veeam

5 Best Practices for Kubernetes Backup

Addressing Data Protection and Management Needs

Contents

Executive Summary

Veeam®, the #1 global leader in data protection & ransomware recovery, is focused on helping customers backup, secure and manage all their cloud, virtual and physical workloads. Veeam's solution now extends to containerized workloads with Veeam Kasten *for Kubernetes*. The Veeam Data Platform helps you achieve the best practices described in this white paper to address cloud native data protection needs for enterprises.

Applications using cloud native and microservice-based architectures have quickly gained traction in the enterprise — an evolution from earlier monolithic and virtualized approaches. Kubernetes has emerged as the dominant, indeed de-facto, container orchestration platform. In parallel with the adoption of containers, newer organizational approaches (usually termed DevOps or ITOps) are being rapidly adopted where software development and IT operations roles are being "combined" to deliver higher agility while maintaining quality. In this environment, developers have more latitude in their tool selection and exert a greater influence on operations.

Data, the most important asset in any enterprise, has seen its value increase further with the widespread adoption of Kubernetes given the artificial intelligence and machine learning stacks being deployed on it. However, this information and the associated software assets are subjected to hacking attacks, often to devastating effect. Privacy concerns have also led to stringent regulations. Taken together, these have made data protection a "front and center" concern for IT teams spanning four distinct operational use cases: backup and restore, disaster recovery, application mobility and ransomware protection.

Operations teams, that manage infrastructure and applications, do not operate in a static world either. The cloud native environments they are now managing are dynamic and complex. The teams need to support a varying mix of traditional, virtualized, and containerized workloads based on the use of various data services, including relational and NoSQL databases. Their deployments can be spread across on-premises, and multiple cloud environments and often rely upon a range of storage solutions from multiple vendors. Any data management solution deployed in these environments will need to balance the needs of operators and developers by being both operations-focused and developer friendly.

Kubernetes further adds to the transformation of the IT landscape as it is fundamentally different from platforms based on earlier technologies. Accordingly, it requires a different approach to backup, one that we describe as a Kubernetes native backup. There are seven reasons why Kubernetes native backup is critical.

7 Reasons Why Kubernetes Native Backup is Critical:

1

Kubernetes Deployment Patterns

The Kubernetes platform is fundamentally different from earlier compute infrastructures. There is no mapping of applications to servers or VMs. A backup solution needs to understand this Kubernetes native architectural pattern and be able to deal with continuous change.

2

DevOps and "Shift Left"

High-velocity application development and deployment cycles are the norm in Kubernetes environments. Consequently, this requires that backup solutions be application-centric and not infrastructure focused.

3

Kubernetes Operator Challenges

Operations requires ease of use to accelerate an IT team's production journey to Kubernetes deployments. Backup solution with CLI access and a clean API along with a powerful yet easy to use dashboard is critical.

4

Application Scale

Kubernetes-based microservices comprise hundreds of discrete components with independent lifecycles visible only to Kubernetes. A Kubernetes native approach to backups, keeping applications as the unit of atomicity for consistent operations is an imperative.

5

Protection Gaps

Relying solely on high availability or replication capabilities can lead to data corruption or catastrophic data loss. A backup solution that works transparently against a wide range of Kubernetes application stacks and deployment methods is required.

6

Security

Kubernetes security features deny access to internal application components and their associated data services from not just outside the cluster but also to other untrusted applications. A well-architected Kubernetes native backup solution that can embed itself into the Kubernetes control plane ensures consistent security operations.

7

Ecosystem Integration

Polyglot persistence, where multiple data services are used within the same application, has coincided with the growth of Kubernetes. A backup solution with workload knowledge to select the capture primitives best suited to the application's requirements as well as interoperation with the rest of the cloud native infrastructure ecosystem is key.

> "Throughout our evaluation, Enterprise Strategy Group determined that [Veeam Kasten] can deliver the data protection and recovery capabilities required by Kubernetes applications. We examined this by comparing how traditional solutions and [Veeam Kasten] address four major use cases: backup and restore, disaster recovery, application mobility, and ransomware protection. [Veeam Kasten] is better suited to these use cases when considering stateful and containerized applications running on Kubernetes, as the platform is application-aware and application-consistent. Since traditional solutions are VM-centric, the data protection offered to Kubernetes applications can be ineffective, inefficient, and incomplete."

Technical White Paper: Dispelling the Myths of Kubernetes Data Protection
ESG, April 2024

Based on our customers' experience and considering the key factors that warrant particular attention above we have arrived at following **five best practices for Kubernetes backups:**

Architecture

The platform used to protect Kubernetes applications needs to automatically discover all the application components running on your cluster and treat the application as the unit of atomicity. The application must include the state that spans across storage volumes, databases (NoSQL/Relational) as well as configuration data included in Kubernetes objects such as configmaps and secrets.

Recoverability

The data protection platforms must allow you to restore the application components you want, and where you want them. You should also have the granular to restore only an application subset such as the data volume. The approach must make restoring simple and powerful by allowing you to select the appropriate point of time copy of the application.

Operations

It is important to ensure that a Kubernetes native backup platform can be used at scale, provide operations teams with the workflow capabilities they require, and meets compliance and monitoring requirements. Operators should be able to give self-service capabilities to application developers without requiring application code or deployment changes.

Security

Controls around identity and access management and role-based access control (RBAC) must be implemented. RBAC allows different personas in an operations team to adopt a least-privilege approach to common tasks such as monitoring. Encryption at rest and in transit must always be implemented to ensure that data is secure whenever it has left the compute environment.

Mobility

Living in a multi-hybrid-cloud world, a cloud native data management platform needs to be able to be flexible in the support for multiple distributions and offer capabilities that allow for the portability of workloads and applications across all these diverse environments. Kubernetes application mobility and migration are the collective capabilities required across multiple use cases including application restoration, cloning, and container migration.

