elastic

Search. Observe. Protect.

Engagement Report

DGARM

Implementation Review

2026 May 8th

Salim Alaeddine

elastic.co

Table of Contents

Document Properties

Version History

Version	Date	Author	Description
1.0	2026 May 8th	Salm Alaeddine	Initial Release

References

Ref. No	Document Name	Location	Description
1	DGARM Design Document - 20250829 v1.2	Email	Design Document

Executive Summary

Overview

Elastic conducted an advisory review of DGARM's Elastic environment to assess the current cluster deployment, validate the implementation against the available design documentation, and identify any gaps, risks, or improvement opportunities.

The review focused on the deployed Elasticsearch and Kibana clusters, including cluster configuration, monitoring, alerting, snapshot and restore readiness, security controls, data lifecycle management, and operational alignment with Elastic best practices.

The engagement was structured around the following areas:

- Current state discovery.
- Cluster configuration review against the documented design.
- High-level health and operational review.
- Alignment to Elastic best practices.
- Monitoring, alerting, backup, and disaster recovery readiness.
- Identification of risks, findings, recommendations, and next steps.

Elastic guidance referenced during the review includes production readiness, dedicated monitoring, index lifecycle management, shard sizing, snapshot and restore, and dedicated master node design. Elastic recommends using ILM to automate rollover and retention, using snapshots as the supported backup method, and using a separate monitoring cluster for production monitoring where applicable.

Participants

Org	Name	Role/Position	mail
Tetra / Infosys	Biswajit Banerjee	Project Manager	biswajit@tetrain.com
	Mukul Mahajan		mukul@tetrain.com
	Muralidharan Subramanian	Lead Architect	muralidharan_S12@infosys.com
	Bhasha Sisodia	Operations Manager	bhasha.sisodia@infosys.com
	Karan Neelkanth	Platform Lead	karan.neelkanth@infosys.com

	Divya Dhyani		divya_dhyani@infosys.com
	Bimlesh Kumar		bimlesh.kumar@infosys.com
	Juhu Sajwar		juhu.sajwar@infosys.com
	Prashant Kashyap		prashant.kashyap@tetrain.com
	Manish Singh		Manish@tetrain.com
Elastic	Salim Alaeddeine	Consulting Architect	salim.alaeddine@elastic.co
	Sri Suba Selvachamy	Delivery Manager	sri.selvachamy@elastic.co

Business Value

The cluster review provides DGARM with a clear understanding of the current Elastic platform state and highlights the actions required to improve platform stability, resilience, governance, and operational maturity.

The key business value includes:

- Improved confidence in Elastic platform health and recoverability.
- Early identification of risks that may impact service availability.
- Better alignment between the deployed environment and the approved design.
- Improved operational visibility through monitoring, alerting, and audit logging.
- Reduced risk of data loss through validated snapshot and restore processes.
- Clearer prioritisation of remediation and optimisation activities.

Outcomes

The engagement produced the following key outcomes:

- Reviewed the current Elastic deployment at a high level.
- Compared the current environment against the available design documentation.
- Identified critical areas requiring immediate attention, including:
    - Production license status.
    - Snapshot repository functionality.
    - Backup and disaster recovery readiness.
    - Monitoring architecture.

    - Alert notification actions.
- Reviewed the current monitoring and alerting setup.
- Identified that Stack Monitoring alert issues related to API keys were resolved after updating the API key.
- Reviewed the current audit logging approach.
- Identified that Elasticsearch audit logs are currently consumed by Imperva but not ingested into Elastic.
- Captured recommended actions to improve alignment with Elastic best practices.

Next Steps

The recommended next steps are:

- Apply the correct production license for the Search Cluster before **8 May 2026**.
- Review and remediate the snapshot repository issue.
- Validate Snapshot Lifecycle Management policies and snapshot success history.
- Perform a controlled restore test to confirm recoverability.
- Move monitoring to a dedicated monitoring cluster where aligned with the design and governance requirements.
- Configure alert actions for critical Stack Monitoring alerts.
- Enable logs collection for Elasticsearch, Kibana, and relevant ingestion components.
- Review audit logging requirements and consider enabling Kibana audit logging.
- Back up non-Elasticsearch configuration artefacts.

Cluster Review

Cluster Architecture & Topology

The following clusters were reviewed or identified as part of the engagement. At the time of the review, no dedicated monitoring cluster was deployed.

