White paper

PLATFORM9

# SaaS architecture for managing private & edge clouds

# Contents


SaaS managed clouds: Market validation 04

Architecture & key IP for SaaS management 05

SaaS managed cloud operating model architecture 05
SaaS Management Plane 05
Extensibility to new cloud services 08
Support for diverse infrastructure environments 09

SaaS management SLA & KPIs 10

Cloud operations SLA 10
Rapid, repeatable deployments 10
Server: Admin ratio 10

How Platform9 works 11

SaaS Management Plane 11
Software profiles standardize deployment and configuration 12
Operational tooling 14

Platform9 and the SaaS Managed cloud competition 17

Google Anthos 17

AWS Outposts 17

Microsoft Azure 18

VMware Tanzu 18

Summary: The future of clouds is SaaS managed 19

# SaaS managed cloud: Market validation

The recent Gartner report "The Many Faces of Private Cloud" provides information about "as-a-service" (XaaS) implementation models. Regarding outsourced private cloud model, the report states that

> These solutions are based on an outsourced control plane (usually VM IaaS or Kubernetes) that is run as a SaaS-style offering by a provider. In this case, however, the physical hardware remains on the customer's premises and is managed through agents or gateways from the SaaS control plane. This type of environment allows the customer to maintain hardware in a private location for compliance or regulatory reasons but outsources the complexity of running the control plane to the provider. It can also be used to support edge use cases when the customer does not want to manage the control plane.

Top characteristics of this solution category include:

1. **The infrastructure flexibility of traditional private clouds.** This includes the ability to use existing customer infrastructure and edge locations thereby avoiding lock-in to a cloud provider's infrastructure

2. **A rich cloud services catalog** and the simplicity of cloud operations of hyperscale public cloud providers

# Architecture & Key IP for SaaS management

## SaaS managed cloud operating model architecture

SaaS Managed clouds are composed of three distinct layers of IP and systems software:

1. The cloud services catalog (Open Source Cloud Services), like that available in the public cloud, is often derived from the open source community, such as Kubernetes.
2. These cloud services are managed by the core, a SaaS Management plane (SaaS Cloud Operations) that automates hyperscale cloud operations without requiring traditional and profession- al services.
3. Infrastructure Management: These automated cloud services can be run on various infrastructure sites, from traditional data-centers to distributed edge locations, or even across public cloud regions. This flexibility is provided by an infrastructure management layer that provides integration into various compute, storage and networking configurations.

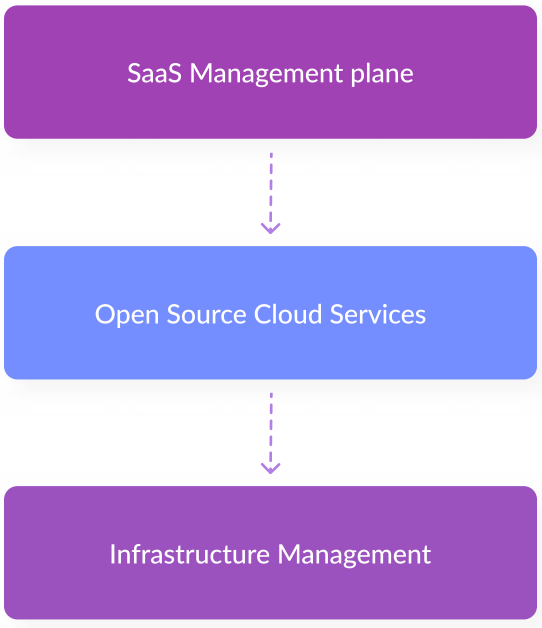


3-layer architecture for SaaS Management of distributed clouds

## SaaS management plane

The SaaS Management Plane provides the operational automation for the consumption of the infrastructure that is managed by the Infrastructure Management layer via the rich cloud services catalog delivered by the Cloud services layer.

