White paper

Datadog CloudPrem Professional Services

Accelerating enterprise troubleshooting with sovereign, full-stack observability

Prepared by
Interkey

Date
September 2026

From Vision to Reality

Executive summary

Interkey offers a CloudPrem-first professional services practice that deploys the full Datadog platform while keeping log data inside the customer's own infrastructure. The goal is one measurable outcome: enterprise teams find and fix application problems faster, without moving regulated log data out of the Kingdom or the customer's cloud tenancy.

Datadog CloudPrem, now marketed as BYOC (Bring Your Own Cloud) Logs, runs Datadog's log engine in the customer's Kubernetes cluster and object storage. Datadog's own architecture notes state that no log data leaves the customer environment during ingestion, and only query results travel to the Datadog UI (architecture). Engineers still get Log Explorer, monitors, dashboards and correlation with APM traces and metrics in one screen.

That hybrid model solves the problem that blocks most GCC enterprises from adopting SaaS observability: logs carry personal data, payment data and security events that regulators expect to stay in-country. CloudPrem removes that blocker for logs; Interkey's services cover everything around it, the Agent fleet, APM, Database Monitoring, Network Device Monitoring, Synthetics, RUM, Cloud SIEM, Workflow Automation, Incident Management and the governance that ties them together.

The practice is delivered in four phases, Assess, Build, Correlate, Operate, and packaged as fixed-scope launch offers, module add-ons, and a managed observability service. Success is measured in mean time to detect (MTTD), mean time to resolve (MTTR), alert noise and log cost per GB, baselined in week one and reported monthly.

What the customer gets	How
Logs stay in customer-controlled storage	CloudPrem / BYOC Logs on customer Kubernetes + S3-compatible object storage
One place to troubleshoot	Unified service tagging across logs, traces, metrics, DB, network and user sessions
Controlled cost	Observability Pipelines filtering, routing and redaction before indexing
Faster root cause	Service maps, trace-to-log pivots, Watchdog and Bits AI SRE investigation

What the customer gets	How
Audit-ready governance	RBAC, Audit Trail, Sensitive Data Scanner, mapped to PDPL, NCA ECC and SAMA controls

Why enterprise troubleshooting is slow today

Most resolution time is spent finding the problem, not fixing it. In a typical GCC enterprise the evidence for one incident is spread across five or six tools owned by different teams, and nobody can join them quickly.

- **Tool sprawl.** Infrastructure in one tool, application logs in an ELK or Splunk cluster, database stats in vendor consoles, network in an SNMP poller, user complaints in the service desk. Each hop is a context switch and often a ticket to another team.
- **No shared identity.** Logs, traces and metrics use different host names, service names and environment labels, so a trace ID in one tool cannot be pasted into another.
- **War-room troubleshooting.** Without a service map, teams debate which tier is at fault. Network blames the application, the application blames the database.
- **Log volume and cost.** Self-hosted log clusters grow faster than budgets. Teams respond by cutting retention or dropping debug logs, exactly the data needed during an incident.
- **Data residency blocks SaaS.** PDPL, NCA ECC and SAMA expectations make security and customer-data logs hard to send to an overseas SaaS region, so the most useful telemetry stays in legacy silos.
- **Alert noise.** Static thresholds on hundreds of hosts page people for symptoms, not causes, and real signals get lost.

CloudPrem changes the residency equation for logs, and the rest of the Datadog platform removes the tool and identity gaps, but only when it is designed and rolled out deliberately. That design and rollout is what Interkey's professional services provide.

The Datadog platform and where CloudPrem fits

CloudPrem is the log tier of the platform; every other module still reports to Datadog SaaS and is joined to CloudPrem logs through shared tags in the same UI. Designing which data goes where is the first architectural decision in every engagement.

What CloudPrem / BYOC Logs is

CloudPrem is a self-hosted Datadog log engine. It indexes, stores and searches logs in the customer's Kubernetes cluster and object storage (Amazon S3, Azure Blob, Google Cloud Storage or S3-compatible storage on-premises), while Datadog SaaS hosts the UI and acts as the control plane (overview). Datadog positions it for petabyte-scale log volumes, data residency and cost control using the customer's own cloud discounts (Datadog blog).

