

Microservices and software development infrastructure upgrade

Stefano Bovina

Agenda

Why an upgrade?

Microservices and related challenges

Infrastructure overview


Infrastructure architecture deep dive

Security management

“New software” highlights

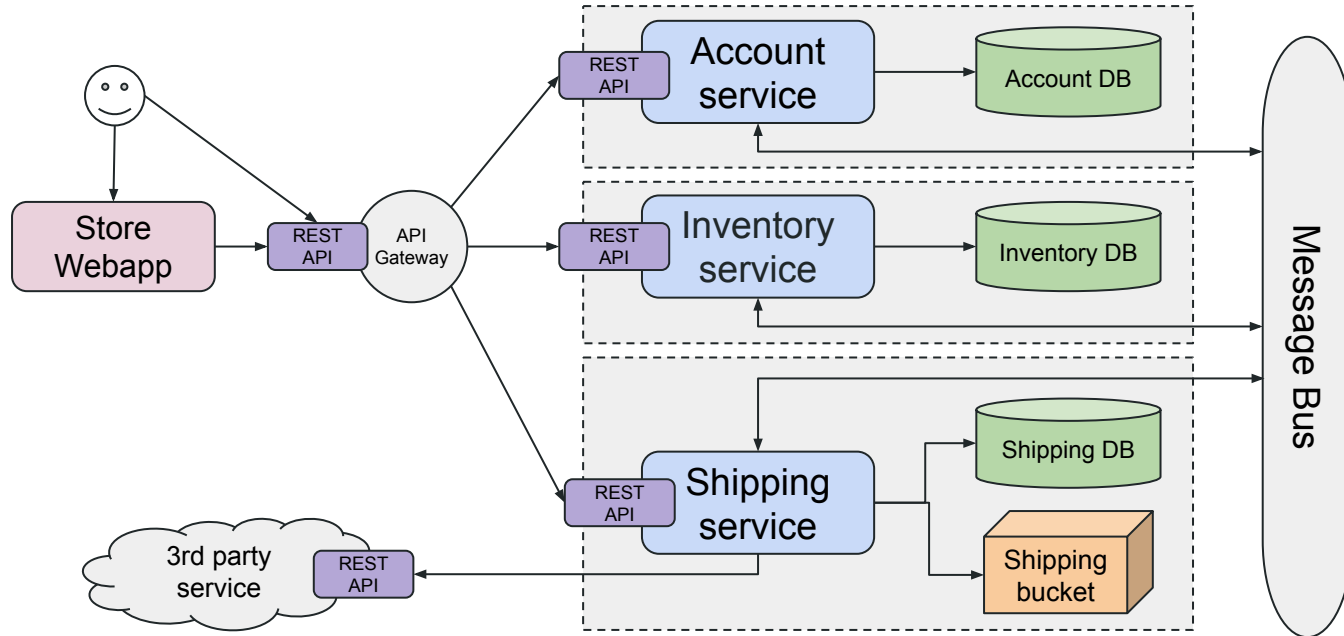
Why an upgrade?

- A lot of legacy applications need to be replaced or redesigned from scratch
- To produce “better” software
- To improve management of security and compliance aspects
- Technological upgrade
- For a better microservice management



Microservices

“...the microservice architectural style is an approach to developing a single application as a **suite of small autonomous services**, each **running in its own process** and communicating with lightweight mechanisms, often an HTTP resource API. These services are **built around business capabilities** and **independently deployable** by fully automated deployment machinery...”



Microservices

Pro:

- Strong Module Boundaries: Microservices reinforce modular structure
- Independent deployment
- Higher degree of organizational autonomy
- Technology Heterogeneity
- Optimized for replaceability
- Scaling independently
- Leads to Improved Fault Tolerance (if we understand and plan for failures)
- Ease of understanding of the codebase of the software system
- Isolation of data and isolation of processing around that data

Cons:

- Distribution: Distributed systems are harder to program, since remote calls are slow and are always at risk of failure
- Eventual Consistency: Maintaining strong consistency is extremely difficult for a distributed system, which means everyone has to manage eventual consistency (Multi-services transactions/changes are complex)
- Operational Complexity: You need a mature operations team to manage lots of services, which are being redeployed regularly
- Global automated testing is more complicated

Distributed computing fallacies and other requirements

F
a
l
l
a
c
i
e
s

Fallacies/requirements	Solutions
The network is reliable	Circuit breaker (Resilience4j), retry and timeout design pattern (Resilience4j), message queues (Kafka)
Latency is zero	caching strategy (Redis), bulk requests, placement/affinity policy (see Kubernetes policy)
Bandwidth is infinite	throttling policy, small payloads
The network is secure	firewall (network policy/micro segmentation), encryption (mTLS, Cert-manager), AuthN/AuthZ (OIDC/Oauth2)
Topology doesn't change	no hardcoded IPs, service discovery tools (see Kubernetes service discovery)
There is one administrator	DevOps culture
Transport cost is zero	standardized protocols like JSON, cost calculation
The network is homogeneous	Circuit breaker (Resilience4j), retry and timeout design pattern (Resilience4j)
Observability	Monitoring system (Prometheus/Grafana/Sensu/InfluxDB), Log aggregation (ELK)
Automation culture/tools	CI/CD platform (Gitlab, ArgoCD), IaC paradigm (Helm/Kustomize/Puppet)
Secret management	Vault

