Bare Metal Cloud: The Key to Getting the Most Out of Kubernetes

White paper by phoenixNAP
June, 2022

INTRODUCTION

As digital transformation propels organizations towards container adoption and into the cloud-native era, we are witnessing Kubernetes going mainstream. The rising popularity of this container orchestration platform is due to its ability to simplify application and microservice management while adding resiliency to container workloads. The further down the container adoption path an organization goes, the more pressing the question of where to run its workloads becomes. While Public Cloud environments are a suitable solution for initial deployments, choosing this as the infrastructure for at scale production Kubernetes workloads comes with a set of limitations.

This white paper offers an overview of the current state of Kubernetes adoption and the challenges of running containerized workloads in production. It also outlines the differences between running containers on Bare Metal Cloud and Public Cloud environments. Industry insights and phoenixNAP's internal knowledge demonstrate why the former is a better choice for Day 2 operations.

I - The State of Kubernetes

A recent survey involving over 3800 global participants showed that **96% of organizations are either using or evaluating Kubernetes**. Of these, 60% are using the platform in production. This not only confirms that Kubernetes is becoming ubiquitous but also that a growing number of enterprises choose it as their preferred tool for orchestrating mission-critical containers.

Another sign of Kubernetes crossing the adoption chasm is the fact that developers are increasingly overcoming its notoriously steep learning curve. According to the latest **StackOverflow Annual Developer Survey** conducted among over 80,000 developers, 72.45% said they loved Kubernetes, putting it in the top 3 favorite development tools.

> 96% of organizations are either using or evaluating Kubernetes and 60% use it in production!

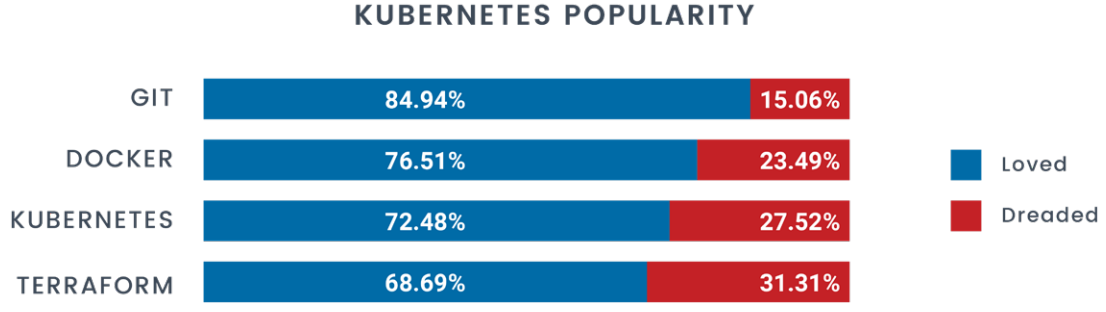

Source: 2021 Developer Survey, StackOverflow, available at:
https://insights.stackoverflow.com/survey/2021#worked-with-vs-want-to-work-with-language-worked-want-prof

II - Adopting and Running Kubernetes: Benefits and Challenges

What made Kubernetes evolve from a mere buzzword to an enterprise-grade solution is the range of benefits it brings to organizations. These include:

Facilitating Cloud Migration

Kubernetes works with just about any type of container runtime or underlying infrastructure. Organizations can use it to effectively lift and shift entire monolithic apps and microservices to hybrid environments, and optimize their code later.

Optimizing Resource Use

While defining a pod in Kubernetes, developers can specify the exact amount of CPU and RAM resources allocated to it. Based on this information, Kubernetes places the pod on a node where these resources are available, making sure there is no overcommitment or waste.

IT Cost Optimization

Autoscaling capabilities help organizations optimize their resource use in real time, lowering their cloud costs and avoiding resource overprovisioning. With workloads running indiscriminately across optimized hybrid deployments, developers can focus on value-added tasks, accelerating time-to-market.

Accelerating Software Release Timeframes

Kubernetes helps developers automate manual container management processes. It checks the health of nodes, containers, and clusters, automating both rollouts and rollbacks. Because it supports microservices natively, teams work in smaller, more productive groups, and code faster.

Apart from testifying to the platform's widespread adoption, the Kubernetes popularity chart given above also shows that almost 30% of the platform's users dread using it.