Chapter 1. Introduction

1.1 Modern IT Environments

The Cloud Native Computing Foundation in their 2023 Annual Survey found that cloud native is the undisputed infrastructure of global technology[1]. IT teams have worked diligently to both refactor existing applications and adopt cloud native architectures as the default for new development. The 2023 Annual Survey found that 66% of providers and consumers were using Kubernetes in production and 18% were evaluating it. This finding shows that Kubernetes continues as the de-facto container orchestration platform with total of 84% either in production or evaluation which, is a 4% increase over the prior year's findings.

84%

of providers and consumers were using Kubernetes in production or evaluation in 2023

According to a recent survey conducted by ESG Research[2], more than 47% of executive decision makers from various industries currently utilize containers in their production systems. Furthermore, an additional 35% of companies have expressed their intent to adopt container technology within the next 12 months. Looking ahead to 2024, it is projected that 82% of organizations will be employing containers. Considering the findings of widespread adoption of Kubernetes container technology, it is anticipated that approximately two-thirds of the containers in production will be based on Kubernetes.

82%

of organizations will be employing containers in 2024

Along with this technology evolution, a change in the organizational and process approaches being adopted by IT teams has also been seen. These are typically described as DevOps or ITOps, and reflect a newer, more agile manner of working with developers exercising greater freedom in their selection of tools and playing a larger role in operational matters. As the mindset has evolved, so have the measures that the teams focus on — with greater attention being paid to metrics such as code release time, deployment frequency, time to restore and change fail rate.

[1] CNCF 2023 Annual Survey, Cloud Native Computing Foundation, April 9, 2024

[2] "Measuring the Current State and Momentum in the Enterprise Market for Kubernetes Protection" Enterprise Strategy Group, Christophe Bertrand, April 2023

1.2 Information is a valuable asset

Not only have IT teams and their organizational philosophy evolved, but IT has become an even greater enabler of competitive advantage than in the past. Entirely new, and very successful, businesses have been spawned that rely on the effective and novel use of IT. This is driven by improved access to scalable computing, networking and storage resources, the growth of and access to data, and the adoption of advanced machine learning and artificial intelligence techniques. The unprecedented growth of data, a firm's ability to extract value from that data, relentless attacks on IT systems by state and non-state actors, and the need to comply with stringent privacy and data protection mandates has led to a clear recognition of corporate information as critical assets that must be protected. All underlying data assets and the associated code need to be protected against both malicious and accidental loss or corruption.

1.3 What the Operations Team needs to address

Typical IT environments in a modern enterprise are not static entities with fixed, standardized single-vendor deployments. The push to continually improve the cost structure and efficiency by adopting new technologies, coupled with corporate events such as mergers and acquisitions, and the need to comply with an ever-evolving regulatory compliance landscape result in a highly dynamic situation that operations teams need to not just address but embrace as the status quo.

Any selected management solution will need to work with:

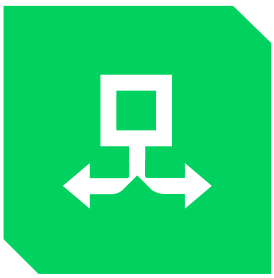

Diversity

Scale

Rate of change

1.3.1 Multi-workloads

Application architectures have undergone a rapid shift over the last few years and have gone from a monolithic design running on bare metal to virtualized applications running in hypervisors and now to containerized applications based on microservice architectures and often running itself on a virtualized infrastructure. Yet, for most enterprises with IT operating at significant scale, it is not reasonable to expect all applications to be on a single point on the evolution spectrum. Instead, the most common observed pattern is to find applications occupying the entire spectrum.

The goal

Should be to deliver a data protection solution that can work against this diversity, but still support refactoring and new development, delivering the benefits derived from cloud native applications, providing an easy transition path for legacy applications, and reducing costs incurred to manage the diverse environment.

1.3.2 Multi-environments

As with applications, the deployment environments one finds are often a mix of on-premises and cloud. It is not uncommon to find enterprises taking a hybrid and multi-cloud approach to application deployment. When using managed Kubernetes or deploying Kubernetes workloads on different clouds, the Kubernetes distributions are also typically different in terms of proprietary extensions and supported features.

1.3.3 Multi-data Services

When it comes to data services underpinning applications, they too typically demonstrate significant variety. Developers are selecting data services that best meet the needs of the task at hand and often multiple data services within the same application. This is resulting in a polyglot mix drawing from Software as a Service (SaaS) offerings, managed services, relational databases, NoSQL systems, message queues and more.

1.3.4 Multi-storage Vendors

The storage infrastructure layer is yet another dimension that shows significant variety. Not only are teams in hybrid environments using storage on-premises and from cloud vendors but depending on their needs, they use different types of storage and potentially within a single Kubernetes cluster. Further, even on-prem, the storage tends to be sourced from different vendors (Dell EMC, HP, etc.), each bringing their own management tools and frameworks.

Chapter 2. Kubernetes Native Backup is Critical

Backup is a long-established discipline with multiple solutions serving the needs for users large and small. Yet, when it comes to backing up and protecting Kubernetes based applications, there are certain reasons that employing a Kubernetes native backup solution is critical. The ebook "7 Critical Reasons for Kubernetes Native Backup" covers this aspect in detail and below is a summary for context.

2.1 Kubernetes Deployment Patterns

The Kubernetes platform is fundamentally different from earlier compute infrastructures. There is no mapping of applications to servers or VMs — Kubernetes manages the distribution of application components across servers potentially collocating applications on servers. Traditional backup systems are challenged to cleanly capture a given application's state.

2.2 DevOps and "Shift Left"

The DevOps philosophy adopted in parallel with Kubernetes cedes control over both infrastructure and deployments to the developer (known as "shift left"). Backup systems must not only integrate with the CI/CD tools the developers use, but they must also automatically detect and protect applications coming online. They should do this in a manner transparent to the developers and employ Kubernetes native APIs that the developers are familiar with. Consequently, this requires that backup solutions be application-centric and not infrastructure-focused.

2.3 Kubernetes Operator Challenges

Teams moving from a vSphere or Linux background to supporting Kubernetes will benefit from a platform that is deeply integrated with Kubernetes and masks its inherent complexity. Ignoring the multitude of resources is not an option because limiting the backup to just infrastructure disks and volumes will result in error-prone recovery and extended recovery times.