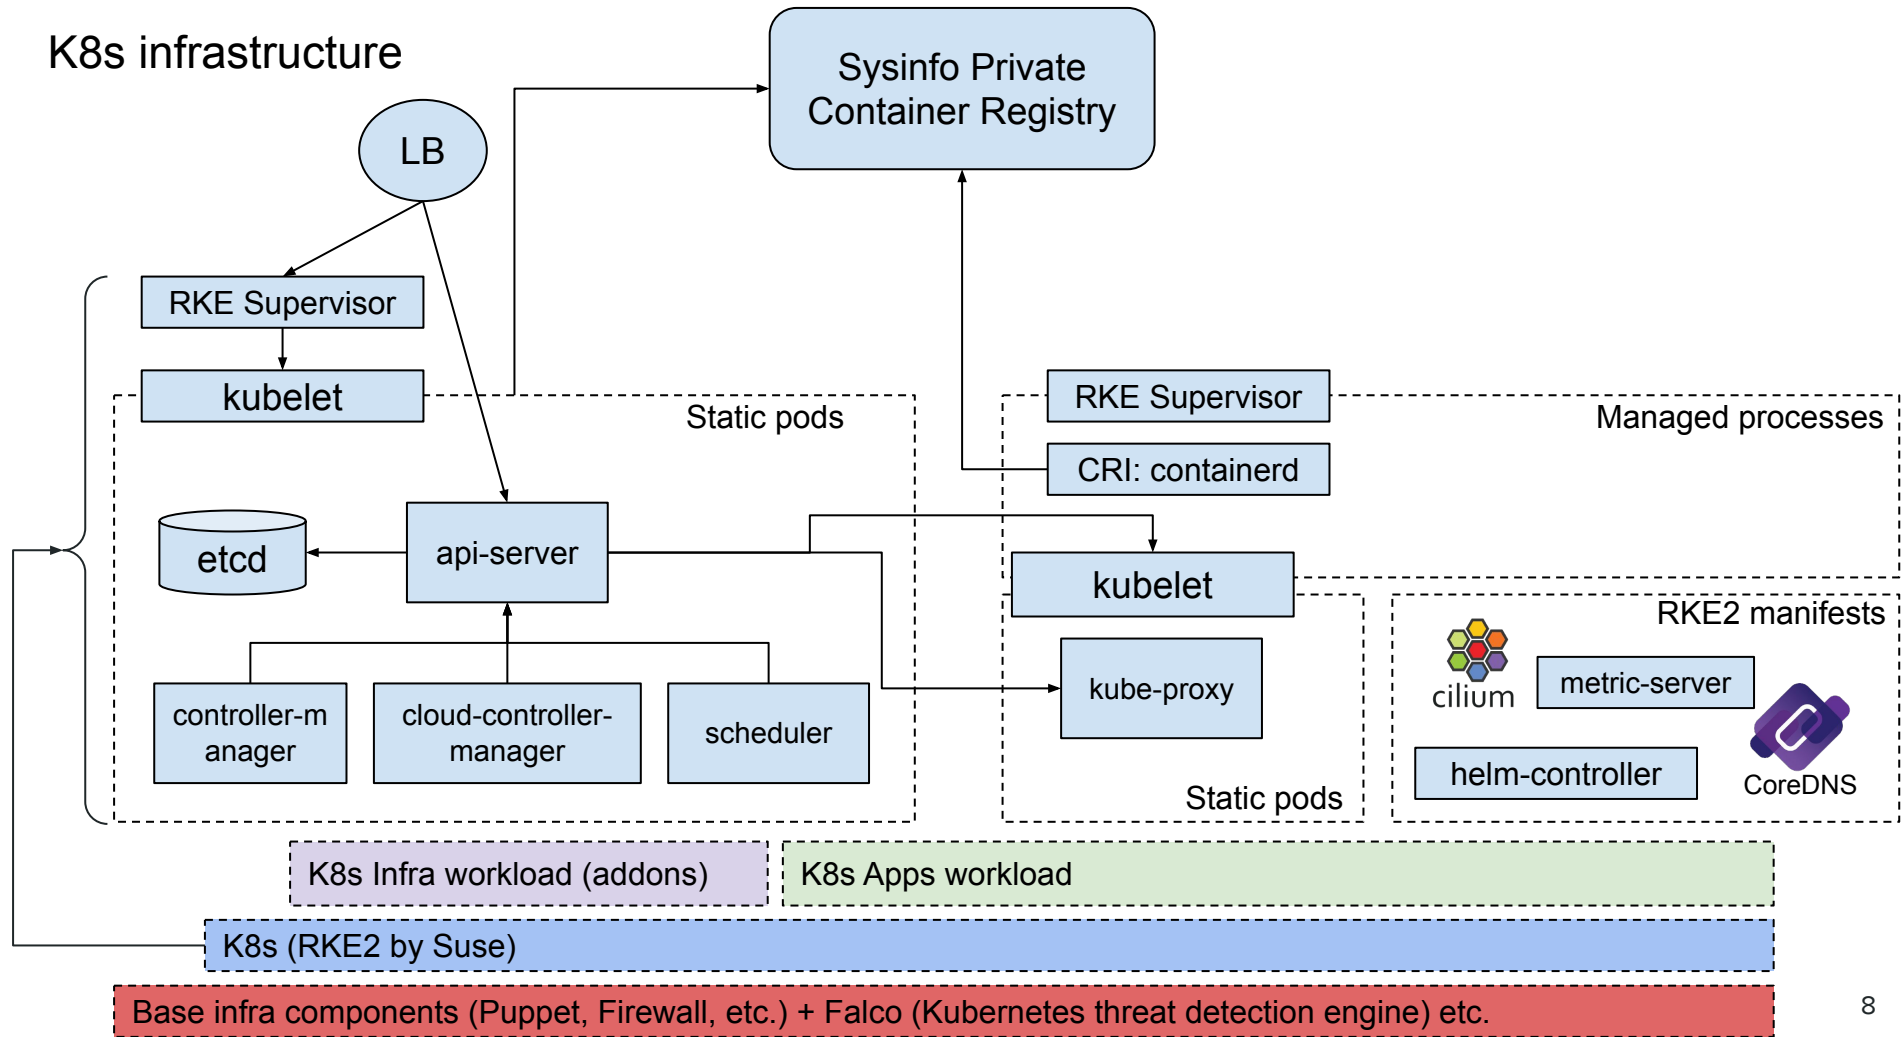
Kubernetes (k8s)



“...a portable, extensible, open-source platform for managing containerized workloads and services, that facilitates both declarative configuration and automation.”

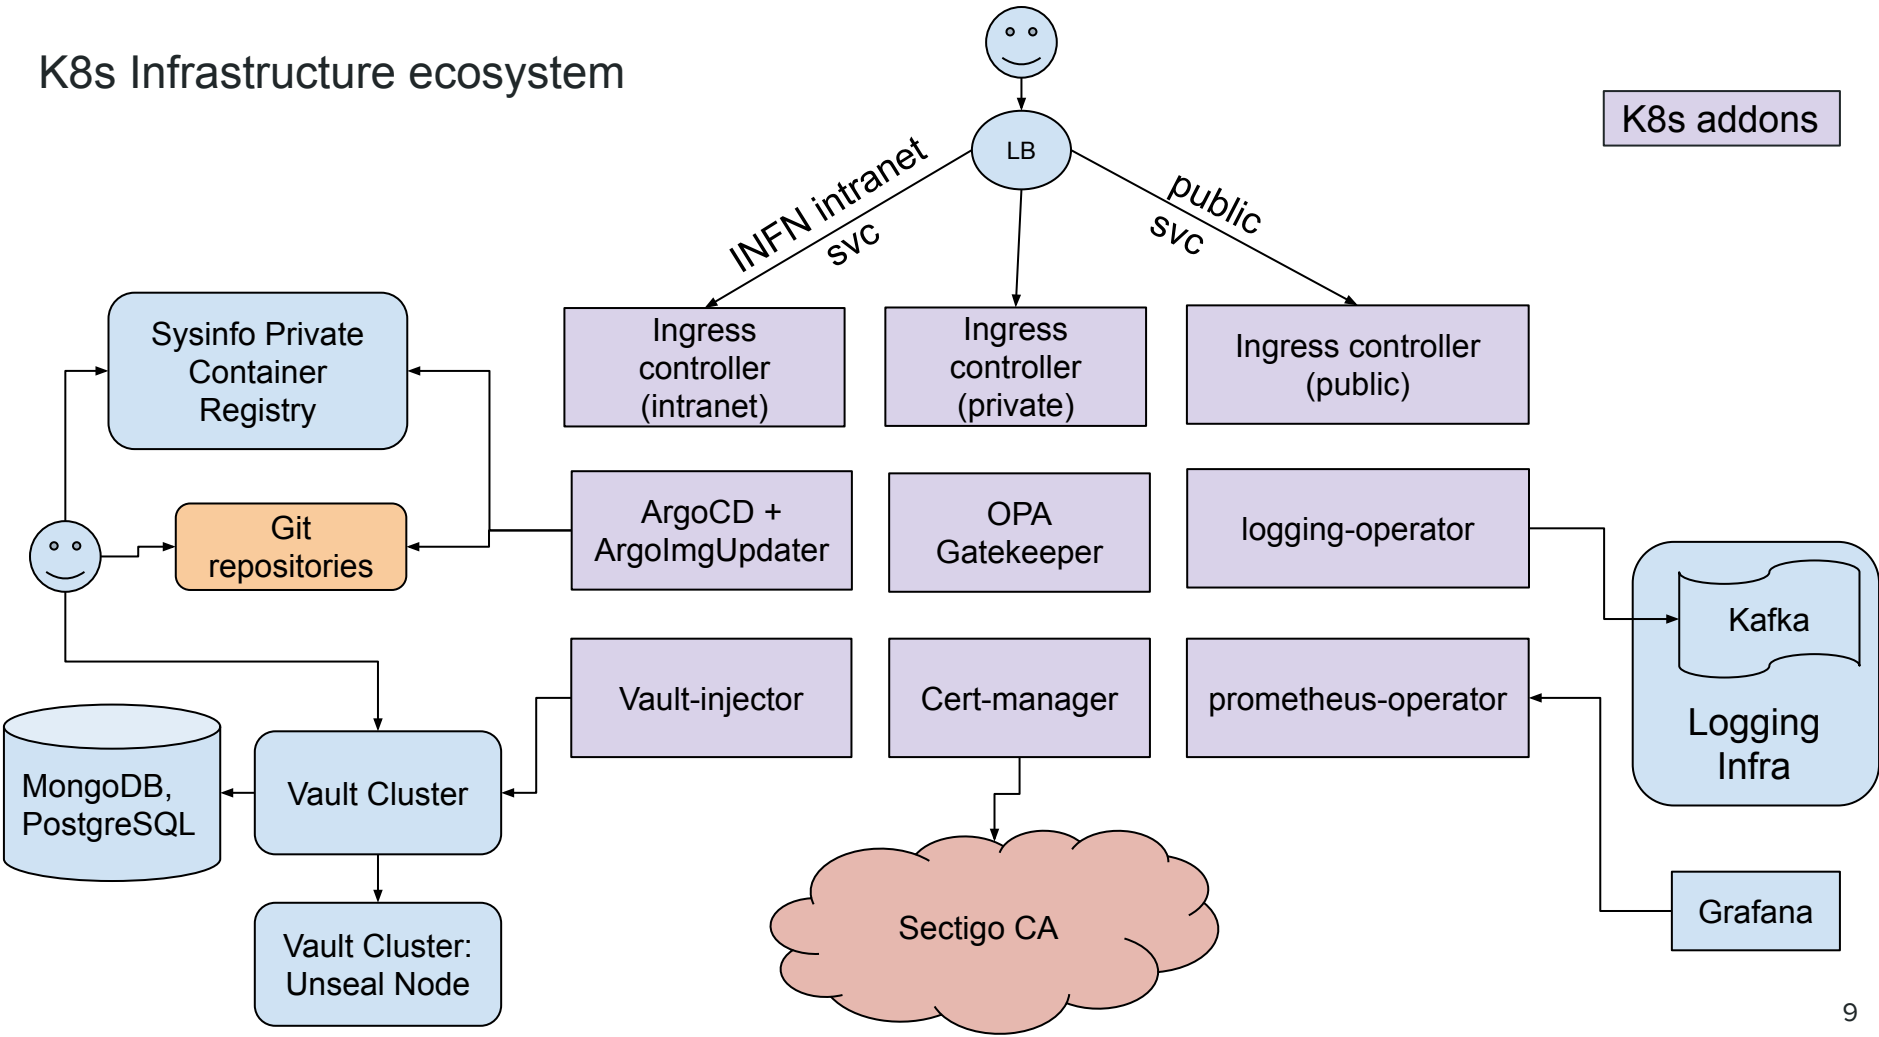
- Google open-sourced the Kubernetes project in 2014
- Features:
 - Service discovery, load balancing, horizontal scaling
 - Self-healing
 - Automated rollouts and rollbacks
 - Secret and configuration management
 - etc.