The product was announced as CloudPrem in 2025 and is now documented as BYOC Logs; this paper uses both names. It is not available on the US government Datadog sites (app.ddog-gov.com, us2.ddog-gov.com).

Module map: what runs where

Module	Where data is stored	Primary troubleshooting role
Log Management via CloudPrem	Customer object storage, customer Kubernetes	Evidence: errors, stack traces, audit and security events
Observability Pipelines	Worker in customer network; routes data	Filter, redact, enrich and route logs before indexing
Infrastructure Monitoring	Datadog SaaS	Host, container, Kubernetes and cloud resource health
APM and Distributed Tracing	Datadog SaaS (sampled traces)	Service map, latency, error traces, trace-to-log pivot
Database Monitoring (DBM)	Datadog SaaS	Query plans, wait events, blocking and slow queries
Network Device Monitoring (NDM) and Network Path	Datadog SaaS	SNMP devices, interfaces, topology, path latency
Synthetic Monitoring	Datadog SaaS	Proactive API and browser tests, private locations inside the network
Real User Monitoring (RUM) and Session Replay	Datadog SaaS	What end users actually experienced
Cloud SIEM and Sensitive Data Scanner	Detection in Datadog; log content governed by residency design	Security detections and PII controls

Module	Where data is stored	Primary troubleshooting role
Incident Management, On-Call, Workflow Automation	Datadog SaaS	Response coordination and automated remediation
Watchdog and Bits AI SRE	Datadog SaaS	Anomaly detection and AI-assisted investigation

Feature parity between CloudPrem and SaaS-hosted logs is still evolving release by release. Interkey validates each required feature, for example a specific log monitor type, Cloud SIEM rule or archive behaviour, against the current release notes during the Assess phase rather than assuming parity.

Reference architecture

Log data is collected, processed, indexed and stored entirely inside the customer boundary; the only link for logs is an outbound connection that carries queries in and results out.