2.4 Application Scale

In the container paradigm, a single application that comprised of just a few VMs may now correspond to hundreds of distinct Kubernetes resources. When considered across all applications in a cluster, this can represent an overwhelming number of components to manage without a Kubernetes native data protection platform. Further, Kubernetes and cloud native applications are designed to scale in response to load. The backup platform must be able to respond effectively to this application scaling across the multi-cluster environments typically being used.

2.5 Protection Gaps

Whether running in the cloud or on-premises, the underlying storage is not failureproof: even AWS's battle-hardened EBS advertises a non-zero failure rate and on-prem volume snapshots may not be resistant to hardware failures, and deletion of a volume usually results in simultaneous and automatic deletion of all related snapshots.

2.6 Security

Kubernetes offers several security features, and to avoid compromising their effectiveness, it is critical that a backup solution be Kubernetes native and embed within the Kubernetes control plane. Also, with developers taking on more of the infrastructure responsibilities ("ITOps" model), it is important to be able to provide fine-grained, role-based, and scoped access using the same roles and tools used by Kubernetes instead of succumbing to the use of additional role management systems and associated increased complexity. Further, to work well with Kubernetes' approach of delegating encryption to storage and backup platforms, the backup system needs to understand Kubernetes certificate management, work with storage-integrated Key Management Systems (KMSs), and support Customer Managed Encryption Keys (CMEKs) through the Kubernetes Secrets interface.

2.7 Ecosystem Integration

The use of polyglot persistence prevalent within Kubernetes environments requires the use of a backup platform that can derive the relationships between the various data services using the Kubernetes metadata. Such a backup solution can then use these relationships and its understanding of the workloads to capture a consistent copy of the entire application stack. The backup solution also needs to fit well with the rest of the Kubernetes cloud native tools set that the operations team uses e.g., for monitoring, alerting, access control (Kubernetes APIs), logging and auditing.

Chapter 3. Data Protection Use Cases

Having touched on the reasons why a Kubernetes native backup is critical and before we dive into Kubernetes backup best practices, let's briefly discuss the key use cases associated with backup.

3.1 Backup and Restore

Naturally, the first use cas e when one thinks about bac kup is protecting against accidental or malicious data loss or corruption. This involves regularly storing copies of the relevant information so that, in case of need, one can restore from an appropriate copy. Traditionally, IT teams have adopted an approach of performing a mix of full backups (perhaps weekly) interspersed with more frequent partial ones. Besides caring that a backup was successfully created, other factors of interest include how long the backup took (this may affect the application's performance while the backup is in progress) and how much storage the backup consumes (solutions with effective deduplication and compression result in lower storage costs). In modern compute environments, the backups can be more frequent with select and stateful data protection set to a near-continuous cadence. In addition to the protection aspect, backups, or 'snapshots', can be used for testing and development purposes i.e., one may take a snapshot of production data and restore it to a separate environment to support development or performance testing needs. Not all approaches to backup are created equal. When considering implementing a backup strategy, it is useful to also bear in mind thenotion of consistency levels. The different approaches that should be considered includestorage-centric, storage-centric but with data service hooks, data services-centric andapplication-centric. More details on these consistency levels can be found in the "Flavors of Data Management in Kubernetes" article[3].

[3] Source: "Flavors of Data Protection in Kubernetes", TFIR, Guarav Rishi, October 28, 2019

3.2 Disaster Recovery

Disaster recovery is the second use case and refers to scenarios where there is a significant-enough local failure (e.g. due to a fire or a flood, or extended power or networking outages) that operations need to be continued at a different location altogether. The entire application stack as well as the associated data need to be made available for use at an alternate location. Besides cost, the key considerations for disaster recovery are usually measured in terms of Recovery Point Objective (RPO), how fresh the data is and how much data loss can be tolerated) and Recovery Time Objective (RTO), how long it takes to get the recovery environment operational). Teams need to take these considerations into account when deciding what kind of data protection platform will be appropriate for their business requirements.

3.3 Application Mobility

Application mobility is the ability to migrate an application from one environment to another, including across clusters, regions, Kubernetes distributions, from on-premises to cloud, or even from one cloud to another. Such moves may be necessitated by changes in an organization's preferred technology platform, cost optimization, or new business requirements. It may also be desirable to restore an application to another platform. This use case also includes the requirements originating from the need to upgrade from one stack to another (e.g., upgrading from OpenShift 3.x to OpenShift 4.x). Addressing the mobility use case requires strong support to handle the many differences between the source and target environments — that the destination may not have the same storage system is just one example. Teams choosing solutions for this use case are mindful of the ability to provide data and application transformations from source to the target when underlying Kubernetes specifications, computer and storage infrastructure can change (e.g., the platform provides policy-based operations to transparently modify application specifications on the fly).

3.4 Ransomware Protection

Ransomware protection is the capability to protect data and applications in the event of a cybersecurity attack which could lead to ransomware demands. In the most basic case, regular backups of applications, critical data and configurations are essential. These practices ensure a quick recovery in case of an attack. Employing the principle of least privilege helps restrict unnecessary access, which helps to minimize the potential impact of a breach. Continuous monitoring of system behavior aids in early anomaly detection, while timely patching and updating of both Kubernetes clusters and protection systems bolsters overall resilience. Implementing network segmentation can contain the spread of a ransomware intrusion within an environment. Additionally, educating personnel on security best practices and embracing practices such as consistency and observability in DevOps can enhance their ability to recognize and respond to potential threats quickly and effectively.

Chapter 4. Best Practices

When we think about backups in a traditional sense, we would expect that in a Kubernetes environment, we would want to be able to backup and restore containers, but this is not as important as having immutable container backups. Protecting the entire application state is the actual goal of a backup in a Kubernetes environment. The goal can be envisioned as creating a recipe that describes in exquisite detail the components and services the application needs to run and take a snapshot of the persistent data and storing it all together. Ideally, this takes place in a fault domain that is different than the Kubernetes cluster. In this case, a restore would involve following the recipe to restore the persistent data and then deploying and configuring the application and associated services to resume normal operations.