K8s infrastructure



K8s Infrastructure ecosystem

K8s addons



Harbor container registry



Projects < sysinfo_apps-test

booking-backend


Info Artifacts

SCAN

ACTIONS



	Artifacts	Pull Command	Tags	Size	Vulnerabilities	Annotations	Labels	Push Time	Pull Time
<input type="checkbox"/>	sha256:1332471f		1638034685-baltig.62854	212.48MB	60 Total - 0 Fixable			11/27/21, 6:38 PM	11/29/21, 9:00 AM
<input type="checkbox"/>	sha256:a1a6319b		1637831143-baltig.62670	212.48MB	60 Total - 0 Fixable			11/25/21, 10:06 AM	11/29/21, 9:00 AM
<input type="checkbox"/>	sha256:585b73ea		1637751272-baltig.62616	212.48MB	60 Total - 0 Fixable			11/24/21, 11:54 AM	11/29/21, 9:00 AM
<input type="checkbox"/>	sha256:3bdb7ed0		1637685825-baltig.62580	212.48MB	60 Total - 0 Fixable			11/23/21, 5:44 PM	11/29/21, 9:00 AM
<input type="checkbox"/>	sha256:4634cb71		1637600113-baltig.62518	212.44MB	60 Total - 0 Fixable			11/22/21, 5:55 PM	11/29/21, 9:00 AM
<input type="checkbox"/>	sha256:6255da2e		1637581609-baltig.62502	195.21MB	60 Total - 0 Fixable			11/22/21, 12:47 PM	11/29/21, 9:00 AM
<input type="checkbox"/>	sha256:c2074d04		1637171965-baltig.62236	195.21MB	60 Total - 0 Fixable			11/17/21, 6:59 PM	11/29/21, 9:00 AM
<input type="checkbox"/>	sha256:4a66b83d		1636803083-baltig.61984	193.32MB	60 Total - 0 Fixable			11/13/21, 12:31 PM	11/29/21, 9:01 AM
<input type="checkbox"/>	sha256:0bf79f91		1636623999-baltig.61840	193.32MB	60 Total - 0 Fixable			11/11/21, 10:47 AM	11/29/21, 9:01 AM
<input type="checkbox"/>	sha256:df5ae4a9		1636567674-baltig.61805	193.32MB	60 Total - 0 Fixable			11/10/21, 7:08 PM	11/29/21, 9:01 AM