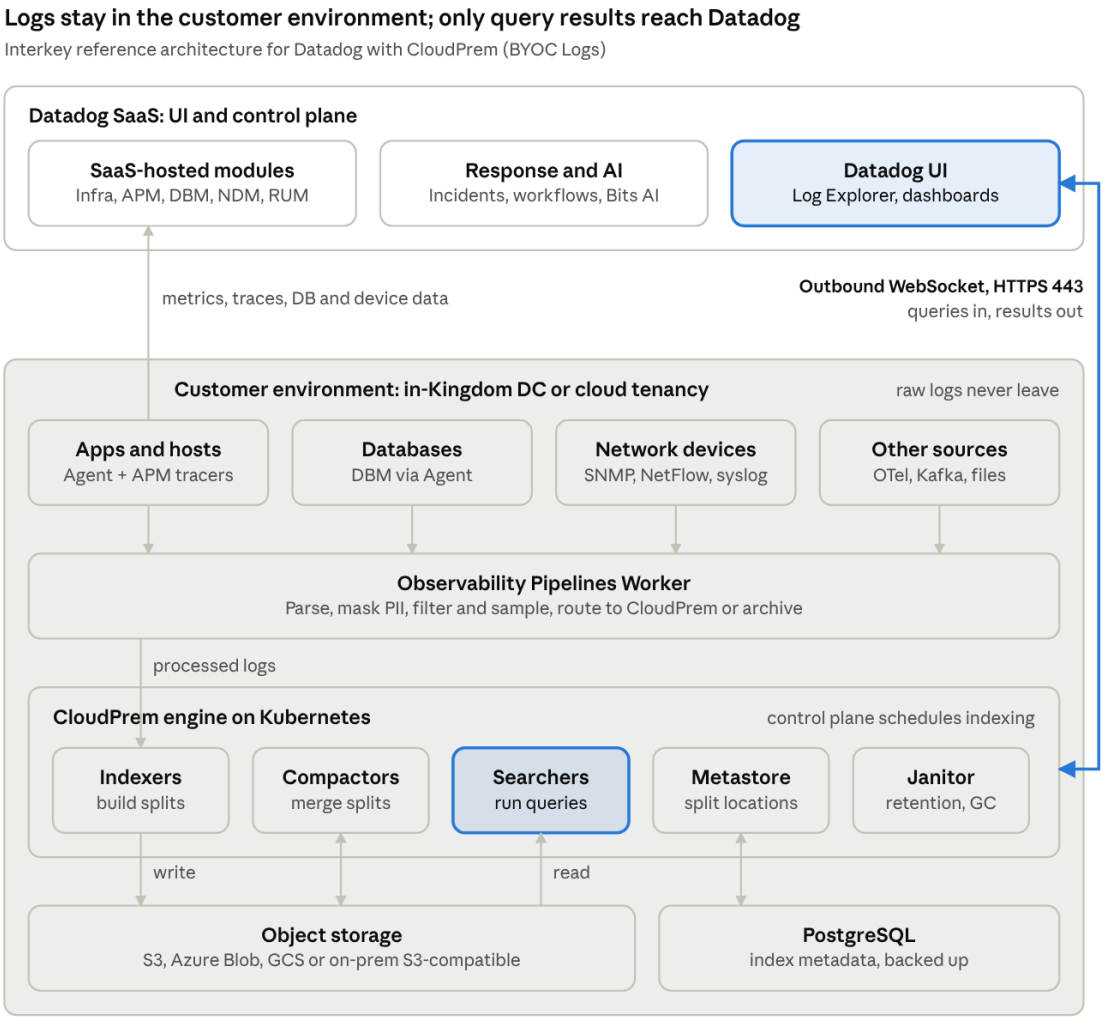


Figure 1. Reference architecture · CloudPrem in the customer environment, Datadog SaaS as UI and control plane

Agents and APM tracers send metrics, traces, database and device telemetry to Datadog SaaS as usual. Logs take a separate path: Observability Pipelines masks and filters them, CloudPrem indexers write them to the customer's object storage, and searchers answer queries from the Datadog UI over the reverse connection. Engineers see both in one interface, joined by shared tags.

System design: CloudPrem engine

CloudPrem separates compute from storage: stateless Kubernetes pods do the work, object storage holds the data, and PostgreSQL holds the index metadata. That separation is what makes it cheap to scale and simple to recover.

Engine components

Component	Role	Design notes from Interkey practice
Observability Pipelines Worker	Receives logs from Agents and other sources, parses, filters, redacts and routes	Deploy at least 2 replicas per zone; this is where PII masking and volume reduction happen
Indexers	Build index files ("splits") and write them to object storage	Need a persistent volume for the write-ahead log; Datadog recommends at least 250 GB per indexer pod on network-attached block storage
Compactors	Merge small splits into larger ones	Standalone service since v0.1.32; schedule on separate nodes from searchers
Searchers	Run queries against object storage	Stateless; size for aggregation-heavy dashboards; keep at least one running or the Datadog UI loses its connection
Metastore + PostgreSQL	Tracks where each split lives	Managed PostgreSQL (RDS, Cloud SQL, Azure Database) or HA Postgres on-prem; read replicas supported since v0.1.32
Control plane	Schedules indexing work	Low resource, single logical instance
Janitor	Retention, garbage collection, deletes	Enforces per-index retention that matches the customer's records policy
Object storage	Holds all log data	S3, Azure Blob, GCS or on-prem S3-compatible (for example MinIO, Ceph, NetApp StorageGRID)

Sources: architecture, production best practices, release notes.

Network model

By default the searchers open an outbound WebSocket over HTTPS 443 to the customer's Datadog site; no inbound ports or public DNS are needed (network). Queries come down that channel and only results go back. An alternative public-ingress mode uses gRPC with mTLS. Standard proxy variables (HTTPS_PROXY, NO_PROXY) are supported, which suits enterprises that force all egress through a proxy.