Search Cluster

Node Role	Hostname	IP	CPU	RAM (GB)	Disk (GB)
Master + Data Node	dcplelkschmdh01	10.150.25.70	16	64	1024
Master + Data Node	dcplelkschmdh02	10.150.25.71	16	64	1024
Master + Data Node	dcplelkschmdh03	10.150.25.72	16	64	1024
Machine Learning Node	dcplelkschml01	10.150.25.73	8	64	200
Kibana Node	dcplelkschkib01	10.150.25.74	8	16	200
Kibana Node	dcplelkschkib02	10.150.25.75	8	16	200
Logstash Node	dcplelkschlgst01	10.150.25.76	8	16	300
Logstash Node	dcplelkschlgst02	10.150.25.77	8	16	300

Observability Cluster

Node Role	Hostname	IP	CPU	RAM (GB)	Disk (GB)
Master Node	dcplelkobsmas01	10.150.25.78	8	16	200
Master Node	dcplelkobsmas02	10.150.25.79	8	16	200
Master Node	dcplelkobsmas03	10.150.25.80	8	16	200
Hot Node	dcplelkobshtdt01	10.150.25.81	16	64	1024
Hot Node	dcplelkobshtdt02	10.150.25.82	16	64	1024
Warm Node	dcplelkobswmdt01	10.150.25.83	16	64	3072
Warm Node	dcplelkobswmdt02	10.150.25.84	16	64	3072
Warm Node	dcplelkobswmdt03	10.150.25.85	16	64	3072

Warm Node	dcplelkobswmdt04	10.150.25.86	16	64	3072
Machine Learning Node	dcplelkobsml01	10.150.25.87	16	32	200
Machine Learning Node	dcplelkobsml02	10.150.25.88	16	32	200
Kibana Node	dcplelkobskib01	10.150.25.89	8	16	200
Kibana Node	dcplelkobskib02	10.150.25.90	8	16	200
Fleet Server Node	dcplelkobsflt01	10.150.25.91	8	16	200
Fleet Server Node	dcplelkobsflt02	10.150.25.92	8	16	200
APM Node	dcplelkobsapm01	10.150.25.93	8	16	200
APM Node	dcplelkobsapm02	10.150.25.94	8	16	200
Logstash Node	dcplelkobslgst01	10.150.25.95	8	16	300
Logstash Node	dcplelkobslgst02	10.150.25.96	8	16	300

Elasticsearch Configuration Files and Bootstrap Checks

As part of the cluster review, the Elasticsearch configuration files were reviewed against Elastic production guidance and bootstrap requirements. The review identified several configuration items that should be cleaned up or hardened to improve operational clarity, production stability, and resilience.

Elasticsearch.yml

Finding	Recommendation	Elastic Reference
cluster.initial_master_nodes entry was found in the Elasticsearch configuration file. This setting is only required during the initial cluster bootstrap and should not remain in the configuration after the cluster has formed.	Remove the entry from the configuration file to avoid confusion during future maintenance, restarts, or node replacement activities.	Bootstrapping a cluster

bootstrap.memory_lock: true is currently commented out. This means Elasticsearch memory may be swapped to disk, which can cause long garbage collection pauses and node instability.	• Disable swap at the operating system level, set vm.swappiness = 1 where applicable, OR •Enable bootstrap.memory_lock: true, and validate that memory locking is active.	Disable swapping / Bootstrap checks
No shard allocation awareness configuration was found. node.attr.rack, node.attr.zone, or similar awareness attributes are not configured.	Configure shard allocation awareness based on the physical topology, such as rack, zone, or data centre, to help distribute primary and replica shards across failure domains.	Shard allocation awareness

Kibana.yml

Finding	Recommendation	Elastic reference
Plaintext password in kibana.yml - elasticsearch.password: ""	Remove from file. Add to Kibana keystore: bin/kibana-keystore add elasticsearch.password	Secure Settings
All three encryption keys exposed in plaintext - xpack.*.encryptionKey	Move all three keys to Kibana keystore.	Security Configuration
Audit logging not enabled - xpack.security.audit.enabled not configured	Enable xpack.security.audit.enabled: true	Audit Logging

Elastic Documentation reference:

- Elastic Production Guidance — Running Elastic Stack in production requires planning for resilience, performance, and scalability.
- Elasticsearch Important Settings — Production configuration items include cluster name, node name, network host, discovery, heap, logging, and backups.
- Elasticsearch Bootstrap Checks — In production mode, failed bootstrap checks prevent Elasticsearch from starting.

Version & Licence

The Elasticsearch clusters are currently deployed on version **9.3.1**.

The Observability cluster has the correct production licence applied, and no immediate action is required.

The Search cluster is currently running on a trial licence, which is due to expire on 8 May 2026. The trial licence should be replaced with the appropriate licence before the expiry date to avoid any loss of licensed features or potential service impact. The severity of this issue has been communicated to the relevant team, and it should be addressed promptly to avoid loss of licensed features or potential service impact.