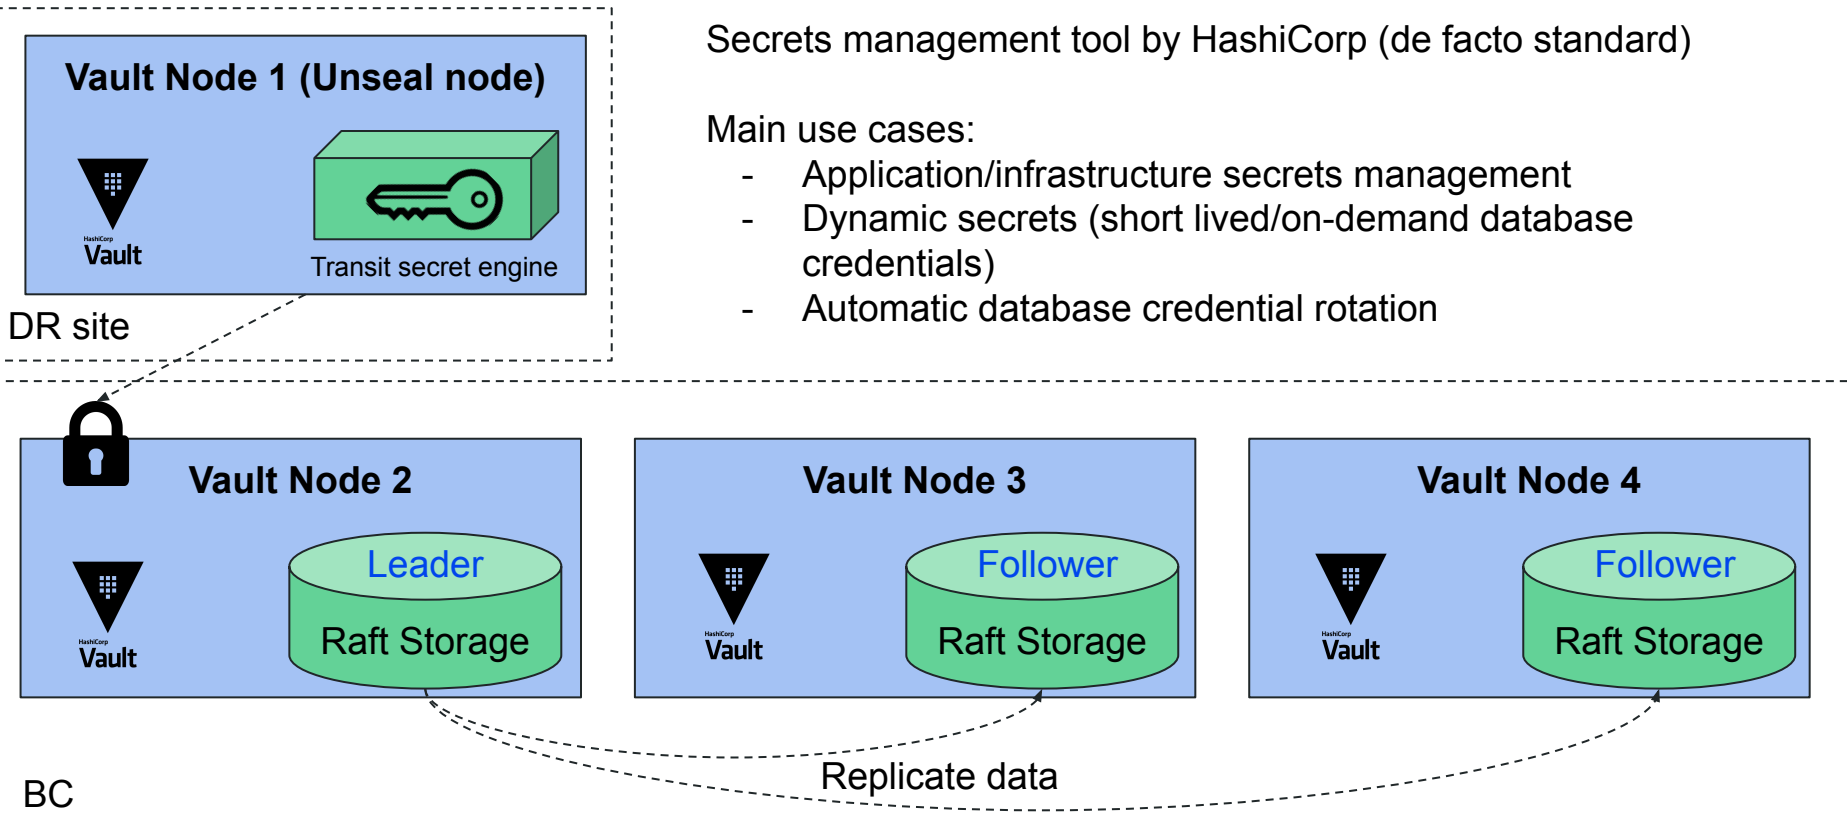


Vault

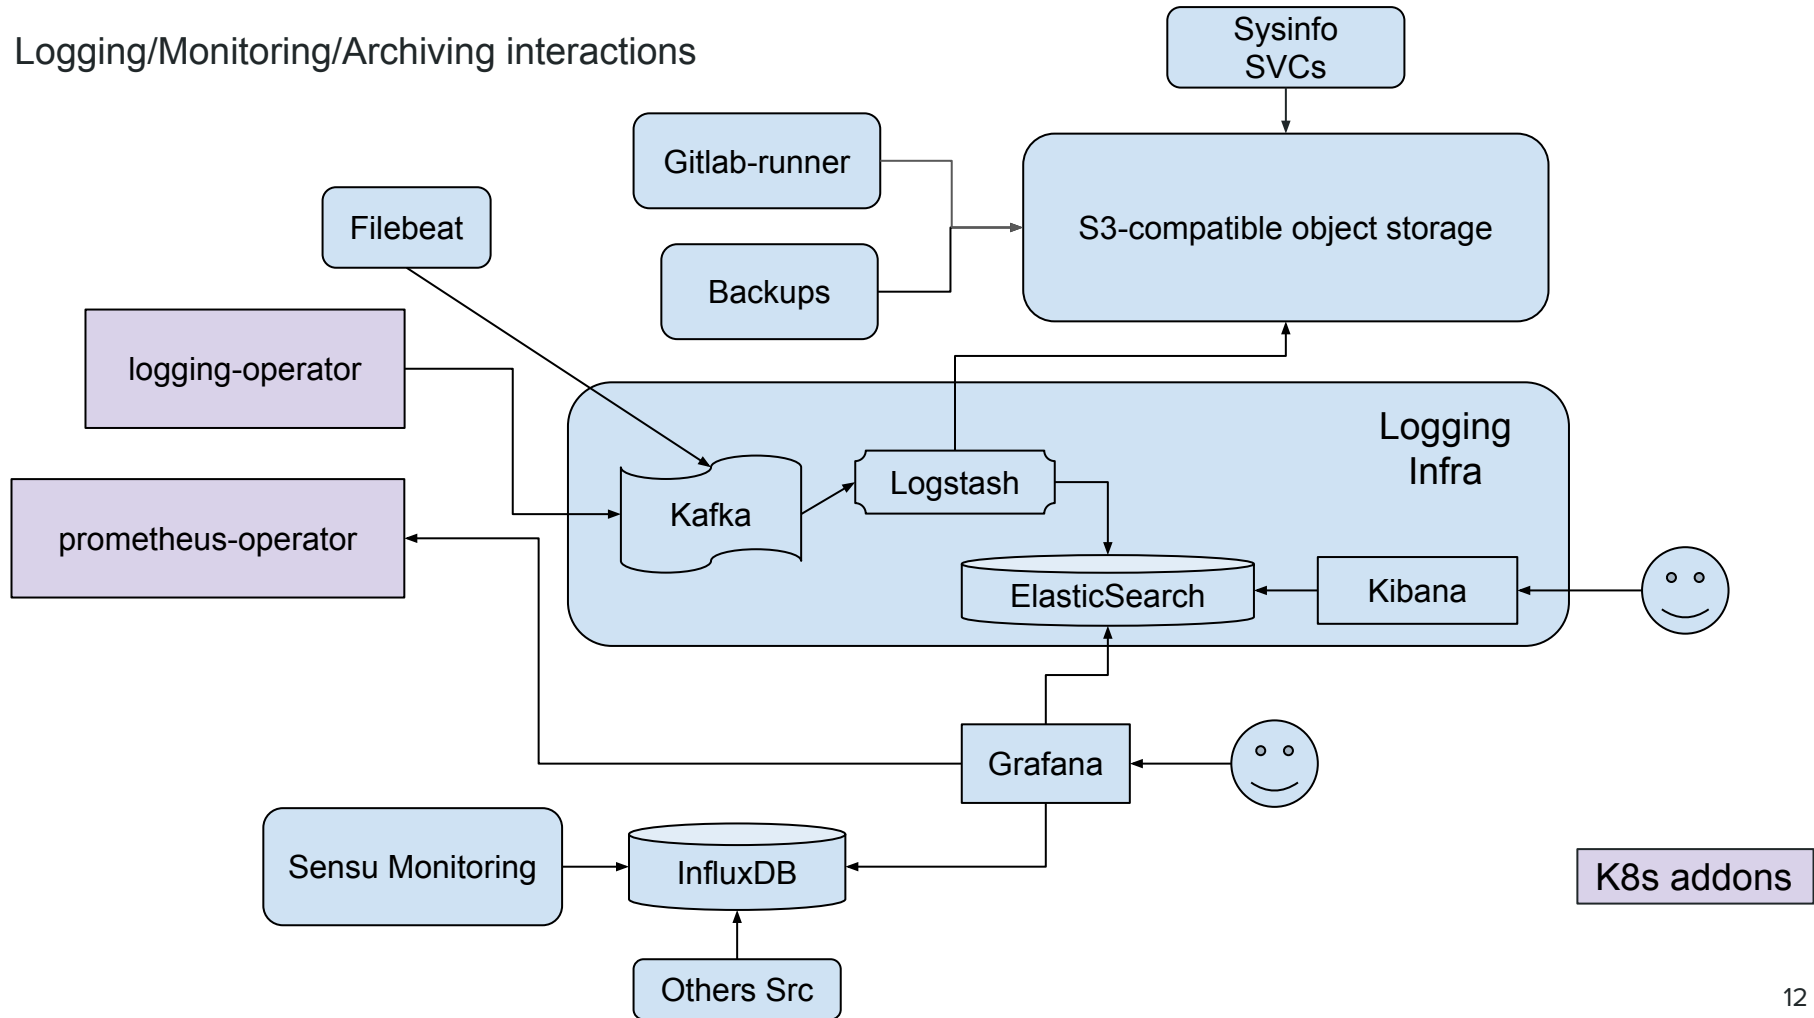
Secrets management tool by HashiCorp (de facto standard)

Main use cases:

- Application/infrastructure secrets management
- Dynamic secrets (short lived/on-demand database credentials)
- Automatic database credential rotation