Sizing approach

Datadog sizes pods through a podSize Helm value; published examples are an xlarge indexer at 4 vCPU / 16 GB RAM and a 2xlarge searcher at 8 vCPU / 32 GB RAM. Interkey sizes each deployment from three inputs measured in the Assess phase:

1. **Peak ingest (GB/hour)** after Observability Pipelines filtering, drives indexer count.
2. **Query profile**: term searches are light, multi-day aggregations and log-based dashboards are heavy, drives searcher count and size.
3. **Retention (days) × daily volume × compression**: drives object storage capacity.

The result is a sizing workbook with node pools, persistent volume sizes, PostgreSQL class and storage capacity, validated in a load test before go-live.

High availability and disaster recovery

- **Zone spread.** Indexers and searchers are spread across availability zones (zone-aware scheduling is on by default since v0.1.34).
- **Stateless recovery.** Losing a searcher or compactor loses nothing; Kubernetes reschedules it.
- **Data durability** comes from the object store. On-prem, Interkey specifies erasure coding or replication across racks or sites.
- **Metadata** is protected with automated PostgreSQL backups and, where required, a standby replica.
- **DR site.** Object storage replication plus a warm standby cluster pointing at the replicated bucket and restored metastore; RPO and RTO are agreed per customer and tested in a DR drill.
- **Self-monitoring.** The engine's own metrics, indexed events, search requests, disk available, feed Datadog monitors, with an alert when indexer disk drops below 20% free as Datadog recommends.

How it makes troubleshooting faster

Speed comes from correlation, not from any single module: every signal carries the same service, environment and version tags, so an engineer moves from alert to root cause in one screen instead of across five tools. Interkey's services are built around making that correlation work end to end.

The correlation backbone

- **Unified service tagging.** env, service and version are set on every host, container, trace, log, RUM session and synthetic test. Interkey defines the tag taxonomy (plus team, business_unit, site) and enforces it through Agent config, Kubernetes labels and CI templates.
- **Trace-to-log injection.** APM tracers inject dd.trace_id and dd.span_id into application logs. Those logs are indexed in CloudPrem, and the trace view shows them inline.
- **Database linkage.** DBM connects APM spans to the exact query, its plan and its wait events.
- **Network context.** NDM and Network Path place device and path health on the same host and site tags.
- **User impact.** RUM links a user session to the backend trace that served it.

Incident walkthrough: checkout latency

A retail customer's checkout API slows down at 20:10 during a campaign. With the platform in place, one on-call engineer works the incident as shown below; each step is one click from the previous one.