## Management plane detailed view

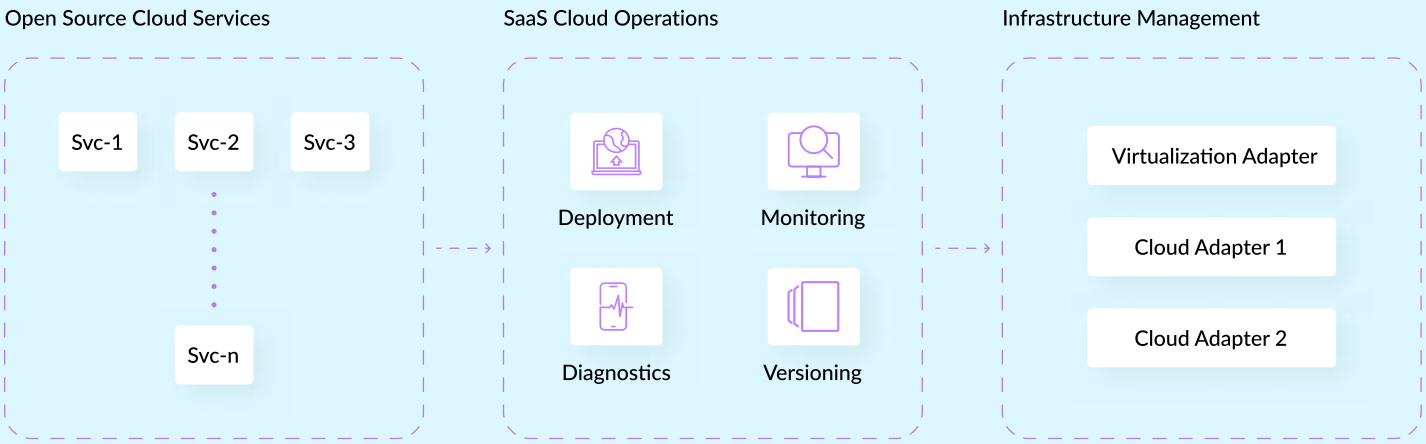


Automation functions of the SaaS management plane

The core functions of the SaaS management plane support operational automation of the cloud service lifecycle.

### Deploying a cloud service

- Modern cloud services, such as Kubernetes, are actually composite services that are them- selves highly distributed. Therefore, deploying a service such as Kubernetes means support- ing a more complex orchestration capability.
- Deploying control plane components - for example, Kubernetes master nodes components such as API server, or etcd - requires addressing the redundancy and high availability capabilities of these components
- In addition to control plane components, most cloud services will also require data plane components - for example, Kubernetes worker nodes components such as Kubelet and docker. These worker nodes make sense only in the context of a certain master node. There- fore, the system must support dependency resolution and integration of data plane components with control plane components.

## Monitoring a cloud service

- Every cloud service and its components can define a set of health metrics which must be monitored on a continuous basis
- Since modern cloud services are highly distributed, a small degradation in certain compo- nents can lead to a larger system-wide degradation in time. In order to simplify troubleshoot- ing (whether automated or human), and in order to mitigate the likelihood of larger problems emanating, these health probes need to be highly granular.

## Applying diagnostics to automate/aid troubleshooting

- Having good health metrics provides a basis to codify resolution of common problems via automated runbooks
- These runbooks can be built for common problems that happen in the course of normal system operation - for example, when a control plane went offline because of an infrastruc- ture failure.
- These runbooks are also effective when there are problems in new versions of cloud ser- vices, or interoperability issues that are found only in the field after deployment at some scale. Since the runbook can be implemented without requiring a new version of the cloud service, immediate mitigation can be provided while a bug-fix or a new version of the cloud service in question is developed. In this way, customers can be operational despite the com- plex, ever-evolving nature of modern open source cloud technologies.