Logging/Monitoring/Archiving interactions



Minio

Software-defined high performance object storage

Features/reasons why we use it:

- Highly available and horizontally scalable
- API compatible with Amazon's S3 (de-facto standard API for business applications to store unstructured data)
- Bucket Versioning
- Object Lock and Immutability - Write-Once Read-Many (WORM)
- Bucket Notifications (i.e. Kafka)
- Server-Side Bucket Replication (BC/DR)
- Object Lifecycle Management (Transition/Expiration)
- Encryption

Use cases:

- gitlab-runner distributed cache
- Long term archiving (es: logs, backups etc.)
- Sysinfo application data



30.4K+

GITHUB STARS

720.9M+

DOCKER PULLS

15.8K+

SLACK MEMBERS


756

CONTRIBUTORS

ArgoCD and GitOps

GitOps: versioned CI/CD on top of declarative infrastructure:

- a Git repository that always contains declarative descriptions of the desired infrastructure state
- infrastructure state versioned in Git
- use of continuous integration/continuous delivery pattern
- automated processes to make the production environment match the described state in the repository
- a change to the infrastructure (e.g. a new application) => modify the repository



ArgoCD

<p>api-gw</p> <p>Project: sysinfo-apps</p> <p>Labels:</p> <p>Status: ♥ Healthy 🟢 Synced</p> <p>Repository:</p> <p>Target Revision: master</p> <p>Path: overlays/api-gw</p> <p>Destination: in-cluster</p> <p>Namespace: api-gw</p> <p><a>SYNC <a>REFRESH <a>DELETE</p>	<p>appman</p> <p>Project: sysinfo-apps</p> <p>Labels:</p> <p>Status: ♥ Healthy 🟢 Synced</p> <p>Repository:</p> <p>Target Revision: master</p> <p>Path: overlays/appman</p> <p>Destination: in-cluster</p> <p>Namespace: appman</p> <p><a>SYNC <a>REFRESH <a>DELETE</p>	<p>booking</p> <p>Project: sysinfo-apps</p> <p>Labels:</p> <p>Status: ♥ Healthy 🟢 Synced</p> <p>Repository:</p> <p>Target Revision: master</p> <p>Path: overlays/booking</p> <p>Destination: in-cluster</p> <p>Namespace: booking</p> <p><a>SYNC <a>REFRESH <a>DELETE</p>	<p>concorsi</p> <p>Project: sysinfo-apps</p> <p>Labels:</p> <p>Status: ♥ Healthy 🟢 Synced</p> <p>Repository:</p> <p>Target Revision: master</p> <p>Path: overlays/concorsi</p> <p>Destination: in-cluster</p> <p>Namespace: concorsi</p> <p><a>SYNC <a>REFRESH <a>DELETE</p>
<p>consuntivi</p> <p>Project: sysinfo-apps</p> <p>Labels:</p> <p>Status: ♥ Healthy 🟢 Synced</p> <p>Repository:</p> <p>Target Revision: master</p> <p>Path: overlays/consuntivi</p> <p>Destination: in-cluster</p> <p>Namespace: consuntivi</p> <p><a>SYNC <a>REFRESH <a>DELETE</p>	<p>fondipersonale</p> <p>Project: sysinfo-apps</p> <p>Labels:</p> <p>Status: ♥ Healthy 🟢 Synced</p> <p>Repository:</p> <p>Target Revision: master</p> <p>Path: overlays/fondipersonale</p> <p>Destination: in-cluster</p> <p>Namespace: fondipersonale</p> <p><a>SYNC <a>REFRESH <a>DELETE</p>	<p>identity</p> <p>Project: sysinfo-apps</p> <p>Labels:</p> <p>Status: ♥ Healthy 🟢 Synced</p> <p>Repository:</p> <p>Target Revision: master</p> <p>Path: overlays/identity</p> <p>Destination: in-cluster</p> <p>Namespace: identity</p> <p><a>SYNC <a>REFRESH <a>DELETE</p>	<p>mail</p> <p>Project: sysinfo-apps</p> <p>Labels:</p> <p>Status: ♥ Healthy 🟢 Synced</p> <p>Repository:</p> <p>Target Revision: master</p> <p>Path: overlays/mail</p> <p>Destination: in-cluster</p> <p>Namespace: mail</p> <p><a>SYNC <a>REFRESH <a>DELETE</p>
<p>preventivi</p> <p>Project: sysinfo-apps</p> <p>Labels:</p> <p>Status: ♥ Healthy 🟢 Synced</p> <p>Repository:</p> <p>Target Revision: master</p> <p>Path: overlays/preventivi</p> <p>Destination: in-cluster</p> <p>Namespace: preventivi</p> <p><a>SYNC <a>REFRESH <a>DELETE</p>	<p>progetti</p> <p>Project: sysinfo-apps</p> <p>Labels:</p> <p>Status: ♥ Healthy 🟢 Synced</p> <p>Repository:</p> <p>Target Revision: master</p> <p>Path: overlays/progetti</p> <p>Destination: in-cluster</p> <p>Namespace: progetti</p> <p><a>SYNC <a>REFRESH <a>DELETE</p>	<p>storage</p> <p>Project: sysinfo-apps</p> <p>Labels:</p> <p>Status: ♥ Healthy 🟢 Synced</p> <p>Repository:</p> <p>Target Revision: master</p> <p>Path: overlays/storage</p> <p>Destination: in-cluster</p> <p>Namespace: storage</p> <p><a>SYNC <a>REFRESH <a>DELETE</p>	<p>titolidistudio</p> <p>Project: sysinfo-apps</p> <p>Labels:</p> <p>Status: ♥ Healthy 🟢 Synced</p> <p>Repository:</p> <p>Target Revision: master</p> <p>Path: overlays/titolidistudio</p> <p>Destination: in-cluster</p> <p>Namespace: titolidistudio</p> <p><a>SYNC <a>REFRESH <a>DELETE</p>