While Kubernetes does streamline certain cloud-native tasks and processes, setting up, optimizing, and managing an enterprise container environment at scale does not come without challenges.

According to **a recent report**, 38% of developers claim that working with Kubernetes leads to extreme burnout and 51% say that building cloud-native apps makes them want to find a new job.

Some of the main challenges behind the struggles to adopt and manage K8s include:

Lack of In-House Skills

Many companies still face difficulties when it comes to building, deploying, and managing their container workflows efficiently with Kubernetes. Experienced staff is limited and expensive, putting a strain on an organization's digital transformation budget. Existing teams can also lack internal alignment when selecting an appropriate K8s distribution, hampering its implementation.

IT Infrastructure Limitations

Having the ability to scale infrastructure quickly and cost-effectively is paramount when designing cloud-native applications, especially when they are stateful. On-prem and legacy infrastructure solutions tend to be poorly equipped for scale and high availability. Also, the complexity of management, absence of automation, and storage constraints all hinder container adoption.

Vendor Lock-In

To decrease the complexity of deploying and managing Kubernetes on their own, organizations often opt for managed public cloud services. However, adding one proprietary service on top of another easily becomes cost-prohibitive and gradually traps businesses in a single vendor's ecosystem. This prevents organizations from using alternative solutions or optimizing costs by adopting hybrid cloud.

Security Concerns

As Kubernetes grows in popularity, cyberattacks aiming to breach its containerized systems become more sophisticated. Consequently, security is one of the main challenges of organizations using Kubernetes. A recent survey conducted by Red Hat states that 93% of respondents experienced at least one security incident in their Kubernetes environments in the last 12 months.

Lack of Observability

To get the most out of Kubernetes, organizations must be able to visualize their environment and see how all its components interact with one another. Monitoring the status of your container infrastructure across layers of your container solution can be quite difficult.

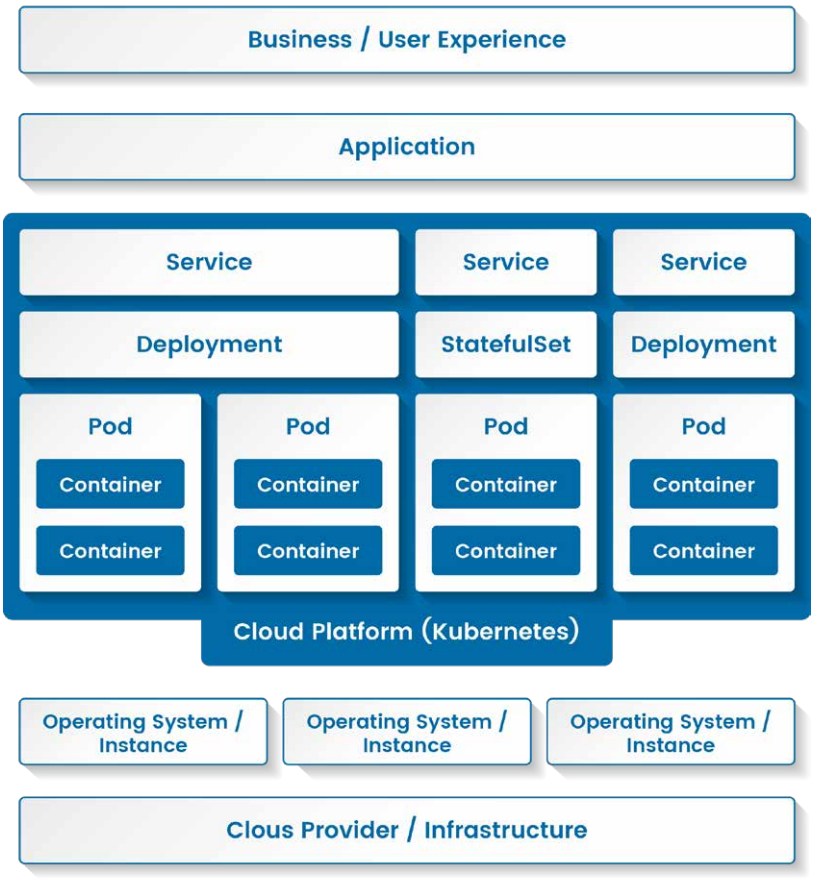

The performance of workloads depends on the infrastructure, as this is the foundation on which every Kubernetes solution rests.