The DR cluster currently has a full production licence applied. This should be reviewed, as the DR cluster was scoped to remain under the appropriate non-production licence unless it is actively being used during a disaster recovery scenario.

During normal operations, the production licence should remain assigned to the production cluster. In the event of a DR activation, the production licence can be moved to the DR cluster and then reverted back once the production environment has been restored.

Kibana Review

The Kibana review covered:

- Kibana availability and sizing.
- Spaces and saved object organisation.
- User access and role-based access control.
- External authentication integration.
- Role mappings.
- Alerting and connector configuration.
- Kibana audit logging.

Findings

Finding	Severity	Observation
Kibana audit logging is not enabled	Medium	Kibana audit logging is not currently enabled.

Alert actions are not configured	High	Critical monitoring alerts may not notify operational teams.
Role mappings require validation	Medium	Access model should be reviewed against the documented security design.

Recommendations

- Validate spaces and saved object organisation.
- Review RBAC design and role mappings.
- Enable Kibana audit logging if required by governance or compliance requirements.
- Configure alert actions for critical operational alerts.

Security and Access Control Review

The security review covered:

- Authentication model.
- External authentication or SSO integration.
- Role-based access control.
- Role mappings.
- TLS and certificate configuration.
- Audit logging.

Findings

Finding	Severity	Observation
Elasticsearch audit logs are not ingested into Elastic	Medium	Audit logs are consumed by Imperva but not available in Elastic for investigation or correlation.
Kibana audit logging is not enabled	Medium	Kibana user activity visibility is limited.
RBAC requires validation	Medium	Roles and mappings should be reviewed against the security design.

Certificate lifecycle management requires validation	Medium	TLS and certificate management should be reviewed.

Recommendations

- Review security design against the live implementation.
- Validate SSO or external authentication if available.
- Review role mappings against least privilege principles.
- Enable Kibana audit logging where required.
- Confirm whether Elasticsearch audit logs should be ingested into Elastic.
- Validate TLS configuration and certificate expiry management.

Snapshot, Backup, and Disaster Recovery Review

Current State

Snapshot configuration is not confirmed as functional. The snapshot repository appears to have issues that require further investigation. Backup and disaster recovery readiness therefore remain a critical risk area.

Elastic snapshots are the supported way to back up Elasticsearch data. Elastic explicitly states that filesystem-level copies of node data directories are not a supported backup and restore method.

Snapshot Lifecycle Management (SLM)

Cluster snapshots can be managed using Snapshot Lifecycle Management (SLM). Note this is different from the searchable snapshots which are managed as part of ILM actions.

When configuring long term snapshot retention, it is important to consider the frequency of the snapshots as well as the retention of a snapshot. Each snapshot has some overhead on the master nodes, so having an appropriate snapshot lifecycle management strategy is important for cluster stability.

The recommended strategy is to apply multiple SLM policies with varying intervals and retention to provide flexibility in both snapshot granularity as well as retention length. The following three policies could be used in conjunction:

Snapshot Policy	Frequency	Retention	Approximate Snapshot Count
Short-term snapshots	Every 15 minutes	24 hours	96
Daily snapshots	Daily	1 month	31
Monthly snapshots	Monthly	1 year	12

Combining the above three policies will ensure that snapshots are taken every 15 minutes, minimising the potential data loss and/or reingestion required if recent data needs to be restored. In addition, it reduces the overall number of snapshots required by taking snapshots at longer intervals. This process is described in detail in the following documentation.

Findings

Finding	Severity	Observation
Snapshot repository issue	Critical	Repository health requires investigation.
Snapshot success history not confirmed	Critical	Snapshot reliability is not yet validated.
Restore test not confirmed	Critical	Recoverability is not proven.
Non-Elasticsearch artefacts require separate backup	Medium	Elastic snapshots do not cover all deployment and configuration files.

Recommendations

- Review and fix the snapshot repository issue.
- Validate Snapshot Lifecycle Management policies.
- Confirm snapshot success and failure history.
- Perform a controlled restore test.
- Monitor snapshot failures and configure alert actions.

- Back up non-Elasticsearch configuration artefacts separately. Configuration artefacts to back up outside Elasticsearch snapshots include:
    - `elasticsearch.yml`
    - `kibana.yml`
    - Logstash pipeline files.
    - Operating system configuration.
    - Certificates and private keys.
    - Other deployment artefacts.

Data Lifecycle, Storage, and Retention Review

The review covered:

- ILM policy design.
- Rollover conditions.
- Retention periods.
- Data tier allocation.
- Shard sizing and shard count.