ArgoCD

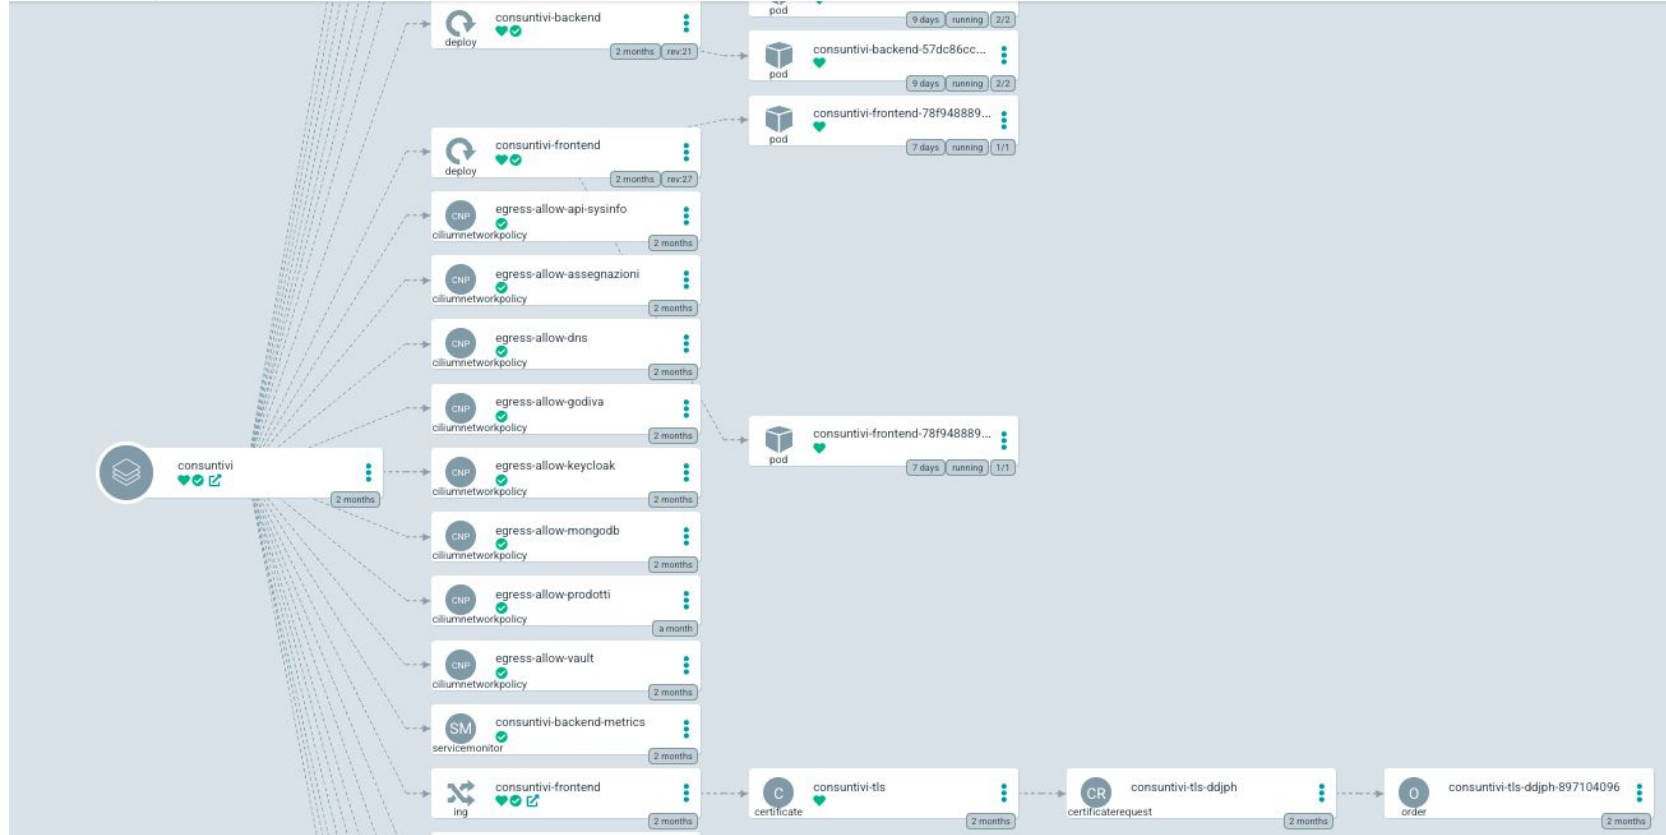
[MORE](#)
To master (c2c4d88)

LAST SYNC RESULT
Sync OK

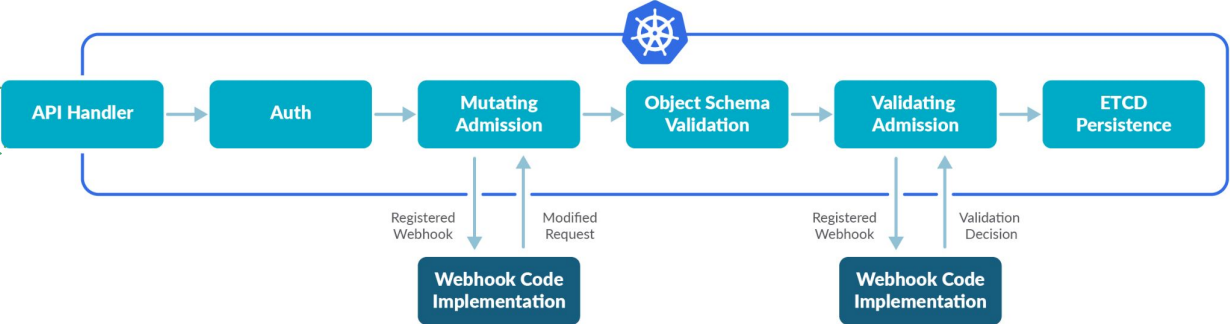
[MORE](#)
To c2c4d88

er <noreply@argoproj.io> -
matic update of consuntivi


Succeeded 7 days ago (Thu May 12 2022 17:10:38 GMT+0200)
Author: argocd-image-updater <noreply@argoproj.io> -
Comment: build: automatic update of consuntivi



OPA Gatekeeper



API request without admission controller



API request with admission controller



```

package k8spallowprivilegeescalationcontainer

violation{"msg": msg, "details": {}} {
  c := input_containers[]
  input_allow_privilege_escalation(c)
  msg := sprintf("Privilege escalation container is not allowed: %v", [c.name])
}

input_allow_privilege_escalation(c) {
  not has_field(c, "securityContext")
}

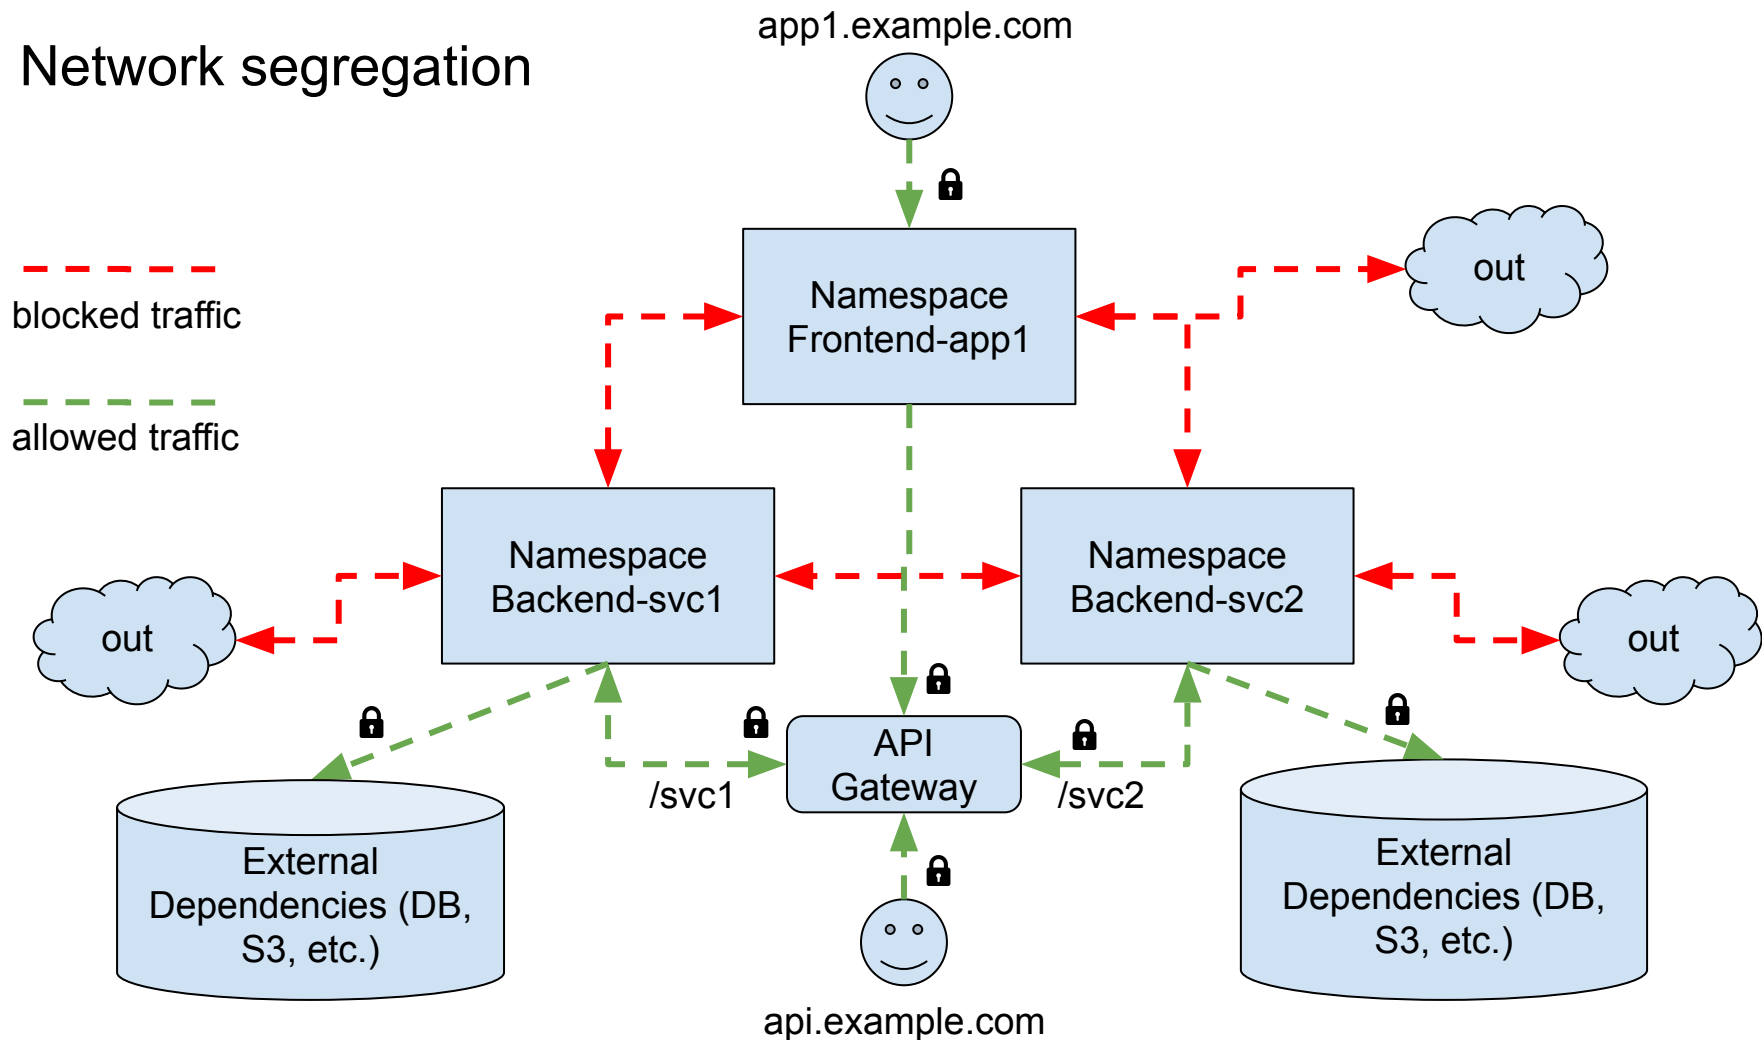
input_allow_privilege_escalation(c) {
  not c.securityContext.allowPrivilegeEscalation == false
}

input_containers[c] {
  c := input.review.object.spec.containers[]
}

input_containers[c] {
  c := input.review.object.spec.initContainers[]
}


# has_field returns whether an object has a field
has_field(object, field) = true {
  object[field]
}
    
```