III - Running Kubernetes: Bare Metal Cloud vs. Public Cloud

From physical on-prem infrastructure and managed solutions to virtualized public cloud services, choosing where to run your containers can greatly impact their performance, security, and observability.

On-prem vs. Data Center

When it comes to finding a suitable infrastructure for Kubernetes, one of the most common questions is whether to invest in setting up the environment on premises or to make use of managed data center services.

While an on-prem infrastructure offers the greatest control over your clusters, it comes with several limitations. Firstly, setting up an on-prem infrastructure optimized for Kubernetes is a daunting task that requires time, skills, and a considerable amount of upfront expenses. Also, in-house infrastructure capacity limits node scaling possibilities. This makes it impossible to geo-distribute deployments

for added resilience unless an organization has its own network of global data centers. Additionally, setting up management, security, monitoring, and infrastructure automation tools from scratch puts a significant strain on your IT staff and can take months to complete.

To avoid the complexities and costs of running Kubernetes on premises, organizations increasingly embrace managed infrastructure services. Enterprise data center service providers usually offer the following:

- High security and global availability
- Power and cooling redundancy
- Hardware scalability and customization on demand
- High bandwidth and global connectivity
- SLA-guaranteed uptime
- Managed services and 24/7 support
- Flexible billing options

Once they opt for managed services, organizations must decide what type of service to use. Depending on their workloads and the toolsets their developers use to manage them, organizations usually narrow down their infrastructure choices to two substrates. They either deploy their containers on popular, turnkey public cloud solutions or choose bare metal for more flexibility and customization options.

A **recent survey** conducted by Canonical showed that among 1300 respondents 50.8% run their Kubernetes clusters in one of the most popular public cloud hyperscaler environments while 19.4% choose bare metal.

Virtual vs. Physical

Most public cloud hyperscalers offer managed Kubernetes services. Such solutions are attractive to businesses mostly because of their ease of use, flexibility, and the availability of resources. However, running containers in virtualized, multi-tenant environments can affect workload performance and lead to vendor lock-in, tying organizations to cost-prohibitive services.

On the other hand, managed service providers offer organizations access to solutions far more versatile than traditional bare metal. One such flexible solution is bare metal cloud.

Bare metal cloud delivers dedicated resources that are deployed and managed in a cloud-like manner. This single-tenant, vendor agnostic environment presents a viable choice for powering Kubernetes workloads as it automates server provisioning and integrates with Infrastructure-as-

Code (IaC) tools. As such, it allows developers to use cloud-native interfaces to quickly boot up their physical infrastructure across locations and consume it on a pay-per-use model.

The main differences between bare metal cloud and public cloud include the following:

FACTOR	Bare Metal Cloud	Public Cloud
Environment	Physical	Virtualized
Hardware Access	All resources belong to a single user	Multiple tenants share resources
Configuration	Transparent, customizable	Non-transparent, no customization
Performance	Consistently high	Inconsistent
Deployment Time	Minutes	Minutes
Scalability	Near-instant	Near-instant
Security	Single tenancy improves security	Multi tenancy may compromise security
Billing	Hourly or monthly	Hourly or monthly
Vendor Lock-In	Freedom to leverage different technologies and environments	Ties users to proprietary technologies and environments

Taking a deeper look into the comparison, we can see that it is necessary to add containers on top of the infrastructure and consider all the layers of both stacks separately.

