Observability improvements and coverage

- Investigated and reduced false-positive alert storms (notably “lost signal” style noise) by adjusting thresholds/time windows and confirming root causes.
- Refined database deadlock alerting so it is more actionable (including separating deadlock signals by application and tightening trigger logic to avoid “bulk noise”).
- Added pipelines to onboard the new LPA environments
- Added configuration to monitor the INT environment

Synthetic monitoring

- Implemented/updated browser synthetic checks for key login and availability journeys across major platforms (e.g., NLIS, eNVD, and MyMLA).
- Extended synthetic coverage into additional non-production environments and aligned those scripts with production behaviour where needed.
- Troubleshooting failing synthetics (including SOAP/API synthetics) to restore reliability of uptime checks.

Logging and telemetry

- Improved forwarding and configuration of server log ingestion where multiple/obsolete configurations existed, reducing confusion and ensuring the correct logs are captured.
- Added targeted logging configurations for critical scheduled/batch components so operational teams can self-serve troubleshooting without manual log pulls.
- Expanded monitoring for key integration endpoints (e.g., eNVD ↔ NLIS integration), focusing on 24/7 availability expectations.

Agent and platform maintenance (stability and consistency)

- Addressed inconsistent infrastructure agent versions and problems with automated monthly upgrades; introduced steps to bring environments back into alignment.
- Performed upgrades in lower environments first, with a validation-first approach before production changes.

Documentation & handover

- Produced practical documentation covering key management patterns (when to use license keys vs user/API keys) and how the observability-as-code pipeline works.

- Provided knowledge sharing to the wider Sysops team

Outcome / impact (high-level)

- Fewer false positives and duplicated alerts, improving on-call and support experience.
- Clearer alert ownership and grouping, improving triage speed.
- Better confidence in service availability via synthetic coverage of critical user journeys.
- Improved troubleshooting speed through better log ingestion and clearer operational documentation.
- Reduced operational risk through standardised agent maintenance and upgrade practices.