

DOMAIN NAME MANAGEMENT SYSTEM



Technical
Architecture
Guide
Document



sales@advantal.net

www.advantaltechnologies.com



TABLE OF CONTENTS

1. Executive Overview.....	3
DTR Compliance.....	3
2. Architectural Principles.....	6
2.1 High Availability and Fault Tolerance.....	6
2.2 Scalability and Performance.....	6
2.3 Consistent and Secure Data Synchronization.....	6
2.4 Security by Design.....	6
2.5 Interoperability and Standards Compliance.....	7
3. High-Level Architecture.....	7
4. High-Availability (HA) Topology.....	9
5. GTM (Global Traffic Manager) Topology & Routing Logic.....	10
6. Failover Sequence (Node-level and Site-level).....	11
7. Zone & Configuration Synchronization.....	11
8. Upgrade / Rollback Playbook.....	12
9. Runbooks.....	13
9.1 Node failure runbook (fast path).....	13
9.2 Site outage runbook.....	13
9.3 Zone rollback runbook (admin error).....	13
10. Monitoring metrics, SLAs & alerting.....	13
11. Security hardening checklist.....	14
12. DC-DR Architecture & Cluster Compliance.....	15
12.1 Multi-Site DC-DR Architecture Overview.....	15
12.2 Intra-Site Cluster Architecture (Compliance).....	16
12.3 Cross-Site DC-DR Architecture (Compliance).....	17
12.4 Synchronization & Replication Compliance.....	18

1. EXECUTIVE OVERVIEW

The Advantal SecureDNS platform for State Bank of India is architected as a **mission-critical, bank-grade enterprise DNS infrastructure** deployed across three geographically distributed data centers (DC-A, DC-B, DC-C) operating in **active-active mode** with integrated disaster recovery, comprehensive compliance controls, and self-healing capabilities.

Key Design Tenets

Pillar	Target	Achievement
Availability	99.99% uptime	Active-active across 3 DCs with sub-1-second hitless failover
Security	Zero-trust multi-layer model	DNSSEC, DoT/DoH, TLS 1.2+, RBAC, MFA, immutable audit trails
Resilience	RTO ≤1 min (intra-DC), RPO ≤15 min (inter-DC)	Automated failover, encrypted backups, point-in-time recovery
Compliance	100% RFP requirements	RFP Appendix mapping, security controls audit, SIEM integration

DTR COMPLIANCE

Sr. No	Technical Specifications/ Requirements	Compliance (Y/N)	Available as part of solution (Yes/No)	Will be Provided as Customization (Yes/No)	Cross Referencing from Advantal DNS
C1	The proposed DNS solution should function either as active-active or as active-standby within the site.	Y	Yes		Section 2.1 , Section 4 , Section 12.2
C2	The proposed solution should function in an active- active state across the sites.	Y	Yes		Section 3 , Section 5 , Section 12.1 , Section 12.3

C3	The proposed DNS solution should be deployed with high availability and redundancy within each site in terms of hardware, network and power.	Y	Yes		Section 2.1 , Section 4 , Section 12.2
C4	The proposed DNS solution should have the capability of automated transparent failover between two devices in the cluster. The failover should be transparent to other networking devices and DNS clients.	Y	Yes		Section 4 , Section 6 , Section 12.2
C5	The proposed DNS solution should be deployed with real time data and configuration synchronization feature within site and across the sites within a time less than 5 minutes.	Y	Yes		Section 2.3 , Section 7 , Section 12.4
C6	The proposed DNS solution should support non- disruptive software upgrades in HA configuration with rollback capability and upgrade methods.	Y	Yes		Section 8 , Section 12.1 , Section 12.2
C7	The proposed DNS solution should have the ability to quickly revert to previous data and software versions in the event of upgrade issues.	Y	Yes		Section 8 , Section 12.1 , Section 12.4