Elastic ILM is designed to automate the lifecycle of time-based indices, including rollover, retention, and deletion, helping optimise performance, reliability, and storage cost.

Elastic also provides shard sizing guidance to avoid excessive shard and index counts, including guidance around the number of indices per GB of heap on master nodes.

Therefore enforcing a strict shard strategy is required. The following guidelines should be followed for shard sizing:

1. Primary Shards should be between 10GB and 50GB in size.
2. In cases where logsdb data stream type is used, 10-30GB is an ideal size.
3. Shards should contain no more than 200 million documents.
4. No more than 3000 indices per 1 GB of master heap.

Findings

Finding	Severity	Observation

ILM alignment requires validation	High	ILM policies need to be compared with business retention requirements.
Shard sizing requires review	High	Shard count and shard sizing should be reviewed as per design
Storage tier alignment requires review	Medium	Data placement should align with data value, search frequency, and retention requirements.

Recommendations

- Review all ILM policies against business retention requirements.
- Confirm rollover conditions are appropriate for data volume.
- Validate delete phases to avoid uncontrolled data growth.
- Review data tiers and ensure data is stored on the correct tier.

Monitoring and Alerting Review

Current State

The current monitoring state appears to be self-monitoring for each cluster, with metrics collection only. Logs are not fully collected into Elastic for monitoring and troubleshooting.

Stack Monitoring alerting had an API key related issue, which was resolved by updating the API key. However, alert actions are not currently configured.

Elastic recommends using a separate monitoring cluster for production environments. This helps ensure monitoring data remains available if the production cluster is unavailable, and it also prevents monitoring activity from impacting production performance.

Findings

Finding	Severity	Observation
Self-monitoring is currently used	High	Monitoring data may be unavailable if the production cluster is impacted.

Logs are not fully collected	High	Troubleshooting visibility is limited.
Alert actions are not configured	High	Critical alerts may not notify administrators.
API key issue was resolved	Closed / Informational	Updating the API key fixed the Stack Monitoring alert issue.

Recommendations

- Move monitoring to a dedicated monitoring cluster in line with the design documentation.
- If a dedicated monitoring cluster is not feasible, assess AutoOps suitability based on governance and compliance requirements.
- Collect both metrics and logs for Elasticsearch, Kibana, Logstash, Elastic Agent, and Beats where applicable.
- Configure notification actions for critical alerts.
- Recommended alert notifications actions could include:
    - Email.
    - Ticketing platform.
    - Microsoft Teams.
    - Incident management tooling.

Recommended alert coverage:

- Cluster health.
- Node availability.
- Disk usage and storage thresholds.
- JVM memory pressure.
- CPU utilisation.
- Shard allocation issues.
- Snapshot failures.
- Ingestion failures.
- Kibana Task Manager issues.

Elastic Stack Monitoring includes preconfigured alerting rules based on Elastic recommendations, which can be tailored to the environment.

Data Source Inventory & Integration Status

The ingestion review covered the current and planned data sources across servers, databases, platforms, applications, storage, network, and security devices.

Elastic out-of-the-box integrations are being used where available, including System, PostgreSQL, IBM MQ, GitHub, Tenable, Palo Alto, Fortinet, Juniper SRX, Aruba CX, F5 BIG-IP, Forcepoint, and Kubernetes/OpenShift where applicable. Elastic integrations provide pre-packaged assets to collect, store, and visualise data from supported sources.

For data sources where no suitable out-of-the-box integration is available, custom ingestion is required using supported ingestion patterns such as syslog, Logstash, API polling, HTTP Endpoint,, or custom Elastic Agent integrations.

A quick overview on the data onboarding progress:

Data Source	Type	OOTB Integration	Current Status / Notes
RHEL Servers	Server	System Logs & Metrics	Onboarded using Elastic System integration for logs and metrics. ~210 host onboarded at this stage
Physical Servers	Server	System Logs & Metrics	
Kubernetes / OpenShift	Platform	Kubernetes / OpenShift integration	Onboarded, but some metrics are not being collected due to OpenShift SCC/context permissions. Requires follow-up.
eMudhra SecurePass / iDAM	Application	N/A	eMudhra Host is being monitored at this stage. Application logs are not onboarded yet.
PostgreSQL	Database	PostgreSQL	Onboarded using OOTB integration.
IBM MQ	Application	IBM MQ	OOTB integration used. Running through some challenges, support case follow up
Neo4J	Application	N/A	Not yet onboarded