## Versioning to automate non-invasive, repeatable upgrades

- The breadth of developers and vendors participating in modern open-source ecosystems means that new versions are constantly being developed, both with bug fixes as well as with security and feature enhancements
- The Management plane makes it easy for customers to stay up to date by fully automating the upgrade to a new version of various cloud services
- These upgrades are typically offered on a granular basis (e.g.: upgrading Service A should be independent of upgrading Service B), which makes change control easier for large scale enterprise deployments
- Finally, these upgrades are ideally offered in a “self-service” manner that enables customers to schedule their own upgrades at a time that is convenient for them, and at a scope of their choosing (e.g.: upgrade the Virginia data-center at 3am on Saturday; but leave Mumbai un- touched for now)

## Extensibility to new cloud services

Modern open-source ecosystems are evolving rapidly both in their breadth and depth. Some of the major ecosystems and their initiatives include:

- The CNCF backed cloud-native suite of technologies includes not just Kubernetes, but also emerging service mesh technologies such as Istio and Linkerd, application monitoring and alerting technologies such as Prometheus, and is poised to extend to other services by using Kubernetes operators as a common operational model.
- OpenStack, KubeVirt and other virtualization technologies not considered core to the CNCF roadmap are in significant demand by enterprises and service providers that need to run traditional applications (that will not be deployed via microservices in the near future)
- Enterprises receive all of these technologies as part of the public cloud provider catalog and expect the same breadth to be available for private and edge cloud deployments. There- fore, the SaaS management plane must be extensible beyond individual open-source frame- works to ensure that the catalog of cloud services grows over time to meet customer use cases:

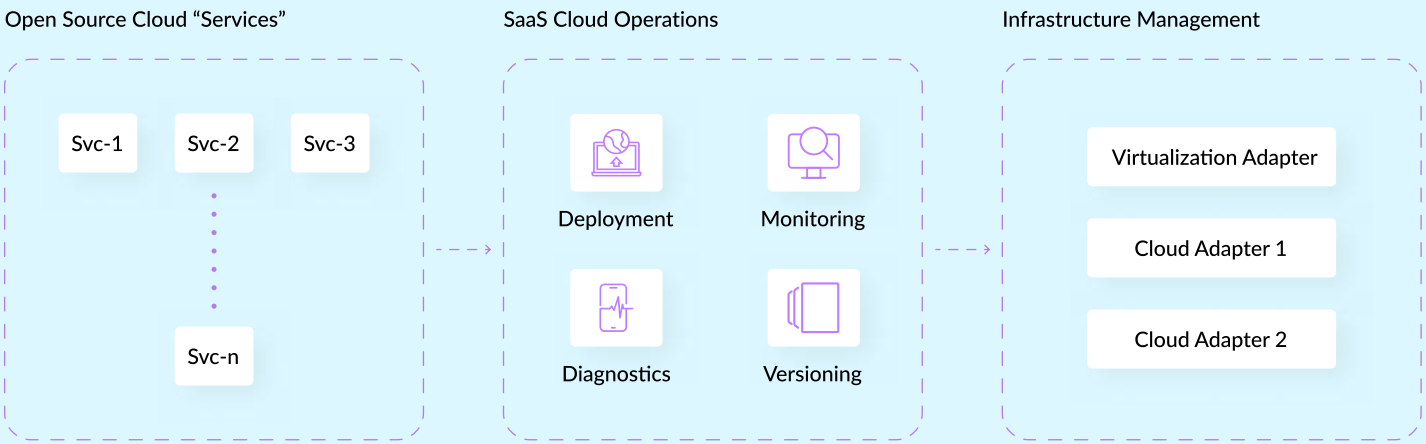


Extensibility enables new cloud services to be integrated from open-source ecosystems

As described above, this requires a high degree of flexibility in the management plane to enable complex orchestration of new cloud services.

## Support for diverse infrastructure environments

Unlike public cloud providers who can standardize their infrastructure deployment SKUs, supporting enter- prise private and edge cloud entails integrating with a large and complex set of configurations:

- A variety of physical infrastructure environments, in particular, storage and networking environments.
- Support for running in centralized core data-centers for private clouds; but ability to operate with very low “touch” for distributed edge environments.
- Support for running on physical machines, or running on existing virtualized environments (e.g. VMware) or deploying virtualization seamlessly using KVM.
- Hybrid deployments combine one or more of the above with one or more public cloud providers (typically, AWS, Azure or GCP).