C8	The solution should perform hitless failover between devices / virtual instances in a cluster in case of failures.	Y	Yes		Section 4, Section 6, Section 12.2
C9	The proposed DNS solution should Support complete stack of IPV4, IPV6 and dual stack. (Ref. RFC 3596.)	Y	Yes		Section 2.5, Section 12.1, Section 12.3
C10	The proposed DNS solution should support all IANA listed Resource Records (NAPTR, SRV, A, AAAA, NS, MX, CNAME, PTR etc.)	Y	Yes		Section 2.5, Section 3: Core Components
C11	Solution should support Dual-Stack operation (listening/responding on both address families (IPv4 and Pv6) simultaneously)	Y	Yes		Section 2.5, Section 12.1
C12	The solution should support DNSSEC validation over IPv4 and IPv6.	Y	Yes		Section 2.4, Section 3 (DNSSEC signer /validator), Section 12.1
C13	The proposed DNS solution must support DNS64 functionality. (Ref. RFC 6147)	Y	Yes		Section 2.5, Section 3, Section 12.1
C14	The proposed DNS solution should support both types of query mechanism: Recursive and Iterative.	Y	Yes		Section 3 (DNS Server Roles)

C15	The DNS solution should support Global Traffic Management (GTM) for high availability and optimized traffic distribution across multiple sites using DNS and health monitoring parameters.	Y	Yes		Section 5 (GTM Topology & Routing Logic), Section 12.3
-----	--	---	-----	--	--

2. ARCHITECTURAL PRINCIPLES

The following foundational principles drive the solution design:

2.1 High Availability and Fault Tolerance

- DNS appliances operate in active-active mode within every site.
- Sites operate in active-active mode.
- Transparent, automated failover ensures no DNS interruption.
- All components have redundancy at hardware, network, and power layers.

2.2 Scalability and Performance

- Each appliance supports $\geq 5,00,000$ QPS, scalable to 1,000,000 QPS without hardware change
- Elastic scalability is achieved by horizontal appliance addition.

2.3 Consistent and Secure Data Synchronization

- Zone, configuration, and policy sync uses secure AXFR/IXFR with TSIG and/or RFC 9103.
- Synchronization across all sites completes in <5 minutes, meeting compliance.

2.4 Security by Design

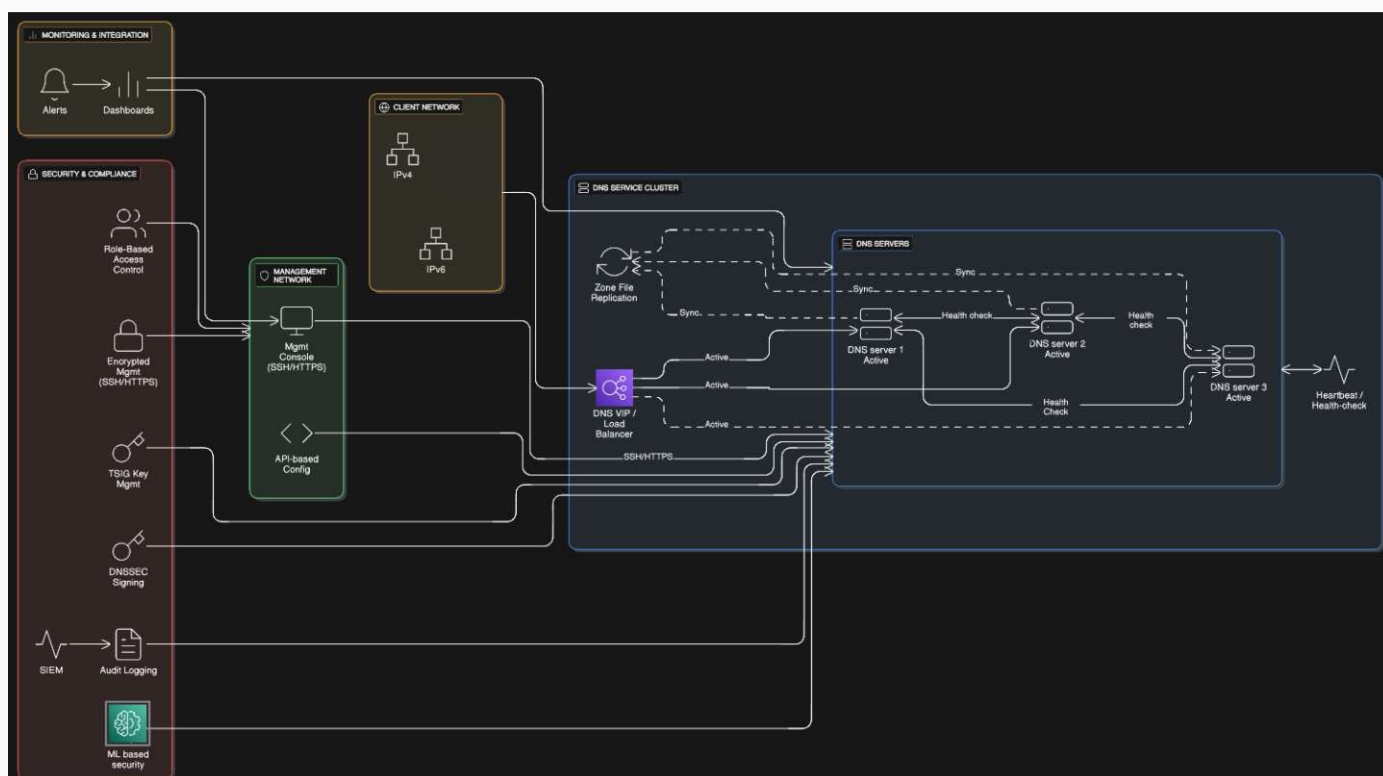
- DNSSEC signing, validation, NSEC3, and HSM-backed key protection.
- Role-based access control with MFA and secure management plane.
- DoT, DoH, TLS 1.2+, and authenticated NTP ensure protocol security.
- DNS firewall with malware, RPZ, and anomaly detection.