As traditional enterprises start to adopt containers in production and at scale, traditional monolithic applications are being "containerized" or migrated to give them a more manageable deployment mechanism. The transition of these traditional applications brings in new considerations around how these workloads are protected and recovered in the event of disaster. When we think about these applications, they traditionally have contained stateful configurations that need to be recovered. Stateful data does not only include the data being stored by the application (e.g., tables in a database). It also includes data related to application configuration (secrets, TLS certificates, etc.) that in the event of a crash, can't be recovered by simply re-deploying the application. These requirements bring in unique challenges for backing up and restoring workloads. When container platforms were primarily used with stateless workloads, applications could typically be spun up quickly and then destroyed and re-deployed as needed. As the adoption of containers is now encompassing stateful as well as stateless applications, the need for a robust backup solution has never been greater.

While most attention is directed to protecting business critical data, ensuring that data is captured with the correct consistency level is as important. From a technical perspective, considerations of crash consistent and application-aware backups are still very much at the forefront. For most workloads, a crash-consistent backup or snapshot is enough, but for workloads that have regular changes like databases, we need to make sure that we can take a consistent point-in-time backup to avoid data corruption, even in the containerized world.

Finally, with the adoption of Kubernetes focused on scale and ease of deployment, automation is a critical component for any management system deployed in the container environment. When looking at backup solutions, traditional backup products have struggled to work alongside or even integrate with the true dynamic scale of Kubernetes infrastructure. Having a backup solution that not only integrates and leverages the Kubernetes APIs but can also extend these APIs and provide greater automation integration is key. Deploying a backup platform as a native containerized system that runs within Kubernetes would the ideal option, rather than having a traditional "backup server" running in a different environment that needs to be separately managed.

Throughout this section, we will discuss the best practices and recommendations for implementing a Kubernetes-aware backup solution and the requirements for a successful backup strategy while also making it simple and flexible to adapt to this rapidly evolving cloud native ecosystem.

4.1 Architecture

When implementing a backup strategy to protect Kubernetes workloads, a deep understanding of how Kubernetes works is critical. The purpose of this whitepaper is not to describe the Kubernetes architecture in depth, but to better understand how a backup strategy should be implemented. A few components need to be discussed.

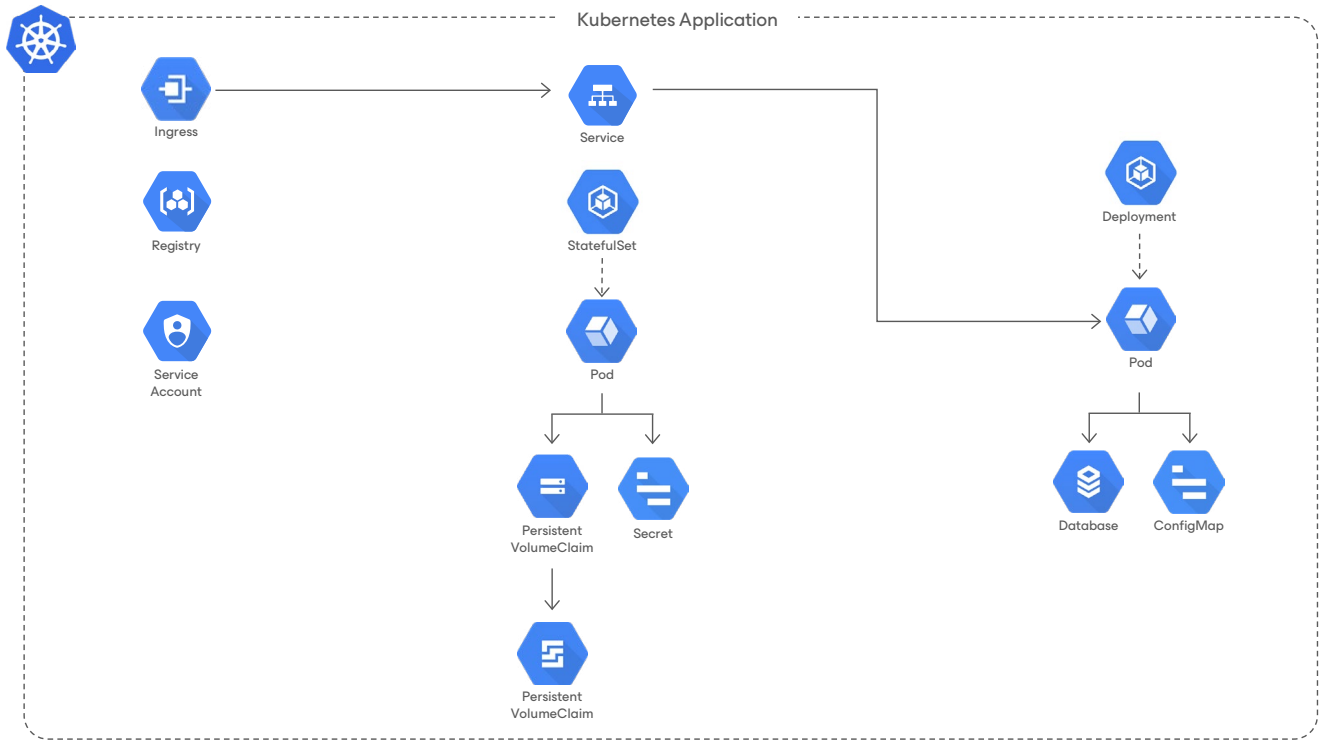

In the diagram above, we see an example of a typical Kubernetes application. It is made up of pods, services, certificates, secrets, persistent volumes, and other components. On average, we observe production applications to be composed of hundreds of these components. With these considerations in mind, it is important to find the correct solutions to be able to not just protect and restore data, but also be able to do the same with all these internal components and at scale.

Once we deploy a backup platform into Kubernetes, the solution can then automatically interface with the Kubernetes control plane via the API server. This integration can be used to not just discover the Kubernetes applications running on the cluster but also integrate with the underlying compute, network and storage infrastructure.