In order to do this, the architecture needs to have extensibility to support a variety of infrastructure “plugins” that abstract the differences in these environments and enables consistent management regardless:

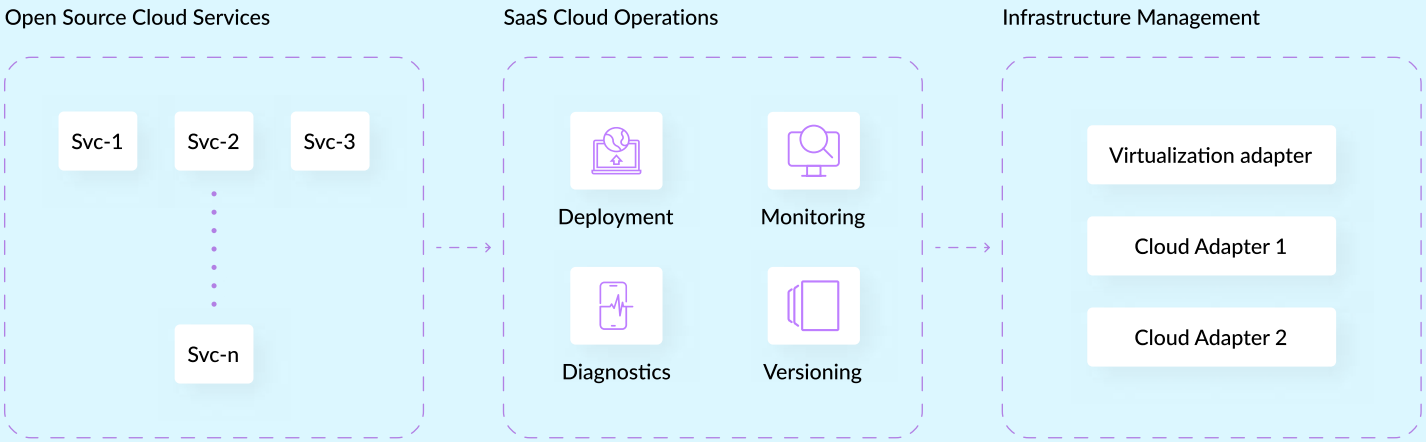


Infrastructure plugins enable integration with a variety of environments found in enterprise private, edge and hybrid clouds

By ensuring consistent management across all of these environments, this architecture delivers the proverbial “single pane of glass” for managing clouds.

# SaaS management SLA & KPIs

As described above, SaaS management entails automating every step of the lifecycle of managing cloud services deployments. This automation provides several benefits:

## Cloud operations SLA

By automating not just deployment, but health monitoring, run book driven resolution of common problems, and streamlining upgrades; it is possible to provide an operational SLA for distributed clouds. Hitertho, this was only available via hyperscale public clouds, but enterprises who desire the simplicity and peace of mind of public cloud computing can now get a similar SLA anywhere

## Rapid, repeatable deployments

Another facet of public clouds is how easy it is for end users to deploy workloads onto their cloud services. In contrast, enterprise private and edge clouds can take a long time to implement. By automating the initial deployment and configuration of environments, this time can be greatly reduced to within minutes or hours (from weeks or months previously).

## Server: Admin ratio

Talent constrained IT teams have long struggled with the complexity of running large scale private clouds. The ratio of servers managed to an admin or automation architect can be as little as 40:1 in private clouds. In contrast, hyperscale public cloud providers have invested significantly in automating the management of their environments, which greatly improves their admin efficiency. It is not unheard of for the server : admin ratio in public clouds to be 4000:1 or more. For enterprise hybrid clouds to be competitive, SaaS management enables a similar ratio but in private, edge and hybrid architectures. This greatly reduces costs and increases the likelihood of success for enterprises

# How Platform9 works

## SaaS management plane

Platform9's cloud operations teams deploy and manage distributed infrastructure sites centrally and remotely, without requiring onsite delivery. This is done by leveraging a centrally deployed management plane, and orchestrating the delivery of various capabilities -- Hypervisor, Storage backend, Network back- end, etc. -- via automation and operational tools to the physical infrastructure sites. This process is depicted conceptually below:

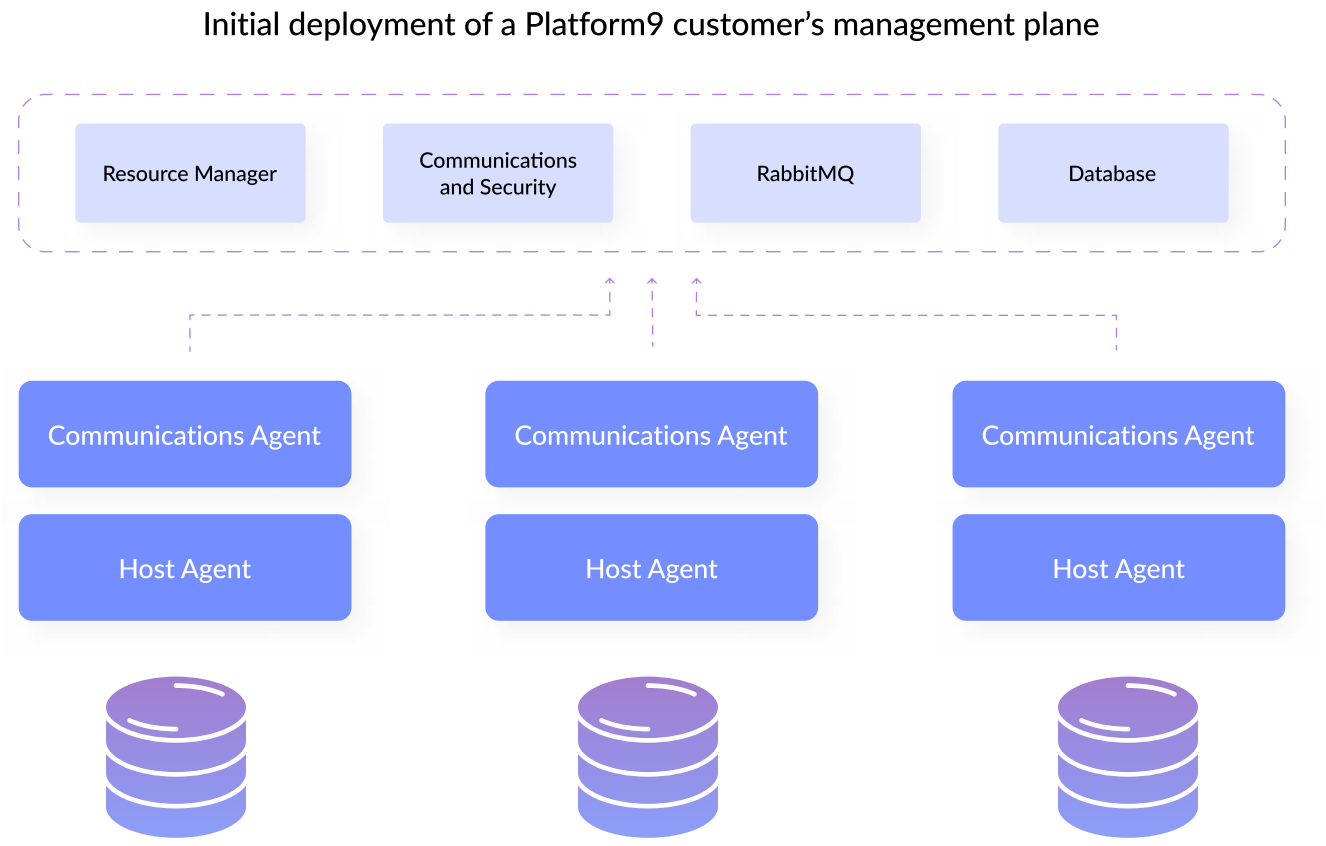


Per-Customer Deployment Unit (Management Plane)

## Software profiles standardize deployment and configuration

The management plane uses software profiles (listing of software packages, versions and dependencies) and configuration profiles (listing of key configuration parameters) to ensure that nodes under management are provisioned consistently, with a high degree of repeatability. This provides several benefits:

- New nodes can be deployed with confidence, since the software profiles and configuration profiles are well known, tested, and cannot drift
- Existing nodes that fail can be rebuilt quickly and reliably
- Manual work, troubleshooting, and firefighting are reduced

