BLESSING PHIRI | TECHNICAL SOLUTIONS ENGINEER

Statement of Work

Design and Deployment of Six Bare-Metal Red Hat OpenShift Clusters with Centralised Logging

Prepared for: A regional financial services group, six countries

Reference	**SAMPLE-OCP-001**
Version	**1.2 (sample)**
Status	**Final**
Prepared by	**Blessing Phiri**

Portfolio sample

Based on a real statement of work. Client identity, people, pricing and commercial terms have been removed or changed. Shared to demonstrate solution design and scoping.

Document control

Item	Detail
Document title	Statement of Work: Design, Installation, Configuration and Deployment of Six Bare-Metal OpenShift Clusters with Centralised Logging Across Six Countries
Reference	SAMPLE-OCP-001
Status	Final
Summary	Professional services to build, harden, validate and hand over six production-grade bare-metal OpenShift clusters and a centralised ELK logging platform.
Estimated duration	Approximately 4 months elapsed

Contents

Section
1. Executive summary
2. Project objectives
3. High-level architecture
4. Scope of work
5. Testing and validation
6. Documentation and knowledge transfer
7. Estimated timeline
8. Delivery model and roles
9. Responsibilities and assumptions
10. Out of scope
11. Deliverables and acceptance criteria
12. Optional enhancements

1. Executive summary

This statement of work defines the professional services required to design, build, install, configure, harden, customise, validate and commission six production-grade bare-metal Red Hat OpenShift Container Platform clusters, one in each country where the client operates.

Every cluster is built to one standard blueprint, so operations, audits and upgrades are consistent across the region. Once the platforms are live, a centralised ELK (Elasticsearch, Logstash, Kibana) platform provides secure collection, indexing, visualisation, monitoring and long-term retention of application, infrastructure and audit logs from all six environments.

The engagement covers infrastructure readiness, networking, storage integration, security hardening, OpenShift installation, Day-2 customisation, logging deployment, testing, documentation and knowledge transfer, followed by a week of post-cutover handholding.

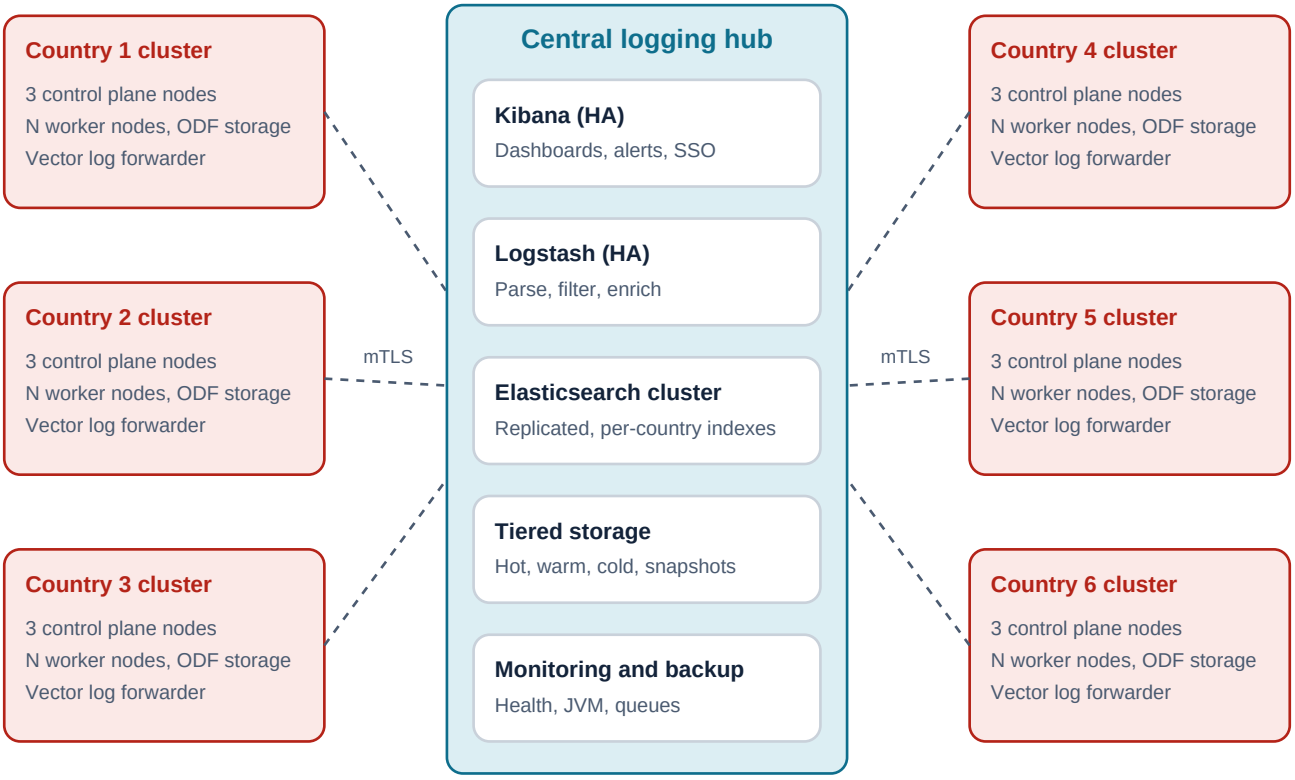
At a glance	
Clusters	6 production OpenShift clusters on bare metal, one per country
Logging	Centralised, highly available ELK hub with per-country indexes
Security	Hardened configuration, enterprise identity, mutual TLS log forwarding
Duration	Approximately 4 months, with phases running in parallel
Handover	Full documentation, 4-day admin training, 12-month support bundle