As a first step, the integration is used to discover the relationship between storage and applications and then determine the best (efficient, performant, consistent) way to capture the application data stored on persistent volumes along with the related application resources. The next consideration is the backup data location including within the storage system for fast recovery or, when running on the major cloud providers, depending on durable snapshots. For most cases though, backup data would be stored data in an object storage system in a different fault domain that could extend all the way to geo-replication for disaster recovery.

When it comes to storage integrations with Kubernetes, there are several key areas that need to be considered. Storage in Kubernetes is represented as persistent volumes that are made available for use to the application containers. Apart from application configuration, this is the key business data that needs to be protected. Another consideration is where to keep that data.

Is it going to be kept on local block storage?

If Kubernetes is running in an on-premises environment or going to be kept off site, use an object storage platform like Amazon S3 or Microsoft Azure blob storage. Retaining flexibility, choice, and ease-of-use in the selection of secondary storage for backup data will be a critical component of implementing a successful data protection strategy.

4.2 Recoverability

Recovery is not as simple as recreating Kubernetes objects and storage volumes. Given the number of components and Kubernetes' complexity, an execution plan needs to be created that first verifies cluster dependencies, creates new Kubernetes views of data that will get restored, and determines the compute infrastructure and Kubernetes cluster where the recovery needs to be initiated (e.g., a cross-availability zone recovery). Once the recovery execution plan is in place, the backup data sources (object storage, snapshots, backups) must be identified, and the destination (storage class remapping, storage platform changes, etc.) storage prepared.

Finally, the plan needs to be transformed as needed (e.g., regeneration of TLS certifications, DNS changes, editing stale secrets, etc.). Kubernetes applications components need to be updated to reflect the new storage resources that will be created as a part of the recovery.

Once this execution plan is in place, the backup platform needs to be able to translate it into relevant Kubernetes API calls to create the required resources (e.g., create a load balancer or recreate a secret). This process ensures that all required Kubernetes resources and microservices that make up a cloud native application are redeployed with the correct configuration. The diagram below outlines this involved restore process.

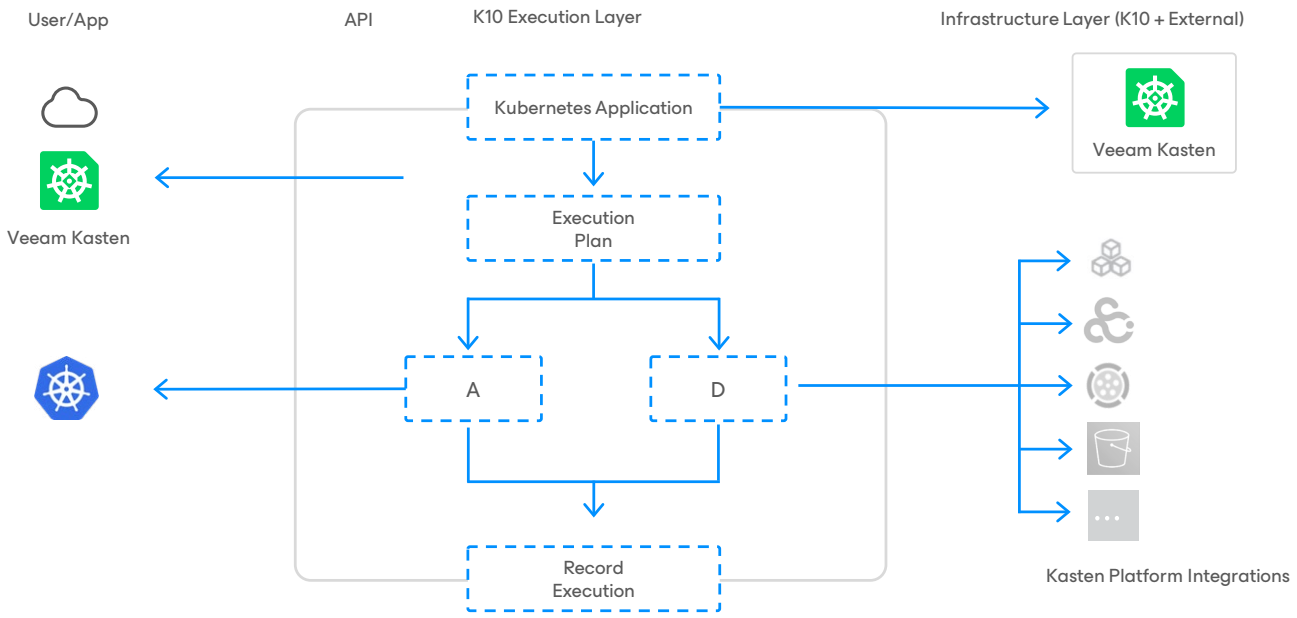

4.3 Operations

Operational best practices are typically the biggest challenge for enterprises, especially when implementing new tools, services and capabilities for an extremely dynamic infrastructure. It is important to ensure that a Kubernetes native backup platform can be used at scale, provides operations teams with the workflow capabilities they require, and meets compliance and monitoring requirements. Another important aspect is the impact, or lack thereof, on developers. One of the greatest benefits of Kubernetes is providing developers with a quick and easy way to deploy applications, roll out upgrades, and the ability to do it at scale. If a backup platform hinders those use cases, developers will find ways to circumvent any processes put in place. There should be zero code, packaging, toolchain, or deployment changes required for developers. At the same time, operators should be able to give self-service capabilities to application developers such as the ability to restore their own application or the option to customize and extend backup operations for their data services (e.g., use of custom or database-vendor tools, cross-service coordination and quiescing, etc.). Further, the ability for all the developer interactions with the backup platform to be API drive is also a must-have requirement. It is therefore essential that any backup platform deployed can meet the needs of both container platform operations teams and developers.

