



Vera C. Rubin Observatory
Rubin Observatory Project Office

Hardening Kubernetes Workload

Gonzalo Seriche

ITTN-079

Latest Revision: 2024-12-19

DRAFT



Abstract

This technical note provides comprehensive guidance for implementing security controls and hardening measures for Kubernetes workloads based on CIS Benchmark recommendations. It covers pod security policies, network security, container image security, and compliance monitoring for on-premise Kubernetes deployments.

Draft

Change Record

Version	Date	Description	Owner name
1	2024-09-13	Initial release.	Gonzalo Seriche
1	2024-12-19	Add primary best practices and security guidelines	Gonzalo Seriche

Document source location: <https://github.com/lstt-it/ittn-079>

Draft

Contents

1 Introduction	1
2 Pod Security Standards	1
2.1 Pod Security Context Implementation	1
2.2 Container Level Security Controls	2
3 Network Security Framework	2
3.1 Default Network Policies	2
3.2 Granular Access Controls	3
4 Resource Management and Quotas	3
4.1 Resource Quota Implementation	4
4.2 Default Resource Constraints	4
5 Monitoring and Compliance Framework	5
5.1 Audit Logging Configuration	5
5.2 Compliance Monitoring Implementation	5
6 Security Best Practices	5
6.1 Image Security Controls	5
6.2 Secret Management	6
7 Implementation Guidelines	6
7.1 Deployment Process	6
7.2 Maintenance Procedures	7
8 CIS Benchmark Compliance Testing	7
8.1 Kube-bench Implementation	7
8.1.1 Version Compatibility	7
8.1.2 Automated Scheduling	8
8.2 Version-Specific Configurations	9
8.3 Results Processing	10

9 Compliance Monitoring and Alerting Framework	10
9.1 Advanced Results Processing	10
9.2 Monitoring Integration	11
9.3 Advanced Alerting Configuration	13
9.4 Compliance Reporting and Analytics	14
A References	15
B Acronyms	15

Draft

Hardening Kubernetes Workload

1 Introduction

This document outlines the security measures, implementation guidelines, and best practices for hardening Kubernetes workloads in accordance with the Center for Internet Security (CIS) Benchmark. The guidance provided here is specifically tailored for on-premise Kubernetes deployments and focuses on establishing a robust security posture while maintaining operational efficiency.

2 Pod Security Standards

2.1 Pod Security Context Implementation

Pod security context defines the privilege and access control settings that are essential for maintaining a secure Kubernetes environment. The following configurations establish the foundation for pod-level security:

```
1 apiVersion: v1
2 kind: Pod
3 metadata:
4   name: secure-pod
5 spec:
6   securityContext:
7     runAsNonRoot: true
8     runAsUser: 1000
9     fsGroup: 3000
10    seccompProfile:
11      type: RuntimeDefault
```

The security context implementation enforces several critical security controls:

- Non-root execution prevents privilege escalation scenarios

- Specific UID assignments ensure precise access control
- Filesystem group configurations maintain proper data access boundaries
- SecComp profiles limit available system calls to reduce the attack surface

2.2 Container Level Security Controls

Container security context configurations provide granular control over container execution parameters:

```
1 spec:
2   containers:
3     - name: secured-container
4       securityContext:
5         allowPrivilegeEscalation: false
6         capabilities:
7           drop: ["ALL"]
8         readOnlyRootFilesystem: true
9         privileged: false
```

These settings establish a defensive posture through:

- Prevention of privilege escalation attempts
- Removal of unnecessary Linux capabilities
- Implementation of read-only root filesystems
- Explicit prohibition of privileged execution modes

3 Network Security Framework

3.1 Default Network Policies

Implementation of a default deny policy establishes a zero-trust networking baseline:

```
1 apiVersion: networking.k8s.io/v1
2 kind: NetworkPolicy
3 metadata:
4   name: default-deny
5 spec:
6   podSelector: {}
7   policyTypes:
8   - Ingress
9   - Egress
```

3.2 Granular Access Controls

Specific allow rules should be implemented for required communications:

```
1 apiVersion: networking.k8s.io/v1
2 kind: NetworkPolicy
3 metadata:
4   name: allow-specific
5 spec:
6   podSelector:
7     matchLabels:
8       app: web
9   ingress:
10  - from:
11    - podSelector:
12      matchLabels:
13        role: frontend
14    ports:
15    - protocol: TCP
16      port: 80
```

4 Resource Management and Quotas

4.1 Resource Quota Implementation

Resource quotas prevent resource exhaustion and ensure fair resource allocation:

```
1 apiVersion: v1
2 kind: ResourceQuota
3 metadata:
4   name: compute-resources
5 spec:
6   hard:
7     requests.cpu: "4"
8     requests.memory: 8Gi
9     limits.cpu: "8"
10    limits.memory: 16Gi
```