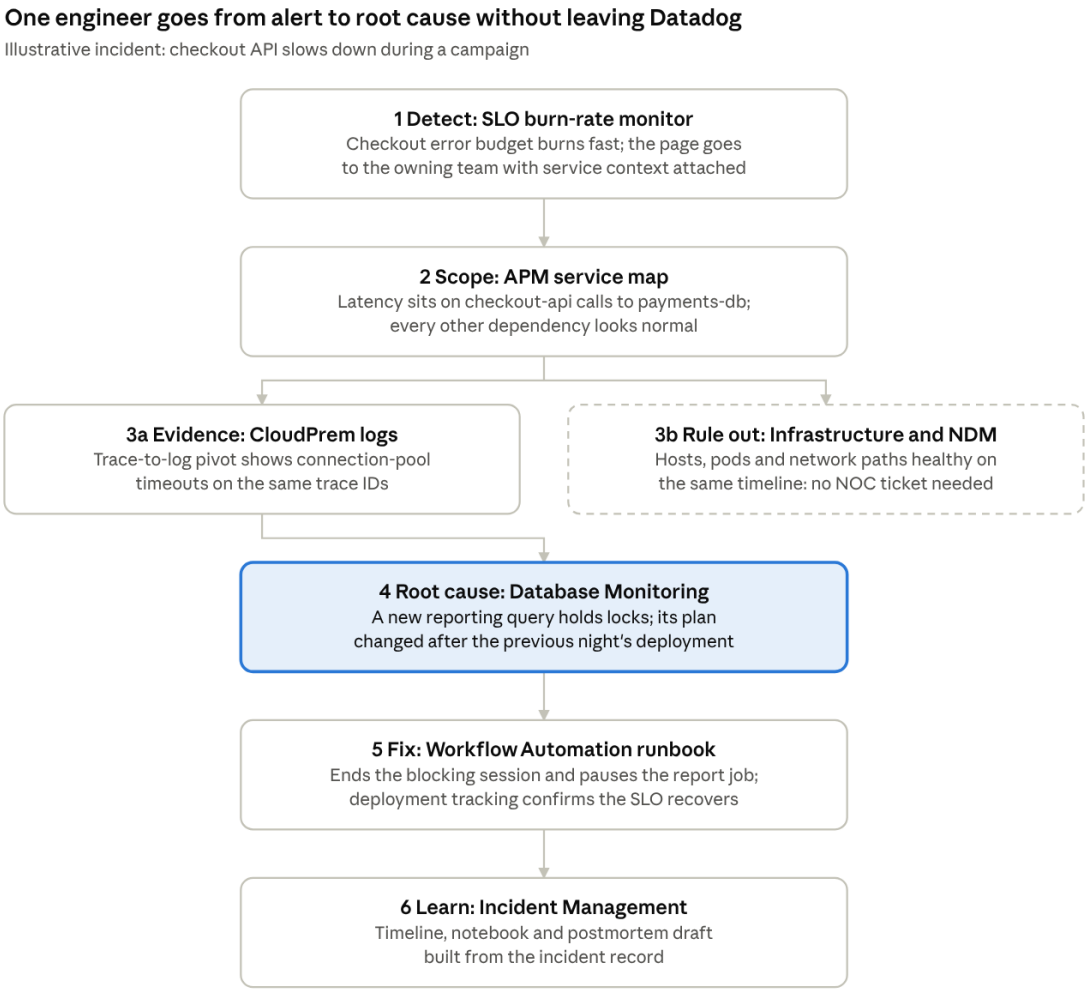


Figure 2. Incident walkthrough · 6 steps, one UI

The network and infrastructure check (3b) runs alongside the log evidence and is closed in minutes because both sit on the same timeline and tags.

Before and after

Troubleshooting step	Without a correlated platform	With Datadog + CloudPrem, delivered by Interkey
Detect	Customer complaint or static CPU alert	SLO burn-rate or Watchdog anomaly on the service, routed to the owning team
Scope	War room; each team checks its own tool	Service map shows the failing dependency and blast radius

Troubleshooting step	Without a correlated platform	With Datadog + CloudPrem, delivered by Interkey
Find the evidence	Request logs from app team; grep across servers	Trace-to-log pivot into CloudPrem logs, already filtered to the trace
Confirm cause	Database team runs ad-hoc queries	DBM shows the query plan and blocking session tied to the span
Rule out infra/network	Separate NOC ticket	Host, container and NDM metrics on the same timeline
Fix and verify	Manual change, watch dashboards	Workflow Automation runbook, deployment tracking confirms recovery
Learn	Postmortem written from memory	Incident timeline, notebook and postmortem generated from the incident record

Interkey does not promise a fixed MTTR figure up front. The Assess phase baselines the customer's current MTTD and MTTR from their ITSM data, and improvements are reported against that baseline.

Interkey professional services catalog

The catalog is modular: CloudPrem and the platform foundation are delivered first, and each Datadog module is an add-on service with its own scope, deliverables and acceptance criteria.

Foundation services (every engagement)

Service	Scope	Key deliverables
Observability assessment	Current tools, critical services, log sources and volumes, ITSM MTTR baseline, regulatory constraints	Assessment report, data-classification matrix, target architecture, sizing workbook
CloudPrem platform build	Kubernetes (EKS, AKS, GKE, OpenShift or upstream), object storage, PostgreSQL, Helm deployment, reverse connection, proxy and TLS	Deployed engine, HA/DR design, runbooks, load-test report
Observability Pipelines design	Sources, parsing, PII redaction, sampling, routing to CloudPrem and archive	Pipeline configs as code, volume-reduction report

