

A Cloud Architect's Guide to Open Distributed Clouds

Leverage the power of public clouds using existing infrastructure anywhere — without vendor lock-in



Contents

Introduction.....	2
A SaaS control plane	3
Deployment	3
Monitoring.....	4
Diagnostics and troubleshooting.....	4
Versioning and upgrades	4
Standardized deployments with profiles	5
Managed open-source services	5
Cloud-native	5
Virtualization	5
Bare metal	5
Support for diverse, distributed infrastructure	6
Platform9's open distributed-cloud service.....	6
Operational tooling.....	7
A cloud-operations SLA.....	8
Key features and benefits of the Platform9 open distributed-cloud service	9

Introduction

Enterprises are expanding their use of cloud computing in ways that increasingly span multiple public clouds, private data centers, and edge locations. A distributed cloud lets companies leverage the power of cloud computing at any of the above locations to deliver greater business value with new use cases previously not feasible.

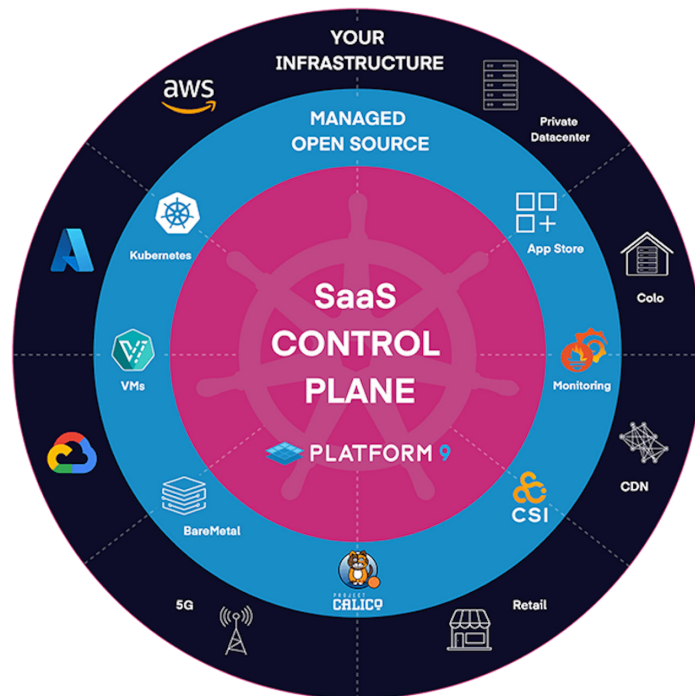
However, implementing and operating distributed clouds is hard if you lack considerable skill and experience. Critical challenges include:

- **Deployment delays** — It can take weeks or months to implement a cloud particularly in a private or edge environment. And integrating that environment with other public clouds is not trivial.
- **Each deployment adds silos** — Implementing a cloud at multiple locations adds to operating overhead unless you can consistently manage these locations with a central control plane.
- **Diverse infrastructure and services compound operational complexity** — Troubleshooting problems and keeping platform software up to date at each location is an ongoing operational burden.
- **Needed talent and skills are hard to find** — Experts in Kubernetes, virtualization, and cloud-native technologies are comparatively rare and expensive.

An ideal solution is to leverage the power of public clouds with the ability to use existing infrastructure at any location — without vendor lock-in.

This white paper describes the three main components of an open distributed-cloud service and introduces Platform9 as the #1 such provider:

- **A SaaS control plane** that makes it delightfully easy to build and operate clouds anywhere.
- **Managed open-source services** that use 100% open-source stacks and components to deliver bare-metal, container, virtualization, and supporting platform services.
- **Support for diverse, distributed infrastructure** with out-of-the-box plugins and integrations for public clouds (AWS, GCP, Azure), public-cloud Kubernetes services (EKS, AKS, GKE), and multiple operating systems such as CentOS and Ubuntu.



A SaaS control plane

This control plane provides operational automation for the consumption of infrastructure and is the heart of a distributed cloud service. It delivers the following benefits:

- **Low management overhead** using a highly automated hyperscale operational model.
- **Reduced maintenance costs** by aggregating all distributed infrastructure behind a single management pane.
- **Rapid and repeatable remote deployments** to 100s or 1000s of distributed cloud locations with consistent, template-based configuration and policy control.
- **An operational SLA** through automated health monitoring, runbook-driven resolution of common problems, and streamlined upgrades.