2. Project objectives

Platform	Operations
Six production-ready OpenShift clusters	Centralised ELK logging platform
Standardised cluster architecture	Secure, encrypted log forwarding
Highly available control planes	Monitoring and alerting
Secure enterprise networking	Operational documentation
Integrated enterprise storage and authentication	Knowledge transfer to the platform team

3. High-level architecture

A hub-and-spoke design: six identical OpenShift clusters at the edge, each forwarding logs over mutually authenticated TLS to a central, highly available logging hub.



Design principles

- Six independent clusters, one per country, for data residency and fault isolation
- Identical blueprint across all sites for consistent operations and upgrades
- Highly available control planes with three nodes per cluster
- Centralised, searchable application, infrastructure and audit logs
- Encrypted log transport with mutual authentication
- Role-based access and multi-tenant indexes per country
- Tiered storage and retention policies with snapshot-based recovery

4. Scope of work

The project runs in twelve implementation phases, followed by one week of post-production cutover handholding.

Phase 1: Discovery and planning	
Activities	Technical workshops, requirements gathering, infrastructure, application, network and security assessment, capacity and storage sizing, risk assessment
Deliverables	Discovery report, project plan, architecture design, risk register

Phase 2: Bare-metal infrastructure preparation (per country)	
Servers	Rack validation, BIOS and firmware upgrades, RAID and disk provisioning, boot order, BMC/IPMI, power redundancy and health checks
Network	NTP, DNS, DHCP where needed, load balancer, switches, VLANs, MTU and bonded interfaces
Validation	DNS resolution, latency, jumbo frames, routing and firewall rules

Phase 3: OpenShift bare-metal installation (per cluster)	
Planning	Cluster name, base domain, machine, service and cluster networks, CIDRs, VIPs and storage classes
Preparation	Pull secret, SSH keys, install-config.yaml, agent configuration and ignition files
Install	Bootstrap, control plane and worker nodes; certificates, API server, ingress, internal registry and monitoring
Validation	API, console, node health, cluster operators, etcd and networking

Phase 4: Post-installation customisation	
Identity	LDAP, Active Directory, SSO and identity providers
Storage	OpenShift Data Foundation where applicable, storage and snapshot classes, dynamic provisioning
Networking	Ingress controllers, routes, certificates, external load balancer, MetalLB where required
Security	RBAC, security context constraints, pod security, secrets and certificate management, TLS, image signing, registry security
Policies	Resource quotas, limit ranges, namespace standards, labels, taints, tolerations and network policies

Phase 5: Platform services	
Install	OpenShift GitOps, OpenShift Pipelines, OperatorHub operators, Prometheus, Alertmanager, Grafana and logging operators

Phase 6 and 7: ELK design and installation	
Design	Elasticsearch cluster, Kibana, Logstash, storage architecture, HA, backup, disaster recovery, capacity and retention
Elasticsearch	Nodes, replication, snapshots, TLS, authentication and index templates
Kibana	Authentication, spaces, dashboards, visualisation and alerts
Logstash	Pipelines, inputs, outputs, filters, parsing and performance tuning

Phase 8: OpenShift logging integration (per cluster)	
Configure	Cluster Logging Operator, ClusterLogForwarder and Vector collectors for audit, infrastructure and application logs
Secure	TLS with mutual authentication, log routing and multi-tenant indexes
Validate	Ingestion, search, dashboards and performance

Phase 9: Monitoring and operations	
OpenShift	Node health, cluster operators, CPU, memory, storage and networking
ELK	Cluster health, JVM, disk usage, pipeline health, index growth and queue depth
Alerts	Node failures, disk utilisation, certificate expiry, Elasticsearch, log forwarding and API failures

5. Testing and validation

Phase 10 proves the platform works before handover. Acceptance is tied to tests the client can witness.

Area	Tests
Infrastructure	Hardware, storage, network and redundancy
OpenShift	Node failover, API availability, ingress, storage and operators
Logging	Infrastructure, audit and application logs; dashboard and search validation
Security	Authentication, RBAC, TLS, certificates and audit trails
Disaster recovery	Node recovery, Elasticsearch recovery, backup restore and snapshot recovery

6. Documentation and knowledge transfer

Documentation (Phase 11)

Infrastructure build guide, OpenShift installation guide, cluster and ELK administration guides, operations manual, backup and disaster recovery procedures, SOPs, runbooks, network and architecture diagrams, configuration records and an asset register.