From the operator's perspective, they should not need to focus on the hundreds of Kubernetes components that make up the application. Backup policies need to not only be completely automated, but they also need to focus on the application and not individual resources or storage infrastructure. Materialization of the policy to concretely define the application components that need to be protected should only happen at policy execution time to ensure that all components in a rapidly changing application are captured without requiring manual policy updates. Similarly, backup policies need to be broad and label-based so that they automatically pick up new applications as soon as they are deployed (e.g., policy that matches all applications that use MongoDB or are deployed via the Helm package management tool). This will not only save the operations teams from having to build out manual change control processes but will also ensure that, as applications are created and destroyed at scale, they never fall out of compliance with backup SLAs.

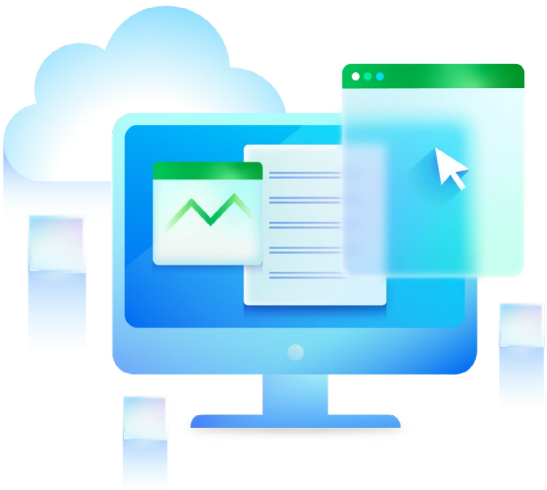

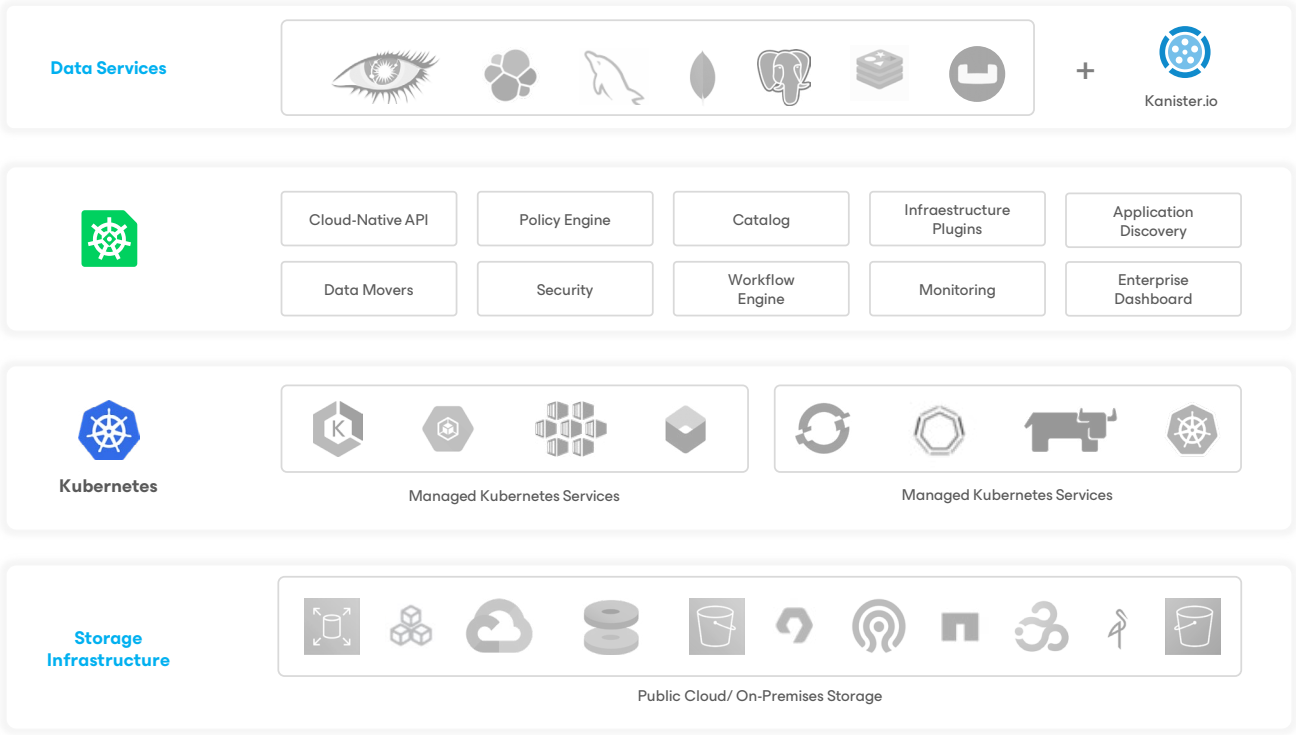

4.4 Security

Security is at the forefront of every product deployed in an enterprise production environment, regardless of whether it is deployed in a public cloud or using an on-premises infrastructure. Controls around identity and access management and role-based access control (RBAC) must be implemented. RBAC gives users and groups specific, and often restricted, user privileges or access privileges into the actual backup platform. This allows different personas in an operations team to adopt a least-privilege approach to common tasks such as monitoring backups, verifying backup success and integrity, and performing requested restores. RBAC also allows for use cases such as granting developers permissions for fast restore and clones from snapshots, but only grants certain team members access to backups stored in off-site secondary storage locations.

In a public cloud, apart from the security requirements described above, a cloud native backup platform also needs deep integration into the cloud's Identity and Access Management (IAM) systems, Key Management Systems (KMSs) and certificate management. Further, a truly Kubernetes native data protection system will integrate not just into the cloud provider's authentication solution (e.g., ODIC, OAuth, Token-based auth, etc.) but will also do so without requiring any extra user or group management, new tools, or new APIs for RBAC policies. All these features will be exposed via a Kubernetes native CRD-based API for a well-architected, cloud native backup platform. The final security aspect that must always be implemented is encryption. Protecting data, whether in transit or at rest, is critical to ensuring that data is secure whenever it has left the compute environment.

4.4.1 Encryption in Motion

When moving or copying data between Kubernetes clusters or storage environments, making sure that the data is encrypted as it leaves one end point and arrives in another is key. Using object storage as an example, an on-premises Kubernetes application deployment that needs to offload backups to AWS S3 will typically transfer data over an external internet connection. The backup platform must always ensure that the data is encrypted using well-known protocols such as TLS when being copied over the internet.