A SaaS control plane includes several components that make it easy to deploy and manage clouds anywhere.

Deployment

Modern cloud services such as Kubernetes are actually composite services that are themselves highly distributed. Therefore, deploying a service such as Kubernetes means supporting a more complex orchestration capability. Deploying control plane components such as Kubernetes master node components (API server, etcd) requires addressing the redundancy and high availability capabilities they need.

In addition to control plane components, most cloud services also require data plane components such as the Kubernetes worker node components kubelet and Docker. These worker nodes make sense only in the context of a certain master node. Therefore, the system must support dependency resolution and the integration of data plane components with control plane components.

Monitoring

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 components can lead to a larger, system-wide degradation in time. To simplify troubleshooting (whether automated or human), and to mitigate the likelihood of larger problems, these health probes need to be highly granular.

Diagnostics and troubleshooting

Having good health metrics provides a basis to codify the resolution of common problems via automated runbooks. These runbooks can be built for common problems that occur during normal system operation such as a control plane going offline because of an infrastructure failure.

These runbooks are also effective when there are problems in new versions of cloud services 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 complex, ever-evolving nature of modern open-source cloud technologies.

Versioning and 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 SaaS control 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 (for example, 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 (for example, upgrade the Virginia datacenter at 3 a.m. on Saturday but leave Mumbai untouched for now).

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. A SaaS control plane and its automation capabilities enable a similar ratio for distributed clouds. This greatly reduces costs and increases the likelihood of success for enterprises.

Standardized deployments with profiles

A SaaS control plane uses software profiles (lists of software packages, versions, and dependencies) and configuration profiles (lists 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 is reduced.

The list of key machine templates includes:

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

Profiles also enable GitOps-style automation at scale by letting administrators define cluster configurations as desired state templates and then applying those templates to see how their environment has deviated. They can optionally remediate to correct a deviation. This can be applied to security policies, networking policies, RBAC configurations, and more.

Managed open-source services

Modern open-source ecosystems are evolving rapidly both in their breadth and depth. We describe some of the major ecosystems and their initiatives below.

Cloud-native

The CNCF-backed cloud-native suite of technologies includes not just Kubernetes, but also emerging service-mesh technologies such as Istio and Linkerd and application monitoring and alerting technologies such as Prometheus. Several other tools and services such as logging, ingress, and load balancers help implement production-ready enterprise deployments of cloud-native environments. A 100% open-source stack delivers the most flexibility and freedom from vendor lock-in.

Virtualization

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 soon).

Bare metal

For private data center and edge deployments in particular, the above open-source stacks need to run on bare-metal infrastructure. However, most bare-metal management today is manual. Open-source technologies like Ironi and Metal³ automate and offload manual bare-metal life-cycle management tasks. This lets users unleash the full performance of their hosts — no matter where they are located — as an elastic and flexible bare-metal cloud where they can rapidly deploy and redeploy any workload at a moment's notice.

A SaaS control plane is extensible beyond individual open-source frameworks to ensure that the catalog of cloud services grows over time to meet customer use cases. This extensibility enables new cloud services to be integrated from open-source ecosystems. This also requires a high degree of flexibility in the control plane to enable complex orchestration of new cloud services.

Support for diverse, distributed infrastructure

Distributed infrastructure spans multiple public clouds, private data centers, and edge locations. Many companies now have deployments across AWS, GCP, and Azure and more recently across hyperscale Kubernetes services such as GKE, AKS, EKS.

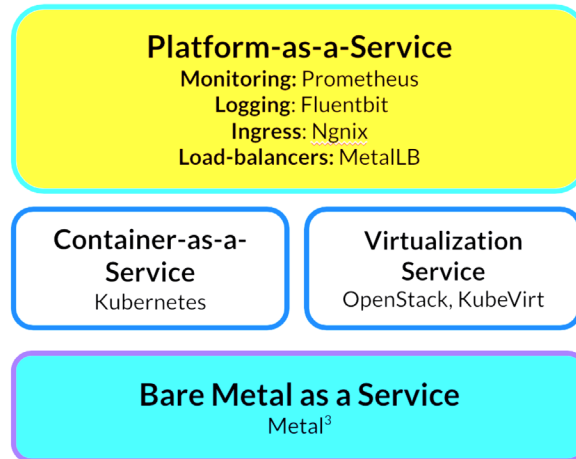
Unlike public cloud providers who can standardize their infrastructure deployment SKUs, supporting enterprise private and edge cloud entails integrating with a large and complex set of configurations:

- A variety of server, storage, and networking environments with varying SSO, security, and certification requirements.
- Support for running in centralized core datacenters for private clouds; but with the ability to operate with very low touch for distributed edge environments.
- Support for running on physical machines, existing virtualized environments (e.g., VMware), and seamless virtualized deployments using KVM.
- Hybrid deployments combine one or more of the above with one or more public cloud providers (typically, AWS, Azure or GCP).

To do this, the architecture needs to have extensibility to support a variety of infrastructure plugins that abstract the differences in these environments and enable consistent management. By ensuring consistent management across all these environments, this architecture delivers the proverbial “single pane of glass” for managing clouds.

Platform9's open distributed-cloud service

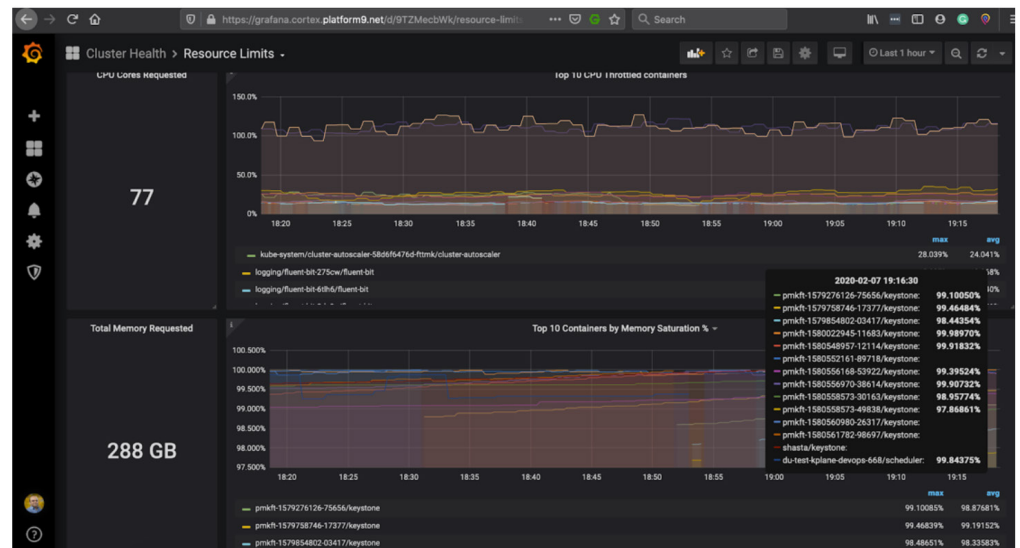
Platform9 delivers full life-cycle management of cloud-native services shown below using 100% open-source stacks and components.



Platform9 cloud operations teams deploy and manage distributed infrastructure sites centrally and remotely without requiring onsite delivery. This is done by leveraging a SaaS control plane and orchestrating the delivery of open-source cloud services via automation and operational tools to physical infrastructure sites.

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 and Grafana and efficiently search across large log sets using log aggregation.



This in turn reduces 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. The following diagram shows how raw metrics and alerts are reduced to actionable alerts and impacted sites (site names are redacted in this picture for customer confidentiality).

PLATFORM9

WHISTLE

Home

Customers

Inventory

DU-Memory

API Analytics

Nova

PMDB

Show:

Production

Time Frame:

the past 24 hours

Alert Type

Update Time

New

Total

Fresh Tier 1 Alerts

kube-not-ok

15m

1

1

s

cindervol-down

1h

1

1

a

(1)

compute-down

1h19m

1

1

c

(1)

kube-role-outdated

3h23m

3

3

s

(3)

not-converged

22h38m

2

2

e

Alert Type

Update Time

New

Total

Fresh Tier 2 Alerts

notifications-fail

1h

6

6

s

mysql (4)

st (2)

external-connection (1)

platform (1)

logging (1)

role-state

22h38m

3

3

e

(3)

neutron-fail

13 days

2

2

s

c (1)

certs-fail

20 days

12

12

j

(1)

s

(1)