NGINX + API Gateway	Application	N/A	Onboarded. Need to verify `sub_status` parsing/visibility.
JBoss Web Server	Server / Application	N/A	Not yet onboarded
GitHub	Application	GitHub	Not yet onboarded
ManageEngine ITSM	Application	N/A	Not yet onboarded
ManageEngine Patch	Application	N/A	Not yet onboarded
Tenable	Application	Tenable	Not yet onboarded
Imperva	Application / Security	Imperva	Not yet onboarded
Cisco SAN Switch / Storage	Storage	Cisco Nexus, subject to validation	Logs : Cisco Nexus integration deployed. Might need to validate if devices are Nexus OS or IOS and use the appropriate integration (Cisco IOS, Cisco Nexus) Metrics: Can use SNMP polling t o get devices health, and metrics.
NetApp Storage	Storage	N/A	SNMP polling is currently used to collect device metrics. This may need to be fine-tuned to ensure all required metrics are covered. Alternative options include using NetApp REST APIs for metrics collection or reviewing NetApp Harvest as a potential metrics collection approach.
Veritas Backup	Backup	N/A	Not yet onboarded
Palo Alto Firewall — Internal	Security	Palo Alto Next-Gen Firewall	Onboarded Logs. Metrics are not yet configured.
Fortinet Firewall — External	Security	Fortinet FortiGate Firewall Logs	Logs are working. Metrics are not yet configured.
Fortinet Sandbox	Security	Fortinet FortiGate Firewall Logs, subject to validation	Not yet onboarded
Juniper Router	Network	Juniper SRX	OOTB integration is deployed. Review the incoming log patterns to confirm alignment with the default integration parsing

			patterns. If the logs do not match the default patterns, create a custom pattern and add it to the ingest pipeline.
Aruba Switches	Network	HPE Aruba CX	Metrics onboarded using SNMP poling
WAF	Security	F5 BIG-IP	Not yet onboarded
DDoS	Security	F5 BIG-IP	Not yet onboarded
Proxy	Security	Forcepoint Web Security	Not working. Requires follow-up.
DLP	Security	Forcepoint Web Security, subject to validation	Not working. Requires follow-up.
F5 Load Balancer	Security / Network	F5 BIG-IP	Data is not coming in. Current issue appears to be that data is being sent as syslog rather than to the expected HTTP endpoint.

Based on the current setup, Elastic is receiving syslog/log data from the devices through the available integrations, but metrics collection is not yet enabled or available for some device types through the same integrations. This is common, as many network, security, and storage integrations are primarily log-focused, while infrastructure and performance metrics may require additional collection methods.

Elastic can provide recommendations around possible ingestion patterns, Elastic-supported integrations, data modelling, dashboarding, alerting, and best-practice alignment. However, the implementation, support, and maintenance of any custom integrations, vendor-specific scripts, third-party tools, exporters, SNMP configurations, API collectors, or non-Elastic-supported components would remain the responsibility of the client or the relevant third-party vendor.

At a high level, the recommended approach would be a hybrid model:

- Continue using Elastic integrations and syslog for logs and security events.
- Use SNMP polling for standard device health and infrastructure metrics where suitable.
- Use vendor APIs where richer platform-specific metrics are required.
- Use Logstash or Elastic Agent/custom integrations where transformation, enrichment, or custom collection is required.

For the listed devices, the possible direction would be:

Storage Devices

For Cisco SAN switches/storage, SNMP polling may be considered for visibility into interface status, throughput, errors, discards, power supply, fan status, temperature, and general hardware health.

For NetApp Storage, metrics can be collected using SNMP polling, NetApp REST APIs, or by exploring NetApp Harvest. SNMP polling may be used for device-level metrics, but it may need to be fine-tuned to ensure all required metrics are covered. NetApp REST API-based collection may provide richer visibility into areas such as volume capacity, aggregate usage, IOPS, latency, controller health, disk/shelf health, and replication status.

For Veritas Backup, API-based collection or scheduled export through Logstash may be considered to monitor backup job status, failed jobs, success rate, backup duration, storage pool usage, and SLA compliance, depending on the available Veritas interfaces.

Security Devices

For Palo Alto Firewalls, the available Elastic Palo Alto integration capabilities should be reviewed to confirm whether metrics collection can be enabled for the PAN-OS environment. Elastic provides a Palo Alto Networks metrics integration that periodically fetches metrics from Palo Alto Networks firewalls and management systems.

For Fortinet Firewalls, SNMP polling may typically be used for metrics such as CPU, memory, interface status, bandwidth, session count, HA status, and hardware health, while syslog continues to provide security and traffic events.

For Sandbox, WAF, DDoS, Proxy, DLP, and F5 Load Balancer platforms, the preferred collection method depends on the vendor, model, supported interfaces, and required metrics. In many cases, SNMP can provide baseline infrastructure health, while vendor APIs may provide richer operational metrics such as policy events, blocked requests, queue depth, backend health, pool member status, throughput, latency, and service availability.