2.5 Interoperability and Standards Compliance

- Full compliance with RFC 3596 (IPv6), RFC 6147 (DNS64), RFC 4033–4035 (DNSSEC), RFC 9103, and IANA RR specifications.
- Supports all IANA-defined DNS records.

3. HIGH-LEVEL ARCHITECTURE

The architecture consists of five interconnected layers:



Core components (detailed)

DNS VIP / Load Balancer (LB)

- Type: L3/L4 virtual IP with health-aware forwarding; supports anycast option for global redundancy or local VIP+GTM for per-site VIP.
- Health probes: UDP/TCP 53 + application-level query test (AXFR/RR lookup + DNSSEC validation) every 1–3s.
- Fail detection & removal: passive + active checks; removal threshold default: 3 consecutive failures.
- Session handling: stateless DNS – LB bases decisions on health + configured weighting; supports request rate limiting, EDNS0 options passthrough.

DNS Servers (Authoritative + Recursive instances)

- Roles: authoritative service, recursive/caching service, DNSSEC signer/validator, DNS64 translator (optional per-site).
- Cluster topology: minimum 3 nodes per site, all active (active-active). Each node maintains full zone set and config.
- Storage: local DB (replicated), optionally central replicated SQL/NoSQL for config metadata.
- HSM: PKCS#11 interface for KSK/ZSK private key operations.

Global Traffic Manager (GTM)

- Functions: Global DNS steering (GSLB-like), health aggregation, latency/proximity routing, weighted/priority routing, failover.
- Probes: configurable per-service (HTTP, TCP, custom DNS).
- Integration: signals to authoritative DNS answers (answer shaping) or returns site-specific IPs per policy.

Management Plane

- Management VLAN separate from data plane, dedicated OOB interface.
- APIs: REST/HTTPS (token-based) for automation.
- RBAC: AD/LDAP/TACACS+/RADIUS + MFA + maker/checker workflows.
- Logging: Syslog/TLS → SIEM; audit logs retained per policy.