4.2 Default Resource Constraints

Limit ranges establish default resource boundaries:

```
1 apiVersion: v1
2 kind: LimitRange
3 metadata:
4   name: resource-constraints
5 spec:
6   limits:
7     - default:
8       cpu: 500m
9       memory: 512Mi
10    defaultRequest:
11      cpu: 200m
12      memory: 256Mi
13    type: Container
```

5 Monitoring and Compliance Framework

5.1 Audit Logging Configuration

Comprehensive audit logging enables security monitoring and compliance verification:

```
1 apiVersion: audit.k8s.io/v1
2 kind: Policy
3 rules:
4 - level: Metadata
5   resources:
6     - group: ""
7       resources: ["pods", "services"]
8 - level: RequestResponse
9   resources:
10    - group: "apps"
11      resources: ["deployments"]
```

5.2 Compliance Monitoring Implementation

The compliance monitoring framework should encompass:

- Automated security scanning tools integration
- Continuous compliance status monitoring
- Policy enforcement automation
- Regular CIS Benchmark assessment procedures

6 Security Best Practices

6.1 Image Security Controls

Image security measures should include:

- Mandatory image signing and verification
- Regular vulnerability scanning
- Base image update procedures
- Registry access control implementation

6.2 Secret Management

Implement robust secret management practices:

```
1 apiVersion: v1
2 kind: Secret
3 metadata:
4   name: application-secrets
5   annotations:
6     vault.security.banzaicloud.io/vault-role: "app-role"
7 type: Opaque
8 data:
9   APP_SECRET: <base64-encoded-secret>
```

7 Implementation Guidelines

7.1 Deployment Process

The security control implementation process should follow these phases:

1. Security posture assessment
2. Gap analysis against CIS Benchmark
3. Control implementation prioritization
4. Gradual security measure deployment

5. Validation and testing
6. Continuous monitoring establishment

7.2 Maintenance Procedures

Establish regular maintenance procedures for:

- Security patch management
- Configuration updates
- Policy refinement
- Compliance verification

8 CIS Benchmark Compliance Testing

8.1 Kube-bench Implementation

Kube-bench provides automated testing for Kubernetes CIS Benchmark compliance. The implementation varies based on Kubernetes versions and deployment methods.

8.1.1 Version Compatibility

Kube-bench supports multiple Kubernetes releases with specific test definitions:

```
1 # Job specification for Kubernetes 1.24+
2 apiVersion: batch/v1
3 kind: Job
4 metadata:
5   name: kube-bench
6 spec:
7   template:
8     spec:
```

```
9     hostPID: true
10     containers:
11       - name: kube-bench
12         image: aquasec/kube-bench:v0.6.15
13         command: ["kube-bench", "run", "--targets", "master,node"]
14         volumeMounts:
15           - name: var-lib-kubelet
16             mountPath: /var/lib/kubelet
17           - name: etc-systemd
18             mountPath: /etc/systemd
19           - name: etc-kubernetes
20             mountPath: /etc/kubernetes
21         restartPolicy: Never
22     volumes:
23       - name: var-lib-kubelet
24         hostPath:
25           path: "/var/lib/kubelet"
26       - name: etc-systemd
27         hostPath:
28           path: "/etc/systemd"
29       - name: etc-kubernetes
30         hostPath:
31           path: "/etc/kubernetes"
```

8.1.2 Automated Scheduling

Implement regular compliance testing through CronJobs:

```
1  apiVersion: batch/v1
2  kind: CronJob
3  metadata:
4    name: kube-bench-schedule
5  spec:
6    schedule: "0 1 * * *" # Run daily at 1 AM
7    jobTemplate:
8      spec:
9        template:
10       spec:
```

```
11     hostPID: true
12     containers:
13       - name: kube-bench
14         image: aquasec/kube-bench:v0.6.15
15         command:
16           - "kube-bench"
17           - "run"
18           - "--targets"
19           - "master,node"
20           - "--outputfile"
21           - "/reports/compliance.json"
22         volumeMounts:
23           - name: var-lib-kubelet
24             mountPath: /var/lib/kubelet
25           - name: reports
26             mountPath: /reports
27     volumes:
28       - name: var-lib-kubelet
29         hostPath:
30           path: "/var/lib/kubelet"
31       - name: reports
32         persistentVolumeClaim:
33           claimName: compliance-reports-pvc
34     restartPolicy: Never
```

8.2 Version-Specific Configurations

Different Kubernetes versions require specific considerations:

For Kubernetes 1.24 and newer:

- Use kube-bench version 0.6.x
- Enable container runtime socket mounting
- Configure specific test suites for newer security features

For Kubernetes 1.23 and older:

- Use compatible kube-bench versions (0.5.x)
- Adjust for legacy API versions
- Include deprecated security controls testing

8.3 Results Processing

Implement automated results processing:

```
1 apiVersion: v1
2 kind: ConfigMap
3 metadata:
4   name: compliance-processor
5 data:
6   process.sh: |
7     #!/bin/bash
8     REPORT_PATH="/reports/compliance.json"
9     if [ -f "$REPORT_PATH" ]; then
10      # Process results
11      jq -r '.Controls[] | select(.status=="FAIL")' "$REPORT_PATH" > failures.
12      json
13      # Alert if necessary
14      if [ -s failures.json ]; then
15        # Send alerts
16        curl -X POST ${ALERT_ENDPOINT} -d @failures.json
17      fi
18    fi
```