Network Devices

For Aruba Switches and Juniper Routers, SNMP polling is commonly used for metrics collection. This may cover device availability, CPU, memory, interface utilisation, interface errors/discards, port

status, temperature, fan/power status, and routing or protocol health where exposed by the device MIBs.

The initial baseline metrics to consider may include:

- Device availability
- CPU utilisation
- Memory utilisation
- Interface status and utilisation
- Interface errors and discards
- Hardware health
- Storage capacity
- Backup success/failure status
- HA/failover status
- Last successful metric collection timestamp

From an Elastic perspective, the metrics should ideally be normalised into dedicated metrics data streams such as network, security, storage, and backup metrics.

For Logstash-based onboarding, it is recommended to use Elastic data streams rather than standalone indices where possible.

Data streams are better suited for time-series data such as logs, metrics, and observability events. They also integrate well with index lifecycle management, rollover, and long-term retention strategies.

Recommended approach:

- Use data streams for metrics and logs.
- Apply consistent dataset and namespace naming.
- Align the data stream naming convention with the design documentation.
- Apply ILM policies based on data type, value, and retention requirements.

Example pattern:

metrics-network.snmp-prod
metrics-storage.netapp-prod
logs-network.device-prod

For SNMP and custom metrics, it is recommended to map important fields to Elastic Common Schema (ECS) where practical.

For example:

Current / Custom Field	Recommended ECS-Aligned Field	Notes
tag.device	host.hostname or host.name	Use keyword mapping for filtering and grouping.
Device IP field	host.ip	Use IP mapping where possible.
Device type	host.type or custom device.type	Helps classify routers, switches, storage, etc.
Interface name	observer.ingress.interface.name, observer.egress.interface.name, or custom interface field	Use consistent naming for network interface dashboards.
Interface status	event.status or custom mapped keyword field	Enables consistent filtering and visualisation.
SNMP source	data_stream.dataset	Helps organise and route data.

The exact mapping should be reviewed based on the source data structure and the dashboards required.

To support dashboard filtering and aggregations, key fields should be mapped correctly.

Recommended examples:

- Hostnames and device names: keyword
- IP addresses: ip
- Interface names: keyword
- Status fields: keyword
- CPU, memory, disk, and bandwidth metrics: numeric fields, for example long, double.
- Timestamps: date

For reference, the following documentation pages may be useful:

- Elastic Agent and Fleet overview: https://www.elastic.co/docs/reference/fleet
- Manage Elastic Agent integrations: https://www.elastic.co/docs/reference/fleet/manage-integrations
- Elastic integrations catalogue: https://www.elastic.co/integrations
- Logstash SNMP integration plugin, including SNMP polling and SNMP trap support: https://www.elastic.co/docs/reference/logstash/plugins/plugins-integrations-snmp
- Logstash SNMP input plugin: https://www.elastic.co/docs/reference/logstash/plugins/plugins-inputs-snmp
- Palo Alto Networks logs integration for Elastic: https://www.elastic.co/docs/reference/integrations/panw
- Palo Alto Networks metrics integration for Elastic: https://www.elastic.co/docs/reference/integrations/panw_metrics

Out-of-Scope Dashboard Assistance

Although dashboard creation was outside the original review and validation scope, the team raised an urgent request for assistance with creating an overall monitoring dashboard. In response, we supported the team remotely via screen sharing and helped structure the dashboard across server, network, and storage monitoring sections.

The server monitoring section was completed where data is available, including CPU, memory, disk utilisation, and top utilisation views. The network monitoring section was structured, but current data availability is limited to Aruba device interface status. Network CPU and memory metrics, other network device types, and additional models are not yet onboarded into Elastic. The storage section was also structured and populated using the available NetApp SNMP data, currently limited to device status and global health.

The overall monitoring dashboard structure was created to cover the main requested areas:

- Server monitoring.
- Network monitoring.

- Storage monitoring.

The dashboard structure now provides a foundation for the required monitoring use cases, with visualisations populated where the relevant data is currently available.

Where data is not available, the required sections were structured so that they can be completed once the missing data sources and metrics are onboarded into Elastic.