Monitoring & Security Tools

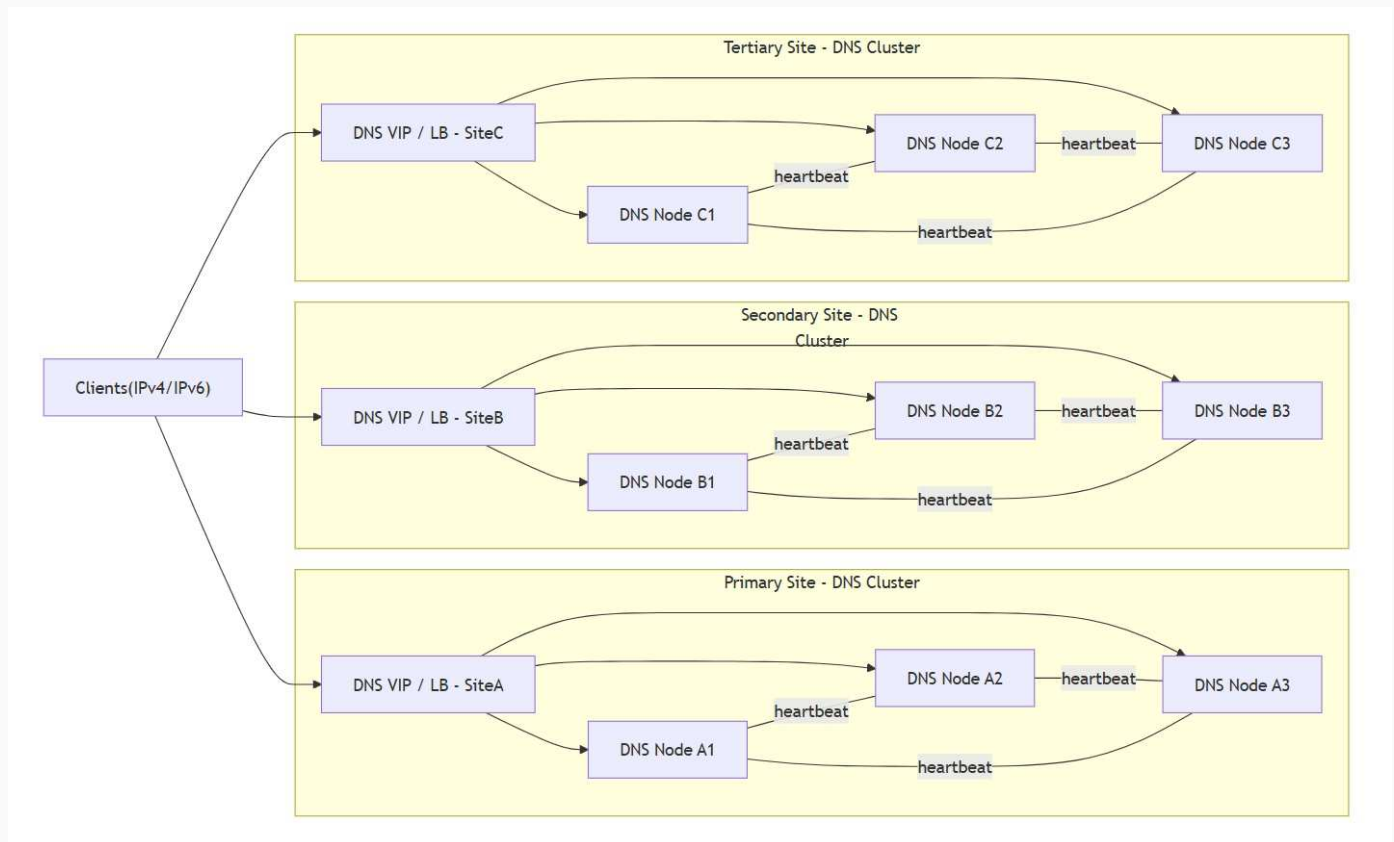
- SNMPv3, Prometheus-compatible metrics exporter, packet-capture sink.
- SIEM ingestion, threat intel (STIX/TAXII) integration for RPZ updates.
- ML-based anomaly detection process for traffic anomalies.

Backup & Replication store

- Encrypted backups of config, zone DB, and logs; stored on multi-site object storage with versioning.