Log aggregation to accelerate troubleshooting and analysis is based on the Elasticsearch Logstash Kibana (ELK) framework:

platform9.loggly.com/search#terms=json.logfile.%20%2Fvar%2Flog%2Fpf9%2Fkubernetes%2Fkubernetes.ERROR%20&from=2020-02-26T14:59:00.537Z&until=2020-02-26T15:00:00.537Z									
solarwinds loggly Search Charts Dashboard Alerts Derived Fields Source Setup Live Tail Metrics/APM New UI									
*Event Stream *eventrouter-crash *Tab 3 *api-server-log *Tab 5 *New									
All Sources json.logfile: "/var/log/pf9/kubelet/kubelet.ERROR"									
Event View Sort: Descending Expand Events More Options Normal view									
<div> <div>2020-02-26 09:58:42.267</div> <div>{ "cluster": "pmkft-prod", "nodename": "ip-10-31-136-200.us-west-2.compute.internal", "log": "E0226 17:58:37.564982 29228 pod_workers</div> </div> <div> <div>2020-02-26 09:58:40.305</div> <div>{ "cluster": "pmkft-prod", "nodename": "ip-10-31-136-125.us-west-2.compute.internal", "log": "E0226 17:58:37.717985 26067 pod_workers</div> </div> <div> <div>2020-02-26 09:58:27.430</div> <div>{ "cluster": "pmkft-prod", "nodename": "ip-10-31-131-6.us-west-2.compute.internal", "log": "E0226 17:58:25.734633 30992 pod_workers</div> </div> <div> <div>2020-02-26 09:58:27.267</div> <div>{ "cluster": "pmkft-prod", "nodename": "ip-10-31-136-200.us-west-2.compute.internal", "log": "E0226 17:58:22.565456 29228 pod_workers</div> </div> <div> <div>2020-02-26 09:58:25.369</div> <div>{ "cluster": "pmkft-prod", "nodename": "ip-10-31-136-125.us-west-2.compute.internal", "log": "E0226 17:58:24.718136 26067 pod_workers</div> </div> <div> <div>2020-02-26 09:58:15.305</div> <div>{ "cluster": "pmkft-prod", "nodename": "ip-10-31-136-125.us-west-2.compute.internal", "log": "E0226 17:58:11.717379 26067 pod_workers</div> </div> <div> <div>2020-02-26 09:58:12.430</div> <div>{ "cluster": "pmkft-prod", "nodename": "ip-10-31-131-6.us-west-2.compute.internal", "log": "E0226 17:58:11.808046 30992 pod_workers</div> </div> <div> <div>2020-02-26 09:58:12.261</div> <div>{ "cluster": "pmkft-prod", "nodename": "ip-10-31-136-200.us-west-2.compute.internal", "log": "E0226 17:58:09.450063 29228 pod_workers</div> </div> <div> <div>2020-02-26 09:58:00.360</div> <div>{ "cluster": "pmkft-prod", "nodename": "ip-10-31-136-125.us-west-2.compute.internal", "log": "E0226 17:57:57.716885 26067 pod_workers</div> </div> <div> <div>2020-02-26 09:57:49.825</div> <div>{ "cluster": "pmkft-prod", "nodename": "ip-10-31-131-6.us-west-2.compute.internal", "log": "E0226 17:57:49.733728 30992 pod_workers</div> </div> <div> <div>2020-02-26 09:57:45.783</div> <div>{ "cluster": "pmkft-prod", "nodename": "ip-10-31-136-125.us-west-2.compute.internal", "log": "E0226 17:57:45.722434 26067 pod_workers</div> </div> <div> <div>2020-02-26 09:57:40.541</div> <div>{ "cluster": "pmkft-prod", "nodename": "ip-10-31-136-157.us-west-2.compute.internal", "log": "E0226 17:57:39.681385 28911 pod_workers</div> </div> <div> <div>2020-02-26 09:57:37.430</div> <div>{ "cluster": "pmkft-prod", "nodename": "ip-10-31-131-6.us-west-2.compute.internal", "log": "E0226 17:57:36.733618 30992 pod_workers</div> </div> <div> <div>2020-02-26 09:57:35.326</div> <div>{ "cluster": "pmkft-prod", "nodename": "ip-10-31-136-125.us-west-2.compute.internal", "log": "E0226 17:57:31.555042 26067 pod_workers</div> </div> <div> <div>2020-02-26 09:57:29.336</div> <div>{ "cluster": "pmkft-prod", "nodename": "ip-10-31-136-157.us-west-2.compute.internal", "log": "E0226 17:57:25.681042 28911 pod_workers</div> </div> <div> <div>2020-02-26 09:57:27.430</div> <div>{ "cluster": "pmkft-prod", "nodename": "ip-10-31-131-6.us-west-2.compute.internal", "log": "E0226 17:57:25.737566 30992 pod_workers</div> </div> <div> <div>2020-02-26 09:57:15.456</div> <div>{ "cluster": "pmkft-prod", "nodename": "ip-10-31-136-157.us-west-2.compute.internal", "log": "E0226 17:57:14.689536 28911 pod_workers</div> </div> <div> <div>2020-02-26 09:57:12.430</div> <div>{ "cluster": "pmkft-prod", "nodename": "ip-10-31-131-6.us-west-2.compute.internal", "log": "E0226 17:57:11.734810 30992 pod_workers</div> </div> <div> <div>2020-02-26 09:57:05.062</div> <div>{ "cluster": "pmkft-prod", "nodename": "ip-10-31-136-157.us-west-2.compute.internal", "log": "E0226 17:57:02.681636 28911 pod_workers</div> </div> <div> <div>2020-02-26 09:56:59.793</div> <div>{ "cluster": "pmkft-prod", "nodename": "ip-10-31-131-6.us-west-2.compute.internal", "log": "E0226 17:56:59.733831 30992 pod_workers</div> </div> <div> <div>2020-02-26 09:56:58.072</div> <div>{ "cluster": "pmkft-prod", "nodename": "ip-10-31-132-169.us-west-2.compute.internal", "log": "E0226 17:56:55.758459 27606 pod_workers</div> </div> <div> <div>2020-02-26 09:56:48.946</div> <div>{ "cluster": "pmkft-prod", "nodename": "ip-10-31-136-157.us-west-2.compute.internal", "log": "E0226 17:56:47.680902 28911 pod_workers</div> </div>									