4.4.2 Encryption at Rest

Continuing with the example above, when the data is finally stored in a secondary location, confirming that it is encrypted is a critical security consideration. Simply implementing RBAC and related security policies at the control layer is insufficient if the data is not encrypted at rest. Using well-proven encryption algorithms such as AES-256-GCM with per-application encryption keys will prevent accidental data leaks or even malicious copying by rogue infrastructure operators or malicious external entities.

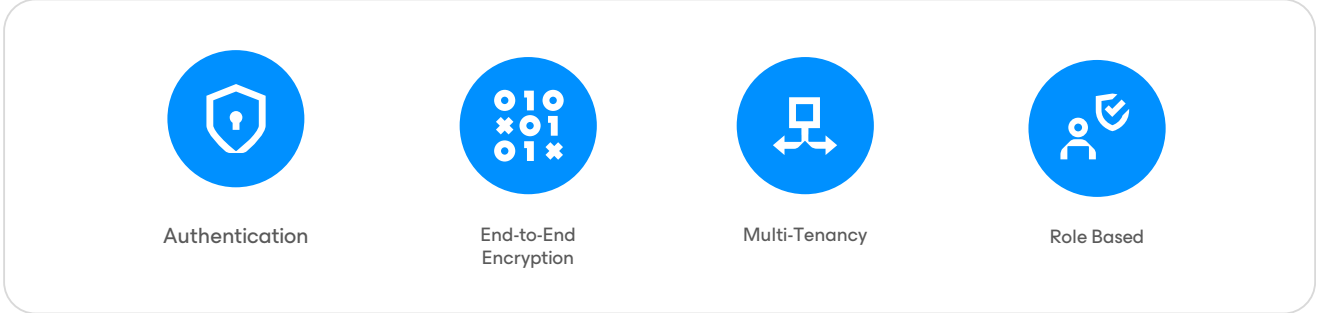

4.5 Mobility

The diagrams below illustrate the powerful platform portability feature offered by Kubernetes. This feature can be leveraged by a backup platform to enable a wide range of use cases. These include transferring data across namespaces within the same cluster, across different storage systems, across various Kubernetes clusters, distributions, and versions. It also extends to spanning availability zones within the same region, bridging regions within the same cloud, connecting cloud or hybrid environments, and even linking test and development environments.