Network segregation




Circuit breaker: the problem


Local failures can propagate all over your architecture and destroy all your systems



- C makes a request to F
- F makes a request to B
- After short time Client got a response



- If B:
 - is down
 - is up but the network is unreachable (network partition)
 - is up but is very slow
- C makes a request to F
- F need to wait until the timeout before came back to C
- During the timeout period requests pile up on the F service




- F became unresponsive

Circuit breaker: the solution

“A service client should invoke a remote service via a proxy that functions in a similar fashion to an electrical circuit breaker.”

- If a call fails, increment the number of failed calls by one
- If the number of failed calls goes above a certain threshold, open the circuit
- If the circuit is open, immediately return with an error **or a default response**
- If the circuit is open and some time has passed, half-open the circuit
- If the circuit is half-open and the next call fails, open it again
- If the circuit is half-open and the next call succeeds, close it



Strategy	Implementations	Fit
Black Box	<ul style="list-style-type: none">● Proxies● Service meshes	Fail fast
White Box	Libraries (e.g. Resilience4j)	Fallbacks relying on business logic

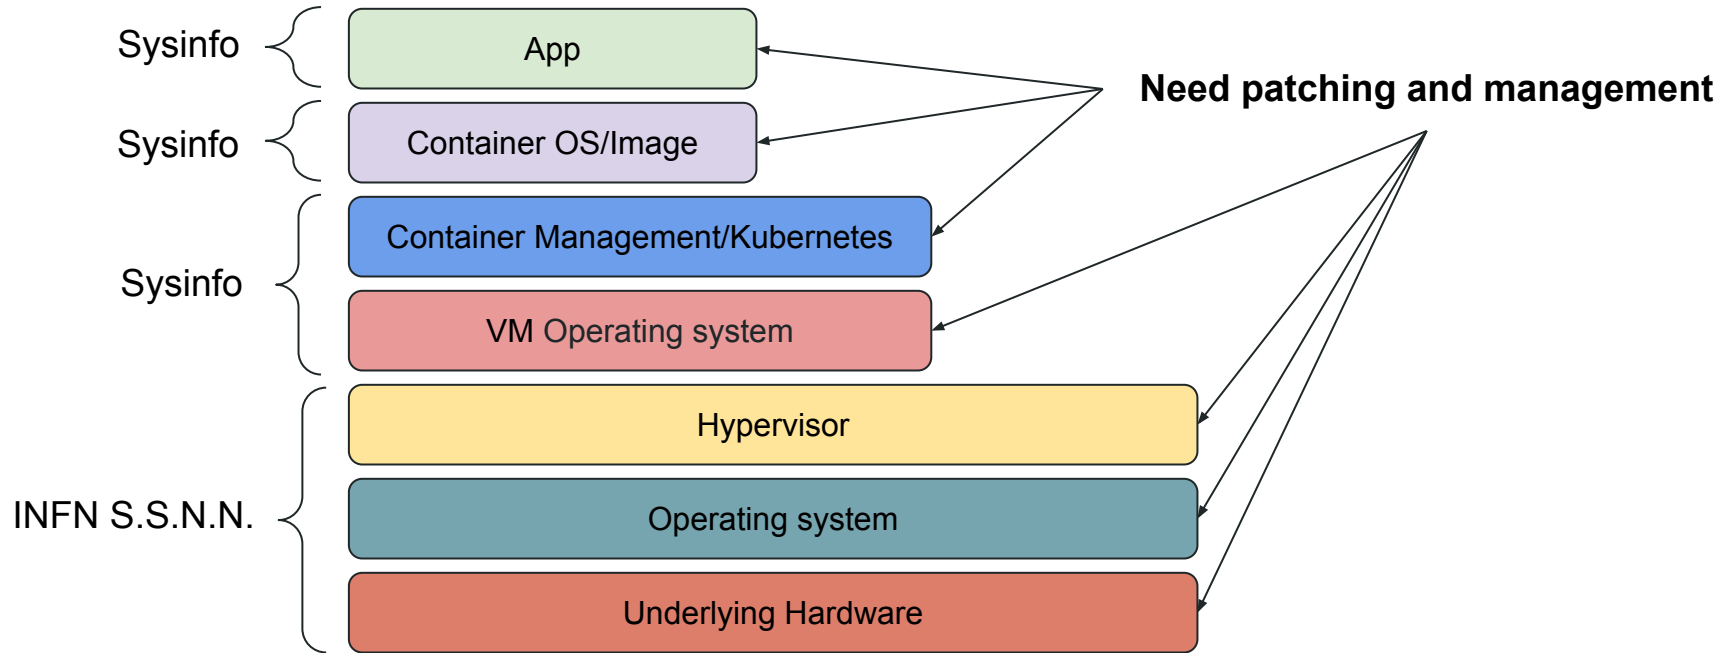
How do we secure this?