Screenshots of the created dashboard:

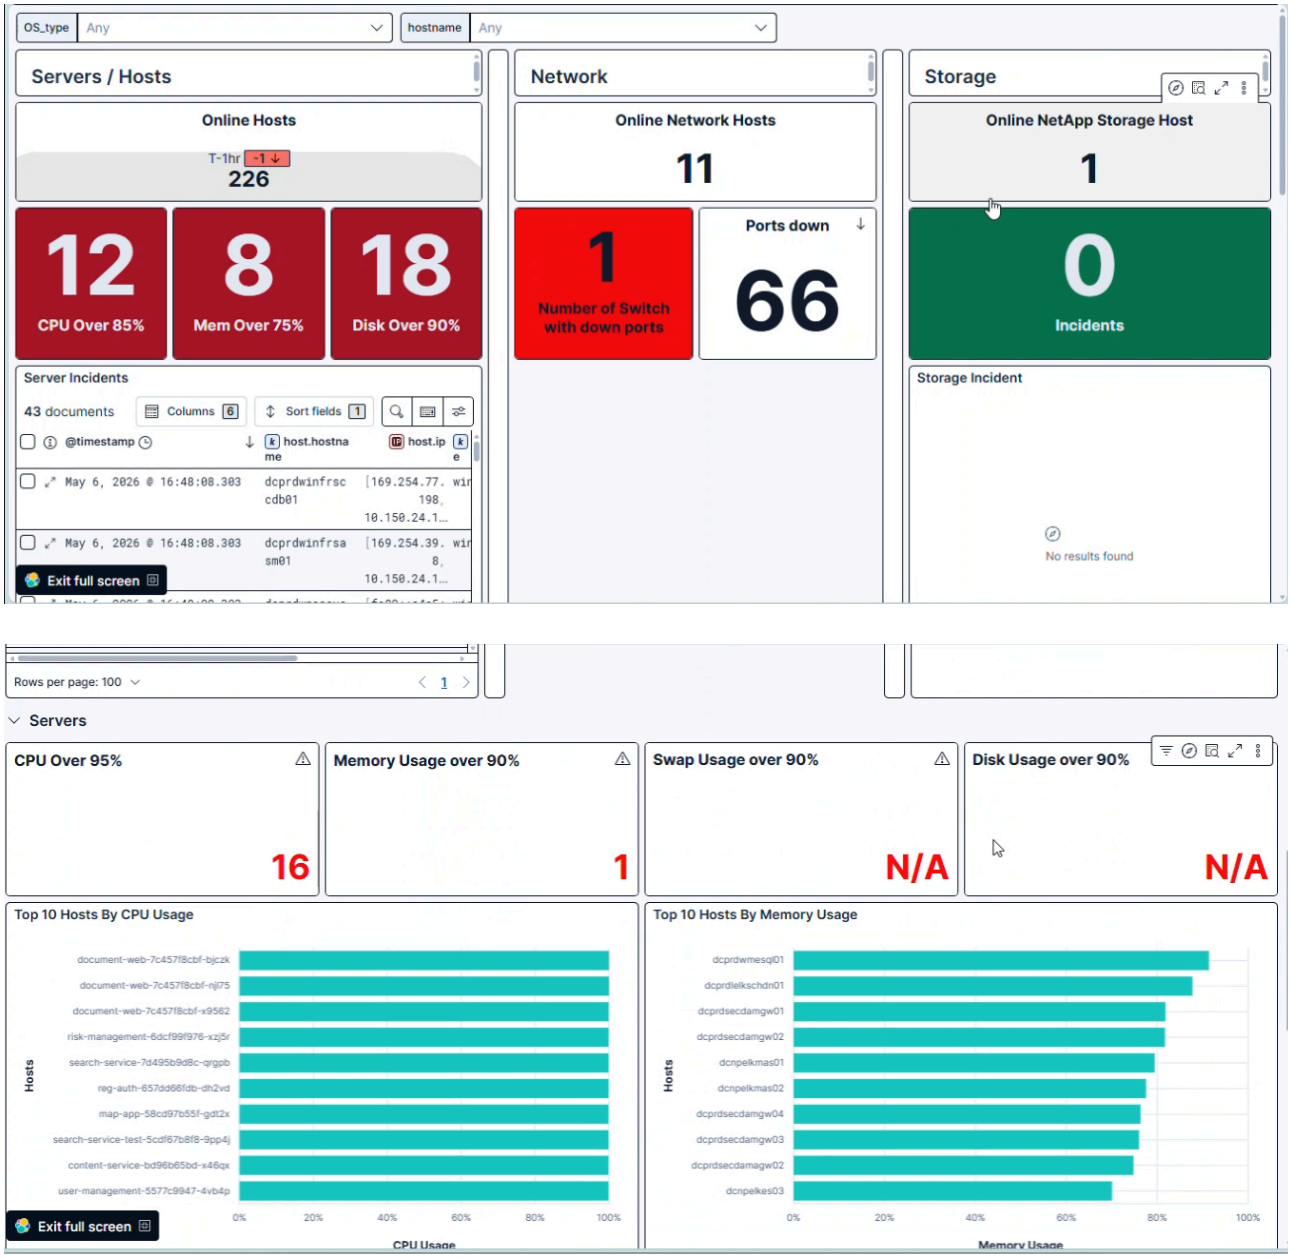

Host	Host IP	CPU usage	Memory usage	Swap Usage	Disk Usage
document-web-7c457f8cbf-hjl	10.150.24.58	99.99%	(null)	(null)	(null)
map-app-58cd97b55f-gdt2x	10.150.24.70	99.89%	(null)	(null)	(null)
reg-auth-657dd66fdb-dh2vd	10.150.24.70	99.88%	(null)	(null)	(null)
search-service-test-5cdf67b8f	10.150.24.70	99.88%	(null)	(null)	(null)
risk-management-6dcf99f976-	10.150.24.70	99.88%	(null)	(null)	(null)
search-service-7d495b9d8c-q	10.150.24.70	99.87%	(null)	(null)	(null)
content-service-bd96b65bd-x4	10.150.24.70	99.80%	(null)	(null)	(null)
watchlist-58785c96f-6b29c	10.150.24.70	99.78%	(null)	(null)	(null)
user-management-5577c9947-	10.150.24.70	99.78%	(null)	(null)	(null)
content-service-bd96b65bd-vn	10.150.24.76	99.70%	(null)	(null)	(null)
dcnpluatapigw02	10.150.20.195	66.20%	31.00%	3.9%	10%
dcproduaclsmgmt02.dgarm.gov	10.150.24.202	49.40%	31.80%	1.5%	9%
dcplelkschlgst02	10.150.25.77	47.90%	11.40%	0%	8%