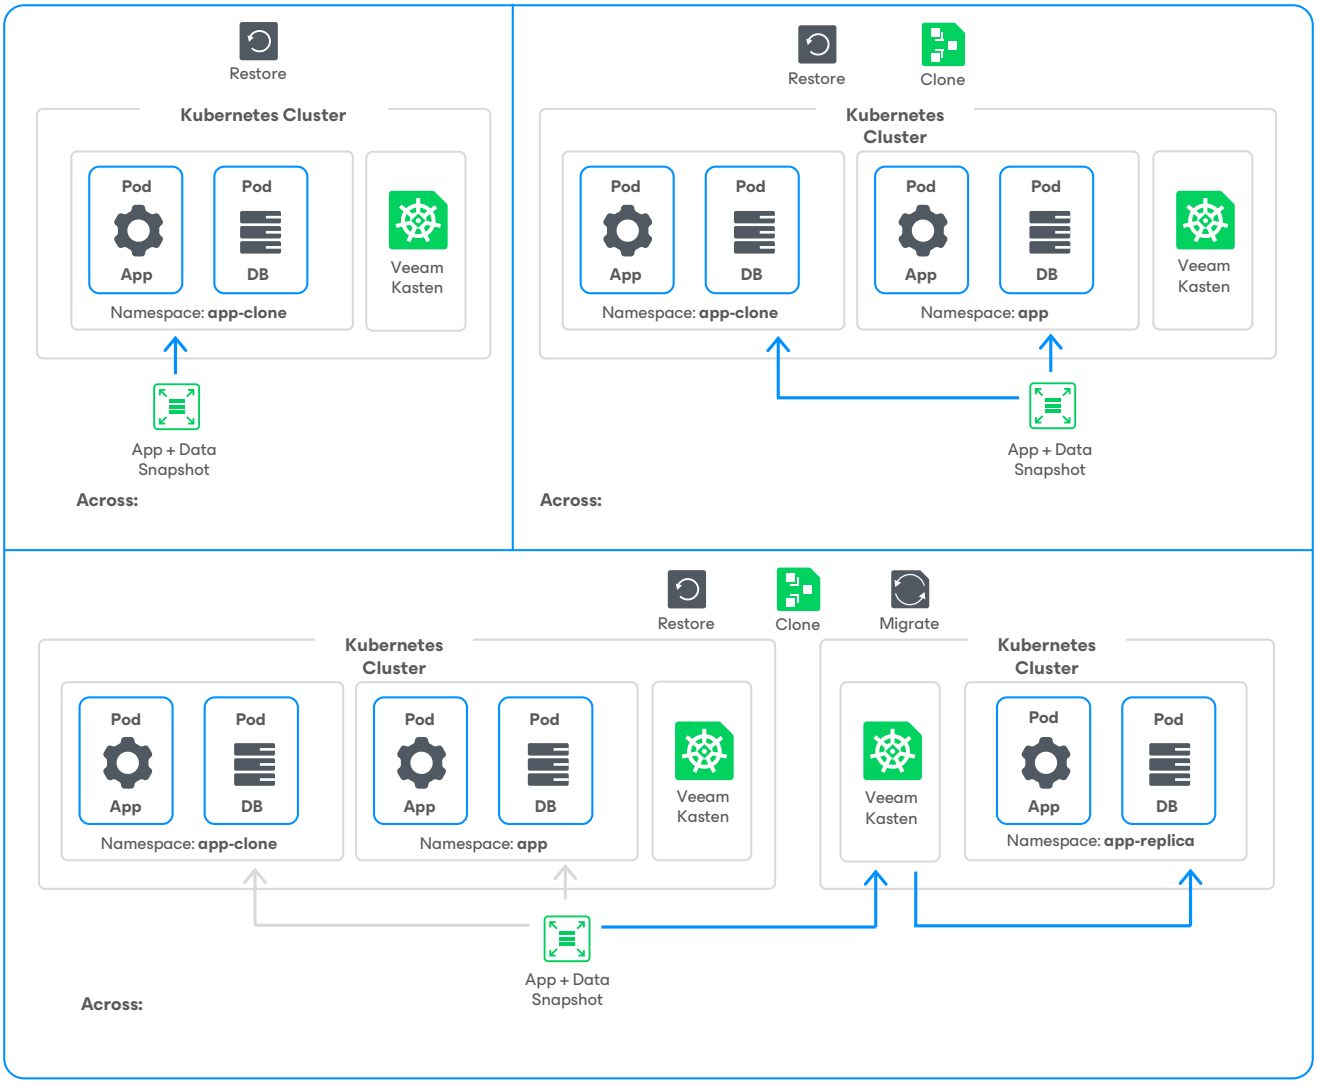

With the ecosystem diversity in Kubernetes offerings available on-premises and in the cloud, it is also critical that a backup and data management solution can migrate Kubernetes applications across arbitrary source and destination clusters that could be running on wildly heterogeneous infrastructures. For example, when migrating a workload from Amazon Elastic Kubernetes Service (Amazon EKS) to Microsoft Azure Kubernetes Service (AKS), you will see the following on each cluster:

```
kind: StorageClass
apiVersion: storage.k8s.io/v1
metadata:
        name: gp2
provisioner: kubernetes.io/aws-ebs
parameters:
        type: gp2
        fsType: ext4
```

```
kind: StorageClass
apiVersion: storage.k8s.io/v1
metadata:
        name: managed-premium-retain
provisioner: kubernetes.io/azure-disk
reclaimPolicy:Retain
parameters:
        storageaccounttype: Premium _ LRS
        kind: Managed
```

These different storage classes are just the tip of the iceberg as far as the differences between distributions go, even though these distributions might be based on the same underlying Kubernetes version. Ensuring that a backup platform can reliably perform restores across these different distributions and infrastructure options while being able to automatically transform the application backup to fit the new restore environment is critical.

Ultimately, when performing migrations of workloads across namespaces, clusters, regions and even Kubernetes distributions, any reliable backup platform being used must be able to understand all application dependencies and translate them across environments. In the same manner a restore is planned and executed in a cloud native environment, a migration plan is needed to ensure that infrastructure (e.g., load balancers), cluster-wide and application dependencies are available or transformed to an equivalent resource for a successful migration execution. It is not only containers and storage volumes that must be migrated, but also FQDNs, secrets and DNS addresses that must be modified in-flight during a migration.

Chapter 5. Conclusion

In the previous section, we discussed the critical requirements and recommendations on how to implement a Kubernetes backup solution. They can be summarized into the following five best practices:

Architecture

The platform used to protect Kubernetes applications needs to automatically discover all the application components running on your cluster and treat the application as the unit of atomicity. The application must include the state that spans across storage volumes, databases (NoSQL/Relational) as well as configuration data included in Kubernetes objects such as configmaps and secrets.

Recoverability

The data protection platforms must allow you to restore the application components you want and where you want them. You should also have the granular control to restore only an application subset such as the data volume. The approach must make restoring simple and powerful by also allowing you to select the appropriate point of time copy of the application.

Operations

It is important to ensure that a Kubernetes native backup platform can be used at scale, provide operations teams with the workflow capabilities they require, and meets compliance and monitoring requirements. Operators should be able to give self-service capabilities to application developers without requiring application code or deployment changes.

Security

Controls around identity and access management and role-based access control (RBAC) must be implemented. RBAC allows different personas in an operations team to adopt a least-privilege approach to common tasks such as monitoring. Encryption at rest and in transit must always be implemented to ensure that data is secure whenever it has left the compute environment.

Application Mobility

Living in a multi-hybrid-cloud world, a cloud native data protection platform needs to be able to be flexible in the support for multiple distributions and offer capabilities that allow for the mobility of workloads and applications across all these diverse environments. Mobility capabilities are required across use cases including application restoring, cloning, and container migration. Ensuring you adhere to the common best practices found in this guide will help you provide a consistent and reliable offering if you face data loss or corruption, or even a complete outage.

Veeam is a recognized leader in backup solutions for virtualized workloads. With the accelerated pace of Kubernetes applications and deployments, Veeam Kasten for Kubernetes offers the #1 data protection and mobility solution to address the cloud native data protection needs for enterprises. The Veeam Kasten data protection platform has been purpose-built for Kubernetes and provides for the backup, restore, disaster recovery, application mobility and ransomware protectionof your entire Kubernetes application while keeping the best practices highlighted above.

Veeam Kasten for Kubernetes

Veeam provides a single platform for modernizing backup, accelerating hybrid cloud and securing your data. Our solutions are simple to install and run, flexible enough to fit into any environment and always reliable. Trusted by the world's largest organizations, Veeam Kasten delivers secure, Kubernetes native data protection and application mobility, at scale, and across a wide range of distributions and platforms. Proven to recover entire applications quickly and reliably, coupled with its core tenetsimplicity, Veeam Kasten gives operations and app teams confidence to withstand the unexpected.

For more information, visit:
veeam.com/kubernetes-native-backup-and-re store

About Veeam Software

Veeam®, the #1 global market leader in data protection and ransomware recovery, is on a mission to help every organization not just bounce back from a data outage or loss but bounce forward. With Veeam, organizations achieve radical resilience through data security, data recovery, and data freedom for their hybrid cloud. The Veeam Data Platform delivers a single solution for cloud, virtual, physical, SaaS, and Kubernetes environments that gives IT and security leaders peace of mind that their apps and data are protected and always available. Headquartered in Seattle with offices in more than 30 countries, Veeam protects over 450,000 customers worldwide, including 74% of the Global 2000, who trust Veeam to keep their businesses running. Radical Resilience starts with Veeam. Learn more at www.veeam.com or follow Veeam on LinkedIn @veeam-software and X @veeam.