Service	Scope	Key deliverables
Agent fleet rollout	Linux, Windows, Kubernetes, VMs; config management via Ansible, GPO or Helm	Agent standards, automated install, fleet health dashboard
Tagging and governance	Unified service tagging, teams, RBAC, SSO/SAML, Audit Trail	Tag taxonomy, role model, governance playbook

Module services

Module	Interkey service	Troubleshooting outcome
Log Management (CloudPrem)	Log pipelines, facets, indexes and retention tiers, log monitors, saved views	Evidence searchable in seconds, kept in customer storage
Infrastructure and Kubernetes	Integrations (VMware, cloud, OS), Kubernetes and container monitoring, host maps	Resource saturation spotted before it becomes an outage
APM and Distributed Tracing	Tracer instrumentation (Java, .NET, Node.js, Python, Go), sampling strategy, service catalog, trace-log correlation	Service map and latency breakdown per endpoint
Database Monitoring	SQL Server, Oracle, PostgreSQL, MySQL, MongoDB setup; query and wait analysis	Slow queries and blocking tied to the calling service
Network Device Monitoring	SNMP v2c/v3 profiles, autodiscovery, topology, NetFlow, Network Path	Network ruled in or out in minutes
Synthetic Monitoring	API and browser journeys, private locations inside the network	Failures found before customers report them
RUM and Session Replay	Web and mobile SDKs, privacy masking, frontend-to-backend linking	Real user impact per release and region
Cloud SIEM and Sensitive Data Scanner	Detection rules, content packs, PII scanning rules	Security and operations share one evidence store
Incident Management and On-Call	Severity model, escalation policies, ITSM integration (ServiceNow, Jira)	Right team paged with context attached

Module	Interkey service	Troubleshooting outcome
Workflow Automation	Runbooks: restart service, scale pods, reboot devices, open tickets	Repeatable fixes executed in minutes
SLOs, dashboards and Watchdog	SLOs per critical journey, executive and team dashboards	Alerts on user impact rather than raw thresholds
Bits AI SRE and AI assistants	Enablement, investigation playbooks, MCP integration review	AI-assisted first pass on incidents

Module availability depends on the customer's Datadog site and licence; Interkey confirms each module against the customer's contract before scoping.

Delivery methodology

Every engagement follows four phases with a sign-off gate between each, so the customer approves the design before build and the platform proves itself before go-live.

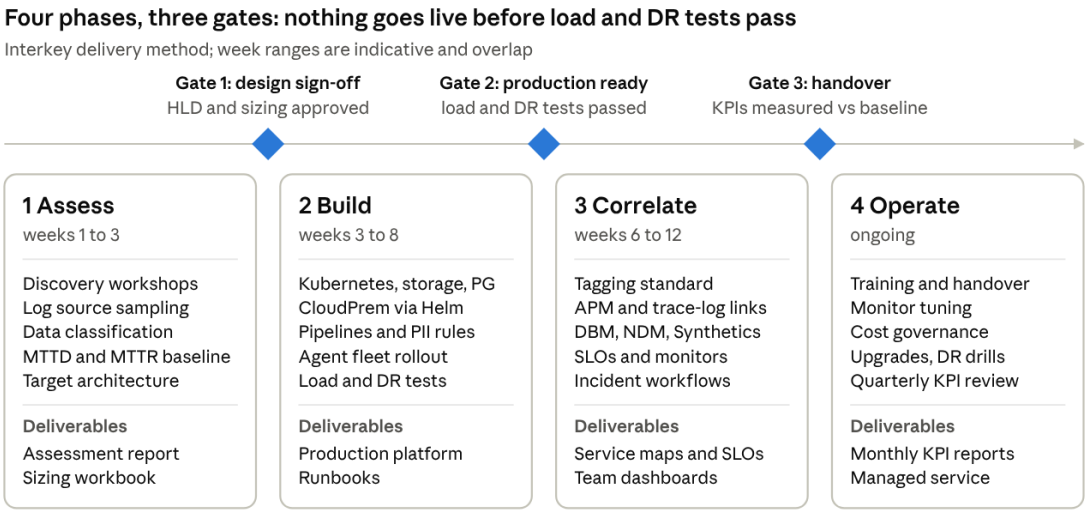


Figure 3. Delivery roadmap · 4 phases, 3 gates

Correlate starts while Build finishes: once the first environment's Agents report, instrumentation and SLO work begins on the critical services. All configuration, Helm values, pipelines, monitors, dashboards, is delivered as code (Terraform and Git) so it can be reviewed, versioned and re-deployed to the DR site.

