ScienceLogic's CloudMapper, you are able to see all load balanced EC2 instances across all regions in one dynamically updated map. In a matter of seconds, you could detect which EC2 instances were causing issues with the load balanced group.

## Avoiding Outages

Let's suppose you have a number of Amazon users and a number of availability zones. The users are merrily instantiating Amazon instances of the services of their choosing. During the deployment of those services, the users are asked to which availability zone they would like to deploy. The users go with the default choice.

Initially, this free-reign, laissez faire practice may not pose a problem. But over several deployments, the lack of oversight and protocol may well lead to the potential for lost money and an increased amount of risk. Perhaps, for example, all instances are deployed in only one availability zone, and that zone just happens to be the most expensive availability zone. In this scenario, two problems have developed at this point: heightened risk has evolved because all of the users' eggs are in one basket and cost has risen by defaulting to the most expensive availability zone.

Visibility across all users' accounts and where



those users' resources are deployed is simply not practical, and in most cases not possible, with legacy tools. MSPs enabled with ScienceLogic, by contrast, have the advantage of having all of these data points automatically detected and tracked by CloudMapper. You are privy to an uber view of everything, including all the specifics of every user account. Each of your customers can instantly see and control in which availability zones their resources (and only theirs) reside. You can determine whether you should expect increased risk/cost based on any of their allocations.

## The Arrival of End-to-End/ Hybrid Application Visibility

Getting peak value out of hybrid IT requires a level of management and visibility that traditional tools, cloud vendor tools, and point solutions simply cannot provide. Case in point: given your current set of tools, could your operations team automatically discover and track the dependencies between your

onsite IT infrastructure and your offsite Amazon-based resources? And can you think of any way they could possibly do the same thing for your customers?

As an application owner in an enterprise, or as a service provider offering an application as a service, you may want to use Amazon on occasion to offload some of your work. Currently, you have no way of looking at the service across both your onsite resources and your offsite Amazon resources, and certainly not in a single view. Furthermore, providing that view to your clients is simply impossible.

ScienceLogic has taken on the challenge of end-to-end/hybrid application visibility, and we have succeeded in developing the tool of insight that enterprises and MSPs — and all their customers — have needed to maximize performance, minimize costs, and maintain total visual control. (See Figure 1.)

---

### Acceleration:

- Speeds up off-prem cloud adoption.
- Enables confident datacenter consolidation.
- Clarity on CAPEX to OPEX cloud transition.

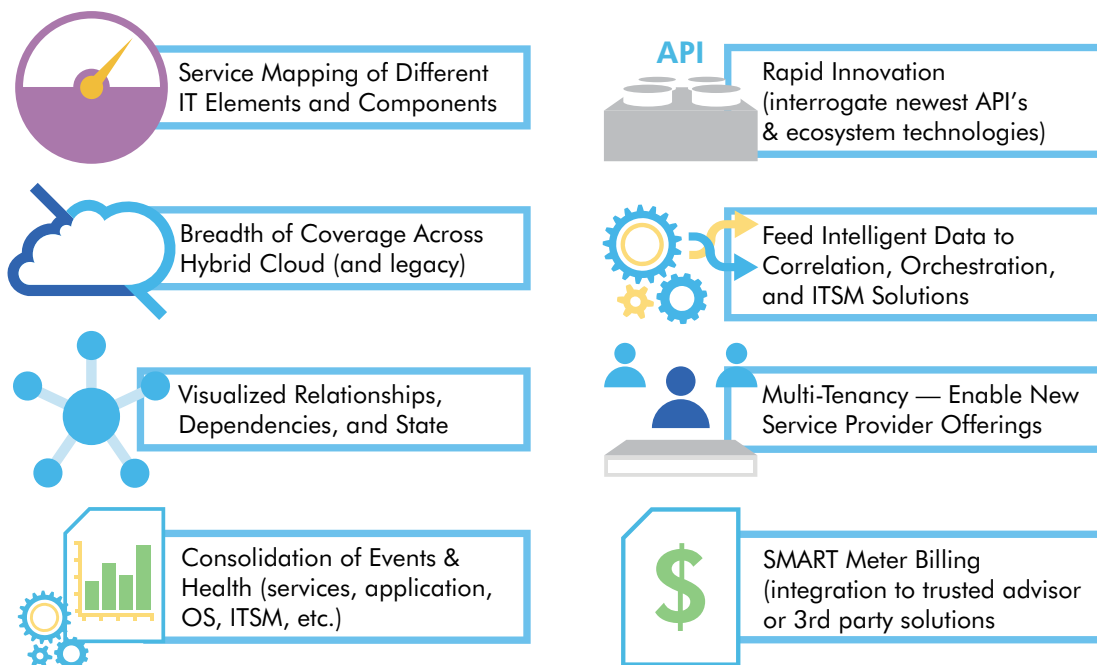
### Validation:

- Relationship between services and changes made on instances.
- Single pane of glass trouble shooting, on all assets.
- Verification for continued Public Cloud Migration.

### Visibility:

- See legacy infrastructure, your hosted environment & AWS Cloud together in one place.
- Multi-tenancy shows services by technology, user, organization, etc. (ITaaS or MSP models).
- Historical views enable intelligent spending and decision making.

**Figure 1: ScienceLogic Benefits for Enterprises and Managed Service Providers in Managing AWS Services**



**Figure 2. ScienceLogic features enhance life in a hybrid environment with end-to-end visibility and complete control over all components, physical and virtual.**


Currently, enterprises and MSP customers cannot attain a holistic view of an application as a service with the tools they have. They must view the service from AWS-only monitoring tools, such as CloudWatch. These limited monitoring tools run separately from their onsite tool, and they prohibit a holistic view of the service and its impacting dependencies. For example, when looking at EC2 instances alone, CloudWatch provides no evident correlation with the application running on top of those instances. Further, since — up until now — no one has offered any “relationship detection” tools, application owners and service providers have not been able to see such particulars as how their storage relates with their compute.

The combination of ScienceLogic’s AWS monitoring and our CloudMapper capability lets you see your entire service, all of its supporting elements, as well as all of the instances supporting it in the public cloud. (See Figure 2.)

End-to-end visibility of all parts of a hybrid service or all instances of a service is essential for an application owner. If the application being offered as a service based in AWS experiences an outage, the owner would not know unless he checked AWS’s CloudWatch, its onsite NMS Tool, and the MSP’s tool — all at once. By contrast, ScienceLogic gives you total visibility on a single pane of glass of your entire hybrid environment end-to-end.

Adding AWS as a managed service or expanding into the cloud with it? ScienceLogic has all the help you need.

We helped debunk misinformation about AWS in Part 1 of this three-part white paper, [Seven Myths about AWS Service](#), and we gave you the lowdown on getting started in Part 2, [How to Get Started with AWS Service](#).



## About ScienceLogic

ScienceLogic delivers the next generation IT monitoring platform for the network of everything. Over 15,000 global Service Providers, enterprises, and government organizations rely on ScienceLogic every day to significantly enhance their IT operations. With over 1,000 dynamic management Apps included in the platform, our customers are able to intelligently maximize efficiency, optimize operations, and ensure business continuity. We deliver the scale, security, automation, and resiliency necessary to simplify the ever-expanding task of managing resources, services, and applications that are in constant motion.

ScienceLogic won InfoWorld’s 2013 Technology of the Year award, Red Herring’s Global 100 Award, Deloitte’s Technology Fast 500™, and MSPmentor 250, among other worldwide recognitions of excellence. For more information, visit [www.sciencelogic.com](http://www.sciencelogic.com).



## Contact Us

### Americas

+1.800.724.5644  
info@sciencelogic.com

### Europe

+44 (0) 203.603.9889  
info-eu@sciencelogic.com

### Asia-Pacific

+61 2 9959 2426  
info-apac@sciencelogic.com