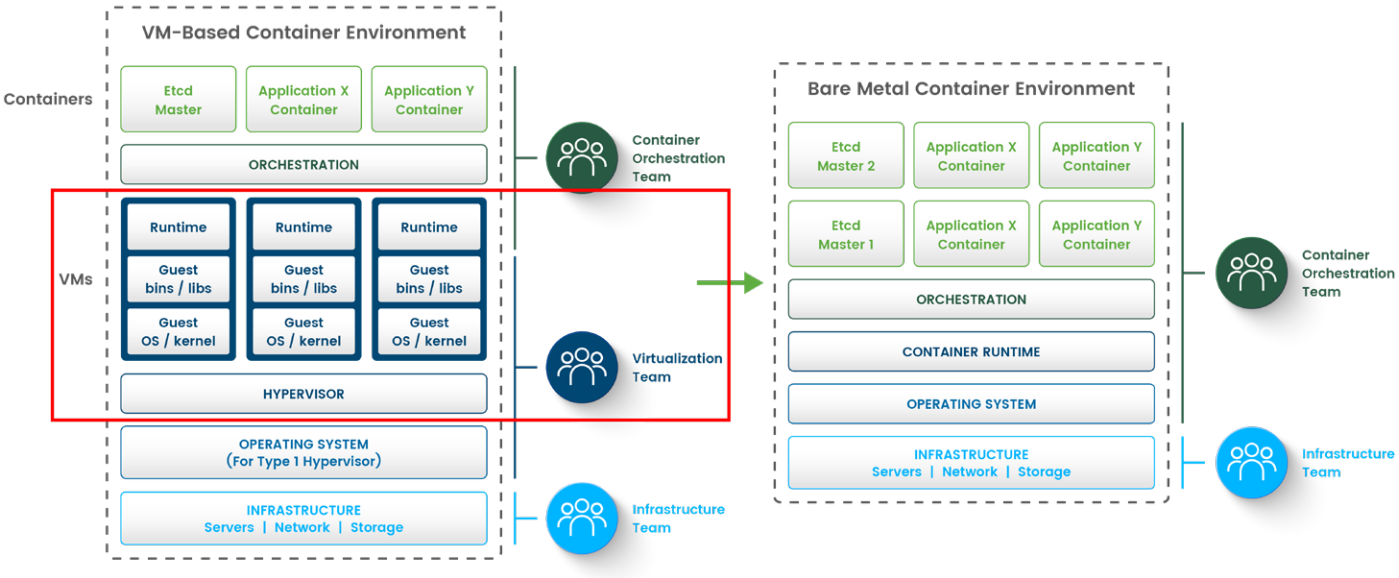

Comparison of a virtual environment and a bare metal stack running containers, showing the virtualization overhead.

The Weight of the Hypervisor

The first thing that sets the two stacks apart is the hypervisor layer. Namely, a virtual environment leaves a significant footprint on the infrastructure's CPU, RAM, and storage. Before an app starts running in a public cloud VM, the hypervisor uses the server's resources to enable virtualization and run its guest OS, impacting the overall workload performance.

This decreases container density since a server that already hosts multiple VMs can hold fewer containers compared to a bare metal server. Paired with the "noisy neighbor" issue that affects resource availability, virtual environments put a strain on the overall performance of containers and their orchestration platforms.

"Noisy neighbor" is a term used to describe a co-tenant that monopolizes bandwidth, storage, CPU, and other shared resources, degrading cloud performance for other users.

VMs are designed to allocate resources to the environment at start-up regardless of whether they are used or not. Because of this, virtual environments additionally contribute to poor resource utilization.

Another burden the additional orchestration layer brings is the need for more staff. In a VM Kubernetes deployment, there typically needs to be a separate team monitoring and managing the virtualization processes and tasks. This not only contributes to the lack of internal alignment and communication complexities between teams but can also lead to the additional expense of hiring expert staff.

Finally, virtualization can take its toll on an organization's IT budget through licensing costs. Certain hypervisors come with paid license plans for virtualization software. Depending on the guest OS used in each VM, licensing costs may go even higher. Unable to move or migrate workloads to more affordable environments due to vendor lock-in, organizations find themselves burdened with constant service price increases.

The Power of Bare Metal

While the advantages of running Kubernetes on bare metal greatly depend on the use case, this dedicated infrastructure exploits the benefits of a container orchestration engine. Firstly, direct access to CPU, RAM, and storage resources enables **lower latency for container processes and maximized resource utilization.** Such consistency in performance is paramount for most production, real-time, or data-intensive workloads.

Typically, container performance on bare metal is about 20-30% better than on a VM.

With its built-in workload acceleration features, bare metal enables organizations to leverage the latest hardware. Its decreased latency and compatibility with next-gen technologies make this infrastructure suitable for **multi-cluster deployments at the edge.** With 5G gaining global momentum, bare metal is often the default infrastructure because it fully supports the new network's capabilities and the modern apps using it.

As regards connectivity, running Kubernetes on bare metal means having a single networking stack instead of wasting resources on moving data between physical and virtual machines. This **simplifies networking setup and maintenance** and boosts overall performance.

A recent study showed that Kubernetes networking on bare metal had **3x lower latency** compared to a VM.