9 Compliance Monitoring and Alerting Framework

9.1 Advanced Results Processing

The compliance results processing system implements a comprehensive pipeline for analyzing kube-bench outputs and generating actionable insights. The system processes both JSON

and YAML outputs, correlates findings across multiple clusters, and maintains a historical compliance database.

```
1 apiVersion: v1
2 kind: ConfigMap
3 metadata:
4   name: compliance-processor-config
5 data:
6   config.yaml: |
7     processing:
8       output_format: json
9       retention_days: 90
10      minimum_severity: MEDIUM
11      excluded_checks:
12        - 1.2.1 # Document exclusion rationale
13        - 2.1.3 # Alternative control implemented
14      alerting:
15        threshold_critical: 85
16        threshold_warning: 95
17        notification_channels:
18          - slack
19          - email
20          - pagerduty
21      reporting:
22        format: html
23        schedule: "0 8 * * 1" # Weekly reports
24        recipients: ["security-team@org.com"]
```

9.2 Monitoring Integration

The monitoring framework integrates with existing observability platforms through a comprehensive metrics pipeline:

```
1 apiVersion: monitoring.coreos.com/v1
2 kind: ServiceMonitor
3 metadata:
```



```
4   name: compliance-metrics
5 spec:
6   selector:
7     matchLabels:
8       app: kube-bench
9   endpoints:
10  - port: metrics
11    interval: 30s
12    path: /metrics
13  - port: compliance
14    interval: 5m
15    path: /compliance
16 ---
17 apiVersion: monitoring.coreos.com/v1
18 kind: PrometheusRule
19 metadata:
20   name: compliance-alerts
21 spec:
22   groups:
23   - name: compliance.rules
24     rules:
25     - alert: ComplianceScoreDegraded
26       expr: compliance_score < 85
27       for: 1h
28       labels:
29         severity: critical
30       annotations:
31         summary: Cluster compliance score below threshold
32         description: Cluster {{ $labels.cluster }} compliance score is {{
33           $value }}
34     - alert: HighSeverityFinding
35       expr: compliance_findings{severity="HIGH"} > 0
36       for: 5m
37       labels:
38         severity: critical
```

9.3 Advanced Alerting Configuration

The alerting system implements a sophisticated notification framework with different severity levels and escalation paths:

```
1  apiVersion: notification.toolkit.fluxcd.io/v1beta1
2  kind: Alert
3  metadata:
4    name: compliance-alerts
5  spec:
6    eventSeverity: info
7    eventSources:
8      - kind: Kustomization
9        name: '*'
10   providerRef:
11     name: slack
12 ---
13 apiVersion: notification.toolkit.fluxcd.io/v1beta1
14 kind: Provider
15 metadata:
16   name: slack
17 spec:
18   type: slack
19   channel: security-alerts
20   address: https://hooks.slack.com/services/YOUR-WEBHOOK-URL
21   secretRef:
22     name: slack-url
23 ---
24 apiVersion: v1
25 kind: ConfigMap
26 metadata:
27   name: alert-templates
28 data:
29   critical.tmpl: |
30     *Critical Compliance Alert*
31     Cluster: {{ .Cluster }}
32     Check: {{ .CheckID }}
33     Description: {{ .Description }}
34     Remediation: {{ .Remediation }}
```

```
35 warning.tpl: |
36   Warning: Compliance Issue Detected
37   Cluster: {{ .Cluster }}
38   Details: {{ .Details }}
```

9.4 Compliance Reporting and Analytics

The reporting system generates comprehensive compliance analytics and trend analysis:

```
1  apiVersion: batch/v1
2  kind: CronJob
3  metadata:
4    name: compliance-report-generator
5  spec:
6    schedule: "0 1 * * 1" # Weekly on Monday at 1 AM
7    jobTemplate:
8      spec:
9        template:
10         spec:
11           containers:
12             - name: report-generator
13               image: compliance-reporter:v1
14               env:
15                 - name: REPORT_PERIOD
16                   value: "7d"
17                 - name: INCLUDE_TRENDS
18                   value: "true"
19                 - name: REPORT_FORMAT
20                   value: "html,pdf"
21               volumeMounts:
22                 - name: report-output
23                   mountPath: /reports
24             volumes:
25               - name: report-output
26                 persistentVolumeClaim:
27                   claimName: report-storage
```

A References

B Acronyms

Acronym	Description
API	Application Programming Interface
CIS	Computer Infrastructure Support
ITTN	IT Technote
JSON	JavaScript Object Notation
PMO	Project Management Office
TCP	Transmission Control Protocol
UID	User Identifier
URL	Universal Resource Locator
YAML	Yet Another Markup Language