### The list of key machine templates includes:

- Hypervisor (uses openstack-nova, libvirt and other)
- Storage backend (uses openstack-cinder)
- Network node (uses openstack-neutron)
- Storage node (uses openstack-ceph)
- Container-visor (uses kubernetes kubelet, kubeproxy, docker and other)
- Storage node (uses openstack-ceph)
- Kubernetes master (uses etcd and other)

Management plane managing a site using Kubernetes (for containers)

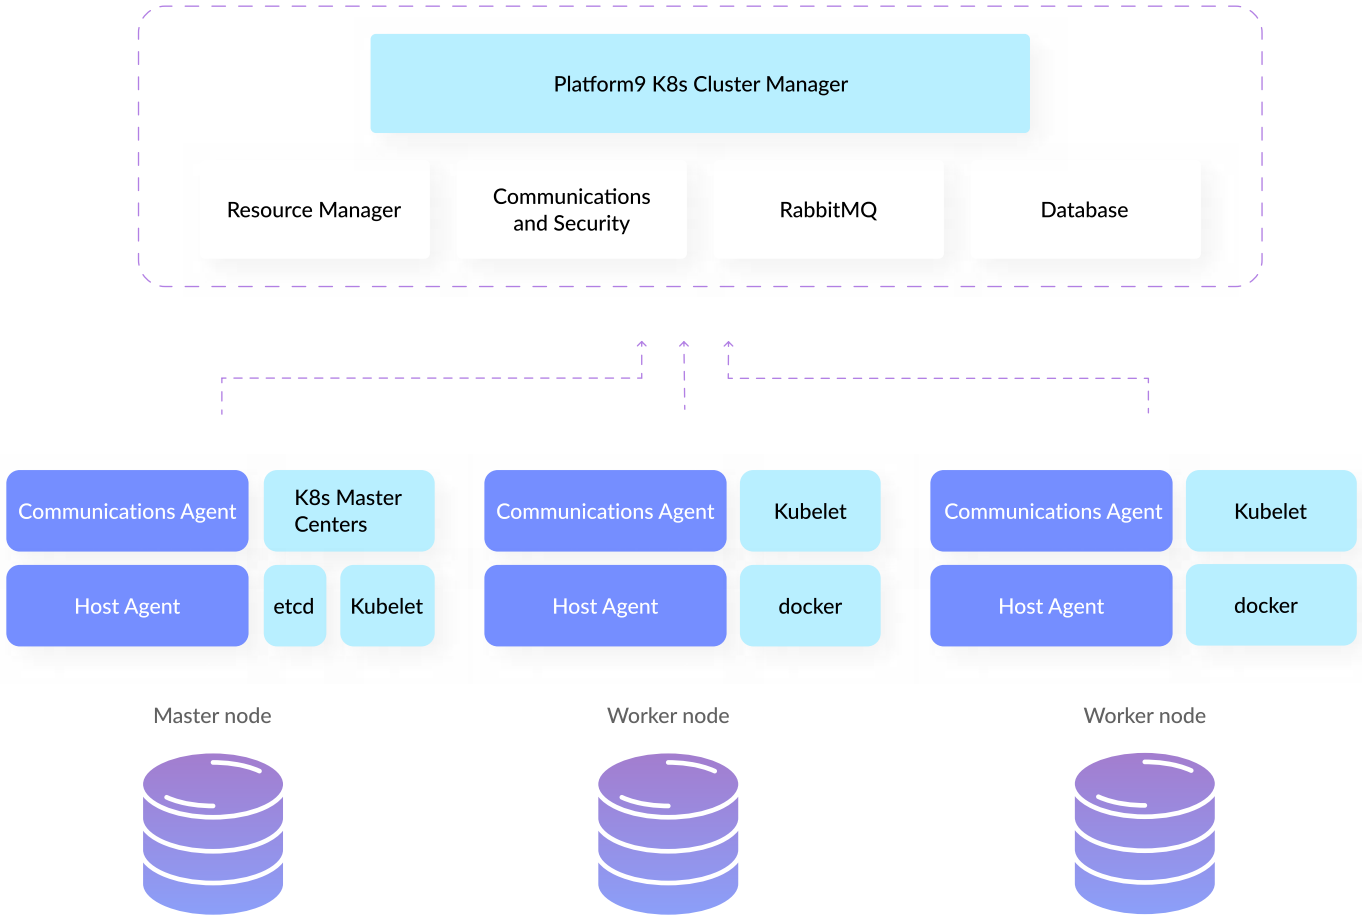


Per-Customer Deployment Unit (Management Plane)