From a data protection perspective, single-tenant environments improve container **security** by reducing the attack surface to a minimum. With bare metal, security pipelines encrypt every node in every cluster.

Kubernetes on bare metal brings simplicity to infrastructure management and access. Administrators can rightsize their environment to fit the needs of their workloads with **optimal efficiency.**

Finally, more available resources per server means **lower TCO** and optimized IT spend. Abstraction of the virtualization layer removes vendor lock-in and brings **lower costs** for licensing and maintenance.

Bare Metal Cloud: Built for Kubernetes

Bare metal cloud is a single-tenant, non-virtualized environment consumable on a cloud model. As such, it allows organizations to utilize the full processing potential of a server's physical hardware while managing and provisioning their infrastructure as code.

phoenixNAP's Bare Metal Cloud provides organizations with access to pre-configured bare metal instances deployable in minutes across multiple locations in the US, Europe, and Asia.

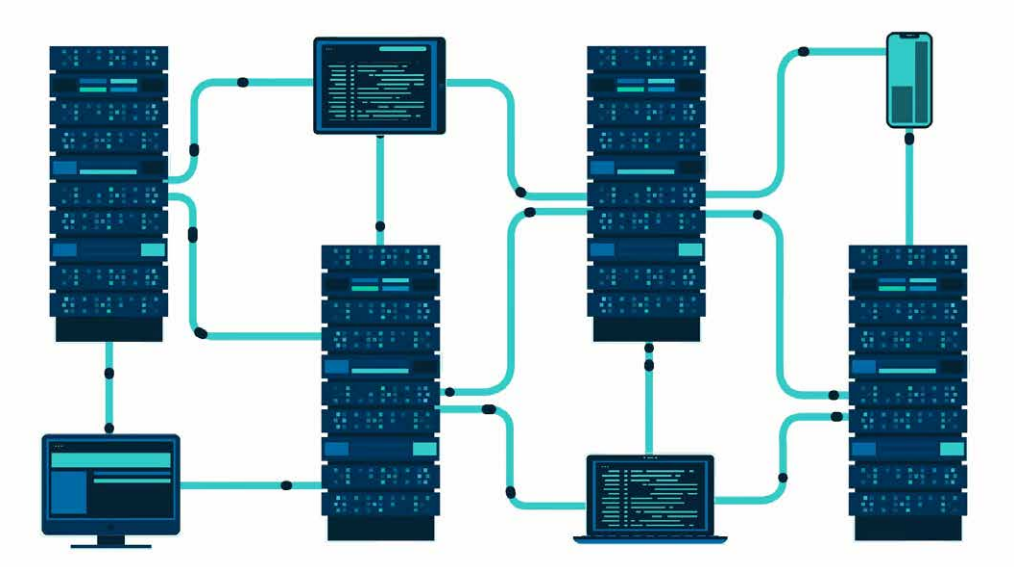

The platform's features include:

- Global and edge bare metal server provisioning in minutes via API, CLI, or WebUI
- Workload-optimized, dedicated server instances powered by the latest hardware technologies
- Automatic installation of a range of OSs (CentOS, Ubuntu, Windows, ESXi, Proxmox, etc.)
- Management via popular Infrastructure-as-Code (IaC) tools (Terraform, Ansible, Pulumi)
- Hybrid cloud connectivity with AWS and Google Cloud on-ramps
- Easy access to petabytes of S3-compatible Object Storage
- GitHub repos, actions, and SDKs (including K8s Controller and Docker Machine plugin)
- Confidential computing capabilities with Intel® SGX
- Up to 50 Gbps bonded network with 20 Gbps free DDoS protection and 15 TB free bandwidth
- SUSE Rancher integration for one-click deployment of a production-ready K8s environment
- Pay-per-use billing and bandwidth models with discounts for reservations

Bare Metal Cloud enables fast deployment of powerful and secure dedicated servers ready for Kubernetes, making it the ideal environment for anything from sandboxing to mission-critical container workloads.

The Importance of the Use Case

Having addressed the pros and cons of deploying Kubernetes in physical and virtual environments, it is necessary to highlight that the choice of infrastructure depends greatly on the use case.