Engagement packages and managed services

Customers can start small with a proof of value and grow into full platform coverage and a managed service. Durations below are indicative and are fixed during scoping.

Package	Best for	Scope	Indicative duration
CloudPrem Proof of Value	Proving residency, cost and search speed	CloudPrem on a non-production cluster, 1 pipeline, 3–5 log sources, 1 critical service traced, success criteria agreed upfront	3–4 weeks
CloudPrem Launch	First production rollout	Production CloudPrem with HA, pipelines, Agent fleet for one environment, tagging standard, log monitors and dashboards	6–10 weeks
Full-Stack Observability	Enterprises replacing several tools	Launch + APM, DBM, Infrastructure, Synthetics and SLOs for an agreed list of critical services	10–16 weeks
Network and Security add-on	NOC and SOC use cases	NDM, Network Path, Cloud SIEM, Sensitive Data Scanner	4–8 weeks
Digital Experience add-on	Customer-facing web and mobile	RUM, Session Replay, browser synthetics, frontend SLOs	3–6 weeks
Automation add-on	Reducing toil	Incident Management, On-Call, Workflow Automation runbooks, ITSM integration	3–6 weeks

Interkey Managed Observability

After go-live, Interkey can run the platform under a monthly service:

- **Platform operations**: CloudPrem upgrades (Helm chart and engine releases), capacity management, DR drills, PostgreSQL and storage health.
- **Configuration hygiene**: tag compliance, monitor tuning, noisy-alert reviews, index and retention tuning.
- **Cost governance**: monthly ingest and usage reports, pipeline tuning to hold log cost per GB.
- **Onboarding factory**: new applications instrumented to the standard within an agreed SLA.
- **Expert on call**: L2/L3 support for complex incidents alongside the customer's operations team.

- **Quarterly service review**: MTTD, MTTR, alert noise and adoption against the baseline.

Enablement is included in every package: role-based training for SREs, developers, NOC and SOC, and a handover pack so the customer can operate independently if they choose.

Security, governance and regulatory alignment

The design principle is simple: raw log content stays in customer storage, and anything sent to Datadog SaaS is classified and minimised first. Interkey documents this as a data-flow register the customer's compliance team can sign off.

Controls built into the design

- **Data minimisation at the edge.** Observability Pipelines masks PII (national ID, IBAN, card numbers, phone numbers) before logs are indexed, and Sensitive Data Scanner rules catch what pipelines miss.
- **No inbound exposure.** The reverse connection needs only outbound 443 through the corporate proxy.
- **Encryption.** TLS in transit; object storage and PostgreSQL encrypted at rest with customer-managed keys where the platform supports it.
- **Access control.** SSO/SAML, role-based access and restriction queries so teams see only their own services' logs.
- **Auditability.** Datadog Audit Trail records who searched, exported or changed configuration.
- **Retention.** Per-index retention in CloudPrem mapped to the customer's records-retention schedule; the janitor enforces deletion.
- **APM and RUM data.** Trace and session attributes are scrubbed of personal data through tracer obfuscation and RUM privacy settings, since these go to SaaS.

Regulatory mapping (KSA and GCC)

Requirement area	How the design addresses it
Saudi PDPL (SDAIA), personal data processing and cross-border transfer	Log content, the richest source of personal data, remains in customer-controlled storage; PII masked before any SaaS-bound telemetry

Requirement area	How the design addresses it
NCA Essential Cybersecurity Controls (ECC) and Cloud Cybersecurity Controls (CCC)	Centralised event logging and monitoring, RBAC, audit trails, defined retention, secured egress
SAMA Cyber Security Framework (banking)	Security event monitoring, incident management workflow, evidence for audits
CST regulations (telecom and cloud)	In-Kingdom storage of logs on customer infrastructure or a local cloud region
PCI DSS v4.0	Log review and retention (Requirement 10), card data masked before indexing

This mapping supports, but does not replace, the customer's own legal and compliance assessment. Interkey provides the evidence pack; the customer's compliance function makes the determination.