Knowledge transfer (Phase 12)

Administrator workshops on OpenShift and ELK operations, plus a customised four-day virtual OpenShift infrastructure administration course for up to six administrators, delivered as 40% theory and 60% hands-on labs.

Section	Modules
Architecture and operations	OpenShift fundamentals; installation and lifecycle management; node administration
Storage, networking and security	OVN-Kubernetes, routes and network policies; persistent storage and CSI; identity, RBAC and SCCs
Monitoring, logging and recovery	Prometheus and Alertmanager; logging and audit; etcd backup, restore and troubleshooting

7. Estimated timeline

Many phases run in parallel, with several engineers working simultaneously, giving an overall elapsed time of approximately four months.

Phase	Estimated duration
Discovery and planning	1 week
Bare-metal infrastructure preparation	3 weeks
OpenShift bare-metal installation	4 weeks
Post-installation customisation	2 weeks
Platform services	1 week
ELK platform design	1 week
ELK installation	1 week
OpenShift logging integration	1 week
Monitoring and operations	2 weeks
Testing	1 week
Documentation	2 weeks
Knowledge transfer	1 week
Post-production cutover handholding	1 week

A 12-month support bundle is included for monthly remote health checks, updates, upgrades and post-cutover assistance.

8. Delivery model and roles

Delivery uses a locally supported, continent-backed model. A certified Red Hat engineer is based in-country, working day to day with the client team, backed by senior Red Hat engineers who provide specialist support. Up to four engineers are expected on the project at any time, with active day-to-day involvement from the client team.

Role	Provided by
Project manager	Provider and client
Enterprise solution architect	Provider and client
OpenShift platform architect	Provider
OpenShift and Linux engineers	Provider
Network, storage and security engineers	Provider and client
ELK engineer	Provider
QA and test engineer	Provider and client
Technical writer	Provider

9. Responsibilities and assumptions

Client provides

- Data centre facilities, power, cooling, rack space and physical security
- WAN and LAN connectivity between all six sites and the central logging platform
- DNS, NTP, DHCP/PXE where applicable, load balancing and firewall changes
- Valid Red Hat OpenShift subscriptions and any required ELK licensing
- Administrative access to BMC/IPMI and network infrastructure
- TLS certificates, backup infrastructure, change approvals, maintenance windows and subject matter experts

Provider delivers

- Discovery, assessment and solution design
- Bare-metal and operating environment preparation
- Installation and customisation of six production OpenShift clusters
- Platform services, centralised ELK platform and log integration
- Monitoring, alerting, testing, documentation and knowledge transfer

Assumptions

- Servers are installed, powered and meet Red Hat hardware certification requirements
- Data centre facilities are operational before implementation
- WAN bandwidth between countries is sufficient for log forwarding
- Subscriptions and licences are available before deployment
- Maintenance windows and change approvals follow the agreed schedule

10. Out of scope

Excluded unless added through a change request:

- Application migration or modernisation
- Custom applications or logging pipelines beyond agreed requirements
- Integration with third-party SIEM platforms
- Ongoing managed services or 24x7 support after completion
- Procurement of hardware, software, networking or licences
- Disaster recovery site build or cross-cluster ELK replication
- Public cloud infrastructure

11. Deliverables and acceptance criteria

Design and build	Operations and handover
Infrastructure assessment report	Monitoring dashboards
High-level design (HLD)	Alerting configuration
Low-level design (LLD)	Complete documentation set
Six production OpenShift clusters	Knowledge transfer and training
Hardened OpenShift configuration	Final acceptance report
Centralised ELK platform and secure log forwarding	

The project is complete when:

- All six clusters are installed, customised and operational
- The centralised ELK platform is fully operational
- All six clusters forward logs, and application, infrastructure and audit logs are searchable
- Dashboards, monitoring and alerting are working, with TLS and authentication validated
- Infrastructure, OpenShift, logging, security and DR testing has passed
- Documentation and knowledge transfer are complete, and the client signs the acceptance certificate

12. Optional enhancements

For a mature multi-country deployment, the client may also consider:

- Elasticsearch high availability across multiple availability zones
- Hot, warm and cold storage tiers with Index Lifecycle Management
- Cross-cluster replication or snapshots for disaster recovery
- Alert routing to PagerDuty or ServiceNow
- Capacity planning based on daily log volume, retention and peak ingestion
- GitOps with Argo CD or OpenShift GitOps for repeatable configuration

About this sample. This document reproduces the structure and technical content of a real statement of work. The client, people, pricing, legal terms and contact details have been removed or changed. Prepared by Blessing Phiri as a portfolio sample.