## Operational tooling

Platform9 operations teams use a variety of tools for service delivery. The teams monitor and troubleshoot customer installations using metrics gathered in **Prometheus & Grafana**; and efficiently search across large log sets using log aggregation.


Prometheus, Grafana and Log Aggregation are used to monitor key metrics, provide related alerts and analyze log traces

This in turn is used to reduce the “actionable alerts” that Platform9’s support teams need to watch, making it easier to quickly identify open issues as well as the different sites that may be affected by them

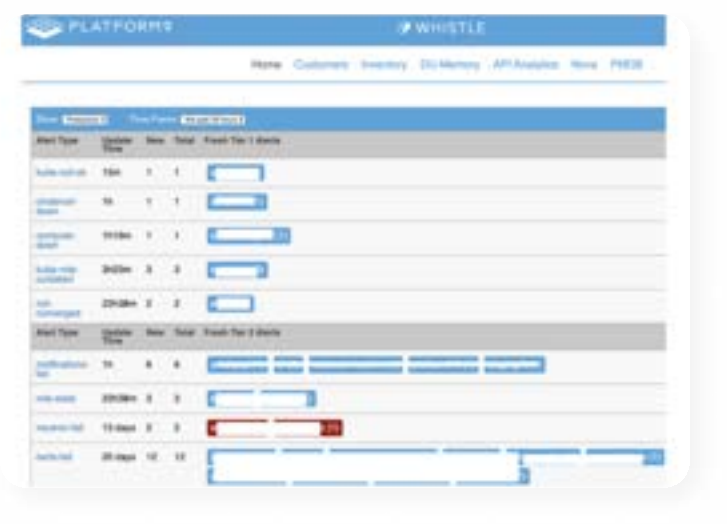

Metrics & alerts reduced to actionable alerts and impacted sites (site names are redacted in this picture for customer confidentiality)

Log aggregation is based on the Elasticsearch - Logstash - Kibana (ELK) framework:

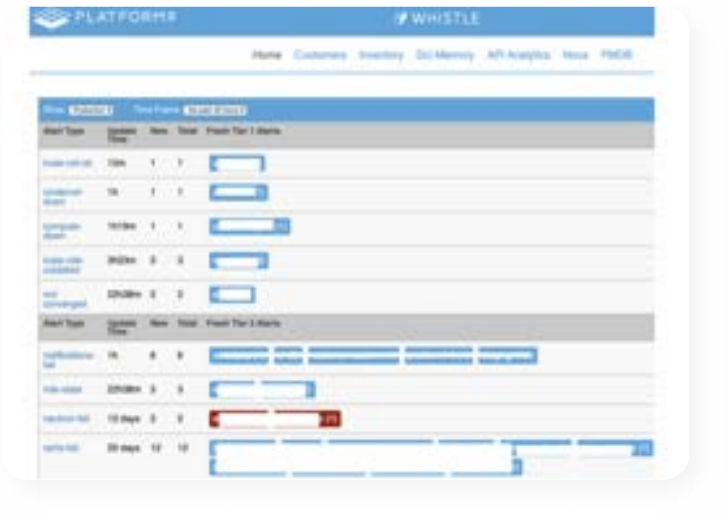

Log aggregation helps accelerate troubleshooting and analysis

## Customer proof points

Kingfisher

“

As we move into the future of containerized microservices, our focus will be to be as cloud-native as possible and have a CI/CD delivery, while freeing ourselves from the burden of recruiting and retaining staff to run the Kubernetes stack. Platform9’s SaaS managed service eliminates the cloud operational burden so we can focus on enabling cloud-native services for the business”

Digital Platform Manager

**Read the case study**

**Data Services company**

“

Between 8,000 to 10,000 VMs could spin up within a 24-hour period and be destroyed within 20 minutes. We spend a lot less time worrying about management and scaling with Platform9 (virtualization). It has been a perfect fit for our on-premises data centers. It just works!"