The image displays a vast collection of logos for cloud-native technologies, organized into several functional categories:

- Database:** KV, V, etc.
- Streaming & Messaging:** cloudvents, HELM, etc.
- Application Definition & Image Build:** HELM, etc.
- Continuous Integration & Delivery:** argo, flux, etc.
- Platform:** Certified Kubernetes - Distribution, Hosted, Installer, PaaS/Container Service.
- Scheduling & Orchestration:** kubernetes, etc.
- Coordination & Service Discovery:** etcd, etc.
- Remote Procedure Call:** gRPC, etc.
- Service Proxy:** envoy, etc.
- API Gateway:** etc.
- Service Mesh:** etc.
- Cloud Native Storage:** etc.
- Container Runtime:** cri-o, etc.
- Cloud Native Network:** etc.
- Automation & Configuration:** etc.
- Container Registry:** etc.
- Security & Compliance:** falco, etc.
- Key Management:** etc.
- Observability and Analysis:** Monitoring (Thanos, etc.), Logging (fluentd, etc.), Tracing, Chaos Engineering.
- Kubernetes Certified Service Provider**
- Kubernetes Training Partner**



Infrastructure patching&management



Infrastructure security management

- GitOps approach
- Static analysis, unit testing and deprecations checks
- Rego policy to enforce K8s best practices and security issues mitigation/prevention:
 - Continuous integration job ---> Trivy scanner
 - K8s API webhook ---> using OPA Gatekeeper
- Vulnerability/compliance scanning
- Minimal base images (updated at least every 24h)
- Image scan before push/after push/periodically (every day)
- Ossec + Falco
- Network Policy
- Log analysis&monitoring

Software security management

- Fully automated Continuous integration/delivery pipelines (target: SLSA level as higher as possible <https://slsa.dev>)
- Automated software security scanning and testing:
 - Decrease the risk of a security breach by automatically blocking known vulnerabilities
 - Critically malicious components and newly released suspicious components are automatically blocked
- Shift-left strategy:
 - Find and prevent defects early in the software delivery process
 - Dev team members are aware of the security constraints and best practices
 - Security by design (cross functional team)
- Software bills of materials (SBOMs):
 - Software “inventory”/ID
 - Generated automatically during the build phase
 - Release of problematic code automatically blocked
 - Enable continuous scanning of released software



Dependency track

“An intelligent Component Analysis platform that allows organizations to identify and reduce risk in the software supply chain.”

- Software bills of materials (SBOMs) as source of truth
- Vulnerabilities/Licenses analysis/management
- Product risk analysis
- Measure and enforce policy compliance

Dependency track



booking-backend ▾ latest

0 0 0 0 0

View Details >

Overview Components 339 Services 0 Dependency Graph 0 Audit Vulnerabilities 4 Policy Violations 0

Show suppressed findings

Search [] []

Component	Version	Group	Vulnerability	Severity	Analyzer	Attributed On	Analysis	Suppressed
commons-compress	1.20	org.apache.commons	NVD CVE-2021-36090	High	OSS Index	10 Nov 2021	Not Set	<input checked="" type="checkbox"/>

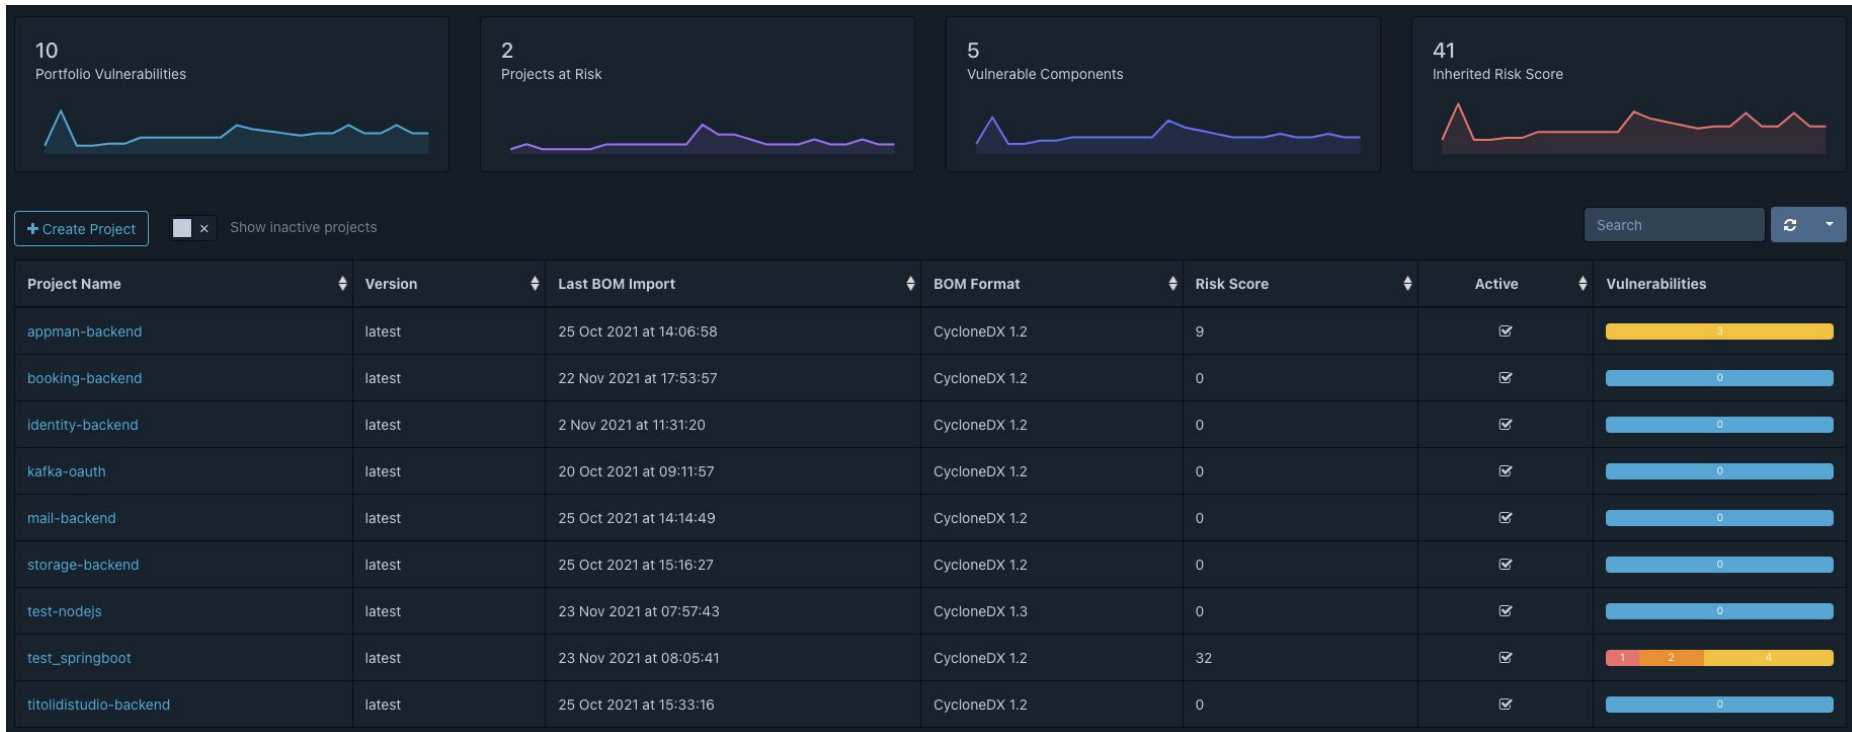
Description

When reading a specially crafted ZIP archive, Compress can be made to allocate large amounts of memory that finally leads to an out of memory error even for very small inputs. This could be used to mount a denial of service attack against services that use Compress' zip package.

Audit Trail

- 17 Nov 2021 at 18:57:29
- 17 Nov 2021 at 18:57:33
Suppressed

Dependency track



"New software" highlights

1. Better security aspects management
2. Improved automated software testing
3. Design for failure
4. Improved/simplified development and release workflow (no more long lived branches)
5. Limited amount of languages/frameworks
6. Database: MongoDB and PostgreSQL
7. File storage: no more NFS or local filesystem
8. Restricted database read-write access:
 - No more "real data" access required for software development
 - No more "physiological" data correction
 - No more manual schema migration
9. No more shared databases/schema
10. No more database raw data direct access

Any organization that design a system will inevitably produce a design whose structure is a copy of the organization's communication structure

(Conway's law, 1967)