dcnpdlphyez03	10.150.20.116	47.40%	46.20%	99.6%	37%

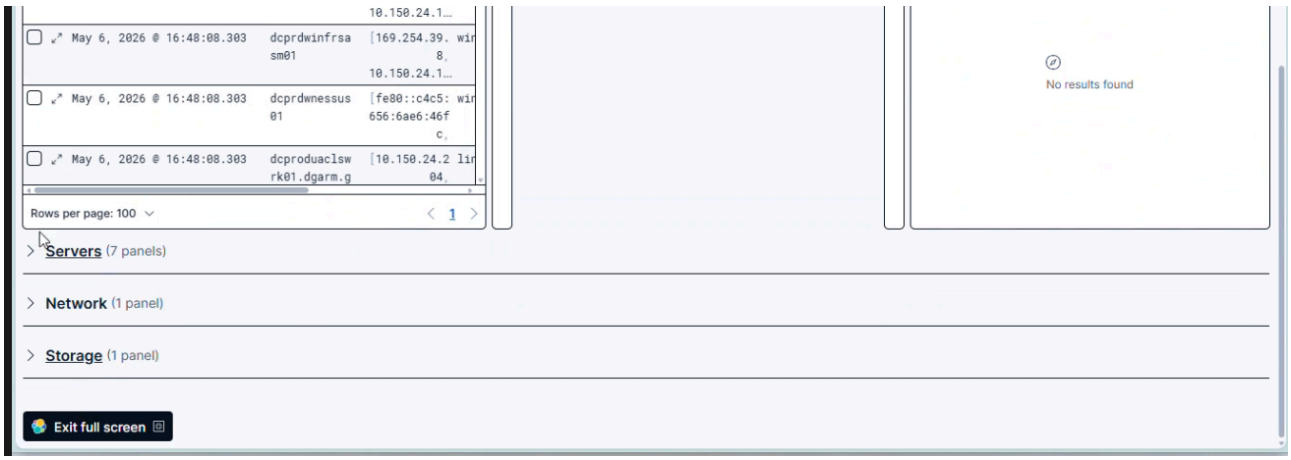

Action Plan - Next Steps

Priority	Action	Owner	Target
Critical	Apply correct production license before **8 May 2026**.	DGARM	Immediate
Critical	Review and fix snapshot repository issue.	DGARM	Immediate
Critical	Validate snapshot policies and snapshot success history.	DGARM	Immediate
Critical	Perform controlled restore test.	DGARM	Immediate
High	Review monitoring architecture against design.	DGARM	Short term
High	Configure alert notification actions.	DGARM	Short term
High	Enable log collection for monitoring.	DGARM	Short term
High	Review ILM, shard sizing, and retention.	DGARM	Short term
Medium	Review and enable Kibana audit logging if required.	DGARM	Medium term
Medium	Confirm whether Elasticsearch audit logs should be ingested into Elastic.	DGARM	Medium term
Medium	Back up configuration artefacts outside Elasticsearch snapshots.	DGARM	Medium term