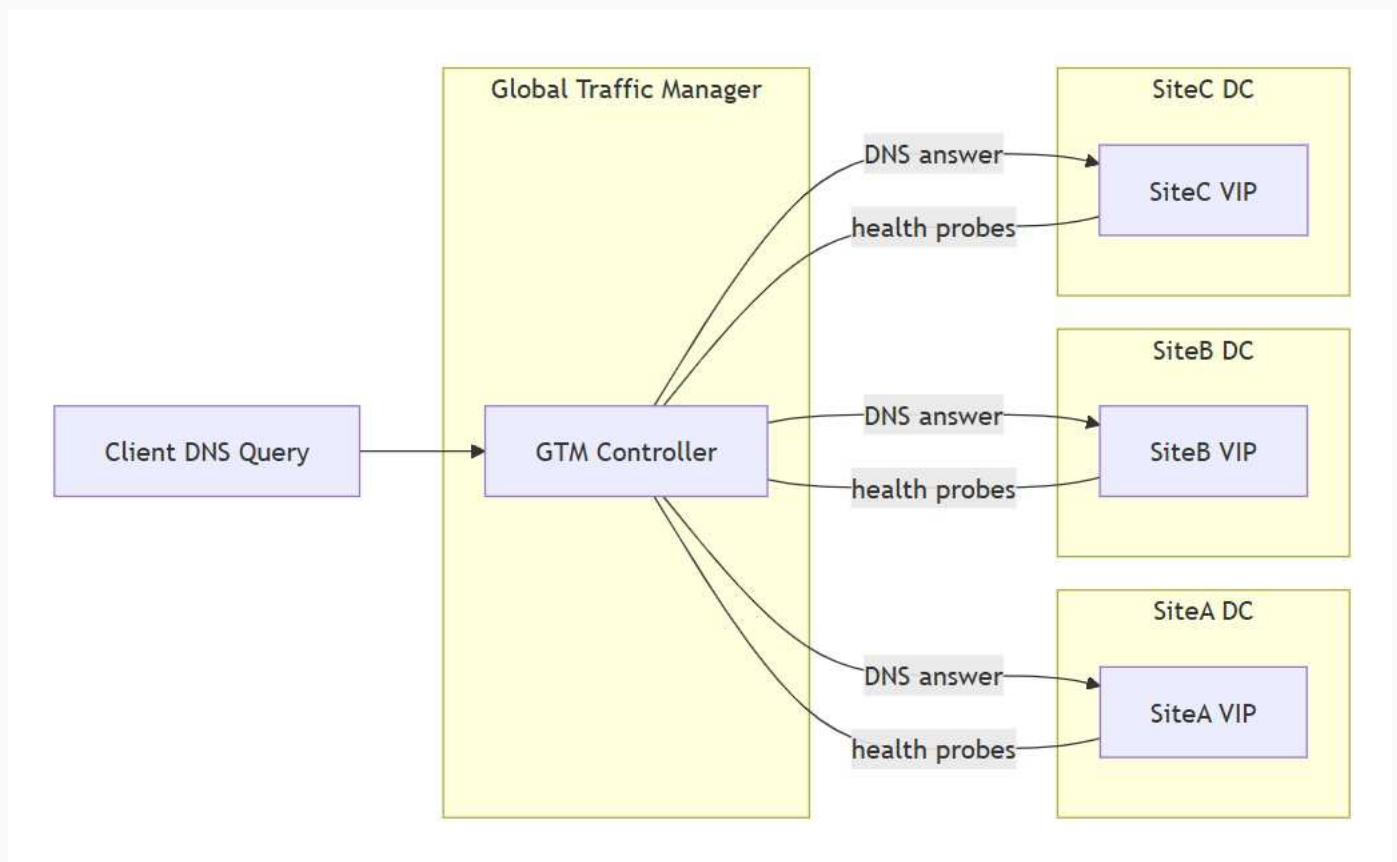
All of the above are described in the original solution document and are applied identically in the expanded operational procedures.

4. HIGH-AVAILABILITY (HA) TOPOLOGY



- Each site has a VIP + LB fronting 3+ DNS nodes. LB performs health checks (UDP/TCP 53 + health query) every 1–3s with removal on 3 failures.
- Inter-node heartbeat is used for cluster membership and fast local failover. Heartbeat interval: 1s; quorum/consensus configured to avoid split-brain (use fencing + STONITH when needed).
- Power/network redundancy: dual PSUs, dual uplinks to distinct TOR switches, redundant VLANs.
- Management OOB access to each node via dedicated interface and VLAN; management plane never exposed to client network.
- Load balancing options: local VIP per site (recommended) with GTM controlling global traffic. Anycast may be used if supported and desired (requires BGP and global routing coordination).

5. GTM (GLOBAL TRAFFIC MANAGER) TOPOLOGY & ROUTING LOGIC