Public cloud has been the go-to solution for container clusters for the last decade or so for a reason. It is a simple, yet efficient way of running Kubernetes and using golden images to spin up pre-configured,

general-use VMs ready for scaling on a single node. As such, it is a suitable environment for the following use cases:

- Running container workloads that are not latency-sensitive.
- Deploying staging environments and Kubernetes labs.
- Running clusters that cannot tolerate multiple node failures.
- When it is necessary to rely on existing services and cloud provider interfaces (CPIs).
- When it is impossible to cost-effectively avoid vendor lock-in of the current infrastructure.

On the other hand, high-availability production containers deployed at scale greatly benefit from running on a bare metal infrastructure.

As far as running Kubernetes on a dedicated infrastructure available as a service is concerned, providers like phoenixNAP offer API-driven bare metal solutions designed for cloud-native workloads.

IV - How to Deploy Enterprise K8s Clusters on Bare Metal in 5 Minutes

SUSE Rancher integration on phoenixNAP's Bare Metal Cloud was built with the goal of helping developers overcome the complexities of manually deploying Kubernetes clusters for Day 2 operations. The solution lets organizations easily deploy a geo-distributed, production-ready Kubernetes environment through a single POST request via the platform's API, or a couple of clicks in the portal.

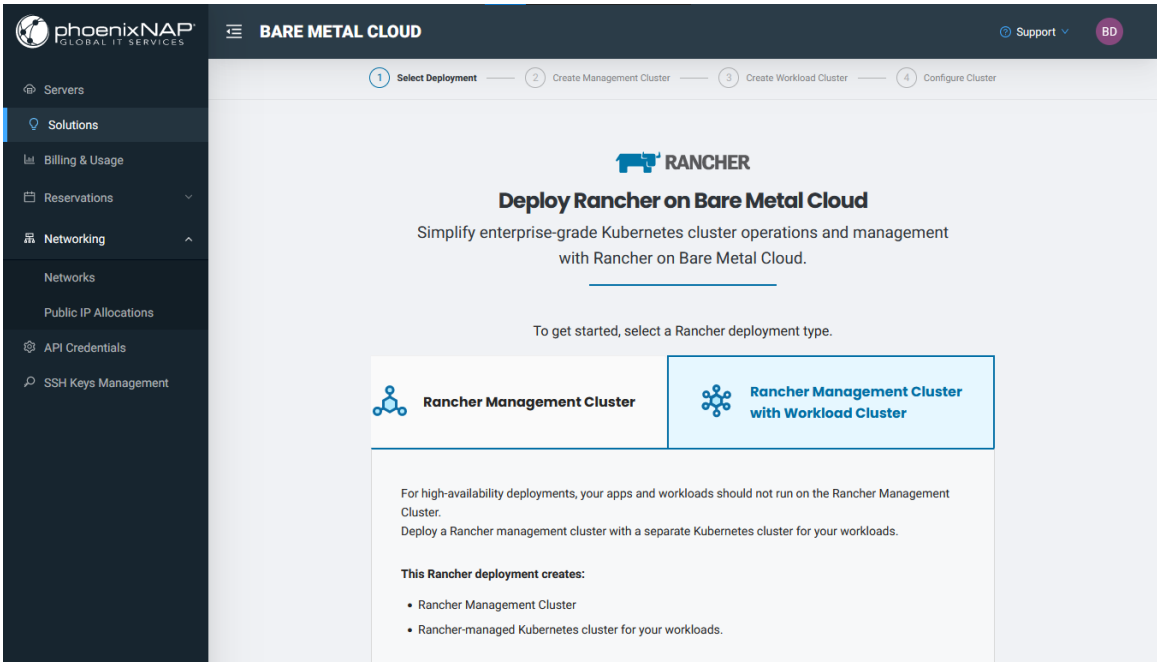

Built on open-source technologies, the Rancher solution on bare metal cloud is completely cloud-agnostic. It supports any certified Kubernetes distribution including RKE for on-prem, EKS, AKS, or GKE for cloud, and K3S for edge workloads.

It takes about 5 minutes to deploy a 3-node Kubernetes workload cluster alongside a HA management cluster on phoenixNAP's Bare Metal Cloud. To do this it is necessary to:

1. Create a Bare Metal Cloud account and log into the portal.
2. Select the locations for your management and workload clusters.
3. Name your clusters and add nodes to them.
4. Choose from pre-configured node instances.
5. Add previously generated SSH keys to your nodes.
6. Deploy your enterprise K8s environment with the click of a button.

The intuitive portal also allows for advanced cluster configuration. Users can set access tokens, add certificates, set TLS settings and node taints, or add an additional hostname and schedule backup snapshots. Post-deployment, both downstream and upstream clusters and nodes can be easily overviewed, managed, and tweaked via phoenixNAP's portal or directly from the Rancher UI.

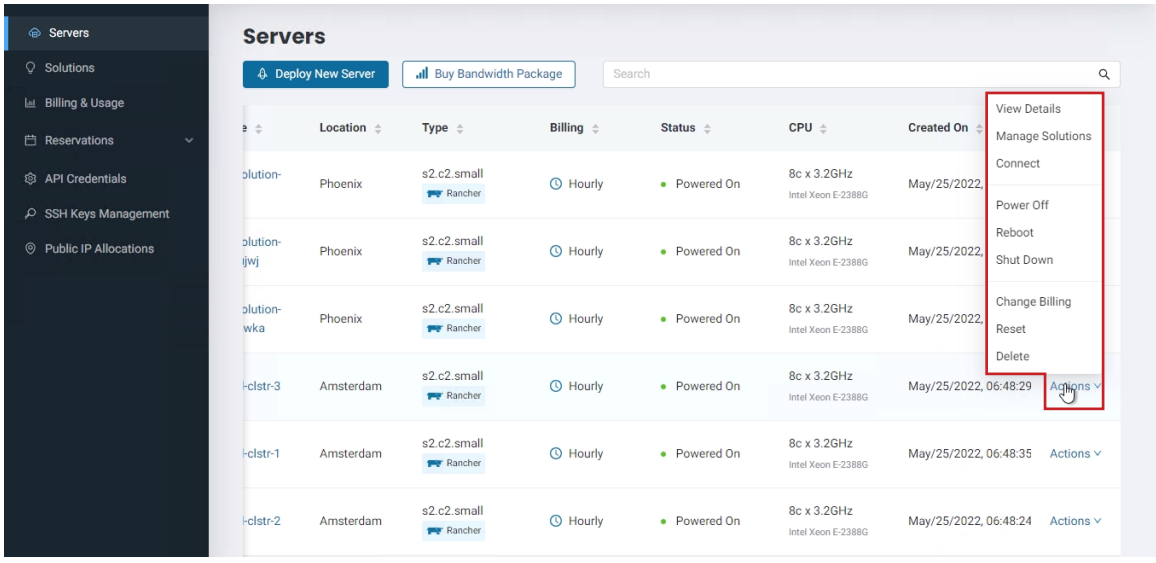

Solutions like Bare Metal Cloud are perfect for organizations adopting Kubernetes or looking to abstract the complexities of building their container service environments from scratch. From hourly-billed,

single-node clusters for testing purposes to enterprise-level, HA clusters across different locations, phoenixNAP's Bare Metal Cloud provides the hardware and software support fast-paced businesses require.

You can deploy phoenixNAP's SUSE Rancher-managed K8s cluster on bare metal cloud for as low as $0.08/hour!

CONCLUSION

Setting up your Kubernetes infrastructure requires serious planning and there is no one-size-fits-all solution. Evidence suggests that containers perform best on bare metal.

Turnkey solutions such as phoenixNAP's Bare Metal Cloud offer an easy way to accelerate time-to-market and save your already strained developers hours of environment setup work.

Resources:

1. CNCF, Annual Survey 2021
2. StackOverflow Annual Developer Survey 2021
3. Red Hat, Kubernetes adoption, security, and market trends report 2022
4. Lumen, Kubernetes on Bare Metal: When Low Network Latency is Key
5. D2IQ, New Study Uncovers Importance and Challenges of Enterprise Kubernetes
6. Canonical, Kubernetes and cloud native operations report 2022

About phoenixNAP

phoenixNAP is a full-service IaaS provider delivering programmable, OpEx-friendly infrastructure solutions from strategic edge locations worldwide. Focused on innovation, cybersecurity, and compliance-readiness, phoenixNAP collaborates with technology industry leaders to make enterprise-grade technologies available on an OpEx-based model. Its cloud, dedicated servers, availability, HaaS, and colocation solutions can be customized to meet any business's requirements.

Contact phoenixNAP at: **sales@phoenixnap.com**
or **855.330.1508 | www.phoenixnap.com**