IT infrastructure Manager

**Read the case study**

# Platform9 and the SaaS managed cloud competition

Platform9 was the first company to offer a "SaaS Managed" cloud operational model for private and edge clouds. The company introduced the service in 2016 and has deployed it in several Global 2000 customers. Significant competitors in this space include:

| Competitor | Solution/Notes |
|---|---|
| Google Anthos | **Solution Summary:**<br>Google Anthos is a distributed cloud platform running on premises and edge.<br><br>**Key Solution Notes:**<br>• Google Anthos extends its public managed Kubernetes service to customer on-premises data centers running VMware environments.<br>• More recently, Google announced nascent support for bare metal. |
| AWS Outposts | **Solution Summary:**<br>AWS Outposts extends AWS public cloud computing services into customer's data centers.<br><br>**Key Solution Notes:**<br>• AWS Outposts requires that the customer buy hardware directly from AWS.<br>• AWS Outposts is also fully-managed, some might even say a "white glove" service, with end to end hands-on support from AWS.<br>• This service approach ensures that the solution's starting and ongoing pricing is very high along with hardware compatibility restrictions that come with it.<br>• AWS Outposts also does not run or manage workloads outside of AWS on oth- er public clouds or any other infrastructure the customer already owns. |

| Competitor | Solution/Notes |
| --- | --- |
| Azure<br>Microsoft | **Solution Summary:**<br>Azure Stack is similar to AWS Outposts in that it essentially extends their Azure public cloud computing services into on-premises data centers.<br>**Key Solution Notes:**<br>• Azure Stack charges fees based on the cloud services consumed. Further, it requires users to purchase hardware that is certified for the platform if they want to use it to host Azure Stack services on-premises.<br>• Although Azure Stack requires users to manage their infrastructure themselves, professional support from Microsoft is available as part of the core platform offer- ing. The customer, however, can't use Azure Stack to manage hybrid workloads hosted on public clouds and onpremises |
| VMware Tanzu | **Solution Summary:**<br>VMWare Tanzu is a Kubernetes offering from VMWare that runs on premises and across multiple clouds.<br>**Key Solution Notes:**<br>• The solution requires a number of other VMware storage, networking, and virtualization products to provide a complete suite.<br>• VMware provides a traditional enterprise class support model, for Tanzu, not a fully managed approach. |

# Summary: The future of cloud is SaaS managed

The success of the public cloud has educated the world about what is possible when infrastructure can be consumed at scale with the speed and simplicity made possible by SaaS management. The days of private clouds being managed in more traditional, labor intensive ways are numbered. The speed, savings and simplification of SaaS based management is significant; and enterprises will choose from either entirely outsourcing the cloud to public cloud providers, or using a SaaS Management model for their private clouds. For the emerging category of edge cloud computing, geographic distribution of infrastructure and work- loads limit the reach of the public cloud. Similarly, distributed edge environments need to be managed centrally, with little to no “touch”. It is clear that SaaS based management will be the de-facto standard for distributed cloud management

Find your customized private cloud strategy

Get a free consultation

PLATFORM9

# About Platform9

Platform9 empowers enterprises with a faster, better, and more cost-effective way to go cloud native. Its fully automated container management and orchestration solution delivers cost control, resource reduction, and speed of application deployment. Its unique always-on assurance™ technology ensures 24/7 non-stop operations through remote monitoring, automated upgrades, and proactive problem resolution. Innovative enterprises like Juniper, Kingfisher Plc, Mavenir, Redfin, and Cloudera achieve 4x faster time-to-market, up to 90% reduction in operational costs, and 99.99% uptime. Platform9 is an inclusive, globally distributed company backed by leading investors.

Follow us on

Headquarter: 800 W El Camino Real, #180, Mountain View, California 94040, US

Phone: 650-898-7369 | Phone: https://platform9.com/contact/

© 2023 Platform9. All Rights Reserved