Success metrics and KPIs

Every engagement is measured against a baseline taken in week one; targets are agreed with the customer rather than assumed.

KPI	Definition	Source	Reported
MTTD	Time from fault start to alert on the owning team	Incident Management timeline, ITSM	Monthly
MTTR	Time from alert to verified recovery	Incident Management, ITSM	Monthly
Tools per incident	Number of separate tools used to reach root cause	Postmortem review	Quarterly
Alert noise ratio	Share of alerts closed with no action	Monitor analytics	Monthly
Tag compliance	Share of hosts, services and logs carrying the required tags	Usage and tag reports	Monthly
Service coverage	Critical services with APM, logs and SLOs in place	Service Catalog	Monthly

KPI	Definition	Source	Reported
Log cost per GB ingested	Infrastructure + licence cost divided by ingested volume	Pipelines metrics, cloud bill	Monthly
Search performance	p95 query latency for standard log searches	CloudPrem engine metrics	Monthly
Proactive detection rate	Incidents detected by monitors or synthetics before a user report	ITSM	Quarterly

Why Interkey

Interkey combines Datadog delivery experience with deep data-platform engineering, which is what a CloudPrem deployment needs: it is as much a Kubernetes, storage and database project as an observability one.

- **Datadog field experience in KSA.** Delivered a retail Datadog rollout covering 75 hosts, 55 APM applications and 20 databases, and branch-wide Network Device Monitoring with FortiGate autodiscovery and SNMPv3.
- **Data-platform depth.** Senior architects with 17+ years in databases, event-driven architectures (Kafka) and enterprise software, the skills needed to size object storage, run PostgreSQL metastores and instrument complex backends.
- **Local regulatory context.** Designs start from PDPL, NCA and SAMA requirements rather than retrofitting them.
- **Reseller and services in one contract.** Licensing, implementation and managed operations under one accountable partner.

Next steps

1. A 90-minute discovery workshop with the customer's operations, application and security leads.
2. Data classification and log-volume sampling for the top 10 log sources.
3. A CloudPrem Proof of Value with success criteria signed off before it starts.

Appendix A: customer prerequisites checklist

- Kubernetes cluster (EKS, AKS, GKE, OpenShift or upstream) with node pools for indexers, searchers and compactors across at least 2 zones
- Object storage bucket (S3, Azure Blob, GCS or S3-compatible on-prem) with lifecycle and encryption policies
- PostgreSQL instance for the metastore, with automated backups
- Network-attached block storage class for indexer volumes (at least 250 GB per indexer)
- Outbound HTTPS 443 from the cluster to the Datadog site, via proxy if required
- Datadog organisation, API and application keys, SSO/SAML integration
- Access to hosts or config management for Agent rollout
- Read-only database users for DBM; SNMP credentials for NDM
- List of critical services and business journeys, with owners
- Data-classification decision on which telemetry may leave the country
- 3 months of ITSM incident data for the MTTD/MTTR baseline

Appendix B: glossary

Term	Meaning
CloudPrem / BYOC Logs	Datadog log engine deployed in customer infrastructure
Split	An index file written by an indexer to object storage
Metastore	Service that tracks where splits live, backed by PostgreSQL
Observability Pipelines	Datadog's log processing and routing layer, run in the customer network
Unified service tagging	Shared env, service, version tags that link all telemetry
MTTD / MTTR	Mean time to detect / mean time to resolve
SLO	Service level objective, a reliability target tied to user experience

Sources

- Datadog BYOC Logs (CloudPrem) overview
- BYOC Logs architecture
- BYOC Logs network model
- BYOC Logs production best practices
- BYOC Logs release notes
- Datadog blog: Store and search logs at petabyte scale with BYOC Logs