A cloud-operations SLA

By automating not just deployment but health monitoring, runbook-driven resolution of common problems, streamlined upgrades, and closed-loop alerting, Platform9 guarantees a financially backed 99.9% SLA for distributed clouds. Hitherto, 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 with distributed clouds.

Key features and benefits of the Platform9 open distributed-cloud service

Key Platform9 capabilities include:

- **Instant deployments** — Spin up a Kubernetes cluster and deploy your cloud-native apps to any distributed cloud in minutes and then scale them up or down on demand.
- **Life-cycle management** — Focus on your apps and let Platform9 handle all aspects of management from monitoring and troubleshooting to repairing.
- **Integrations** — Built-in monitoring that integrates with Slack; auto log forwarding to your aggregator; integration with your SSO provider; and much more.
- **1-click upgrades** — Upgrades and security patches are fully-automated. Choose from several upgrade options and the convenience of multi-version support.
- **Guaranteed uptime SLA** — Closed-loop automation, monitoring, alerting, and self-healing backed by our site reliability engineers and DevOps teams to ensure our financially backed 99.9% uptime SLA.
- **Expert support and solution architects** — Certified Kubernetes Administrators (CKA) and cloud experts work as an extension of your team and are available 24/7.

Business benefits our customers have achieved with this approach include:

- **Transformations to cloud native 4x faster** — Create production, enterprise-grade, cloud-native environments 4x faster with our SaaS-based deployment model, pre-built integrations, and access to experienced solution architects.
- **Increased operational productivity gains** — Guaranteed 99.9% uptime with 24/7 monitoring and self-healing, full life-cycle automation, and an on-demand 100% CKA team.
- **Maximized freedom and flexibility** — Choose to run on public clouds, on-premises, or at the edge with a cloud-agnostic SaaS control plane and pre-built infrastructure providers.

For the emerging categories of distributed and edge cloud computing, geographic distribution of infrastructure and workloads 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 management will be the *de facto* standard for distributed cloud management.

Learn more at [Platform9.com](https://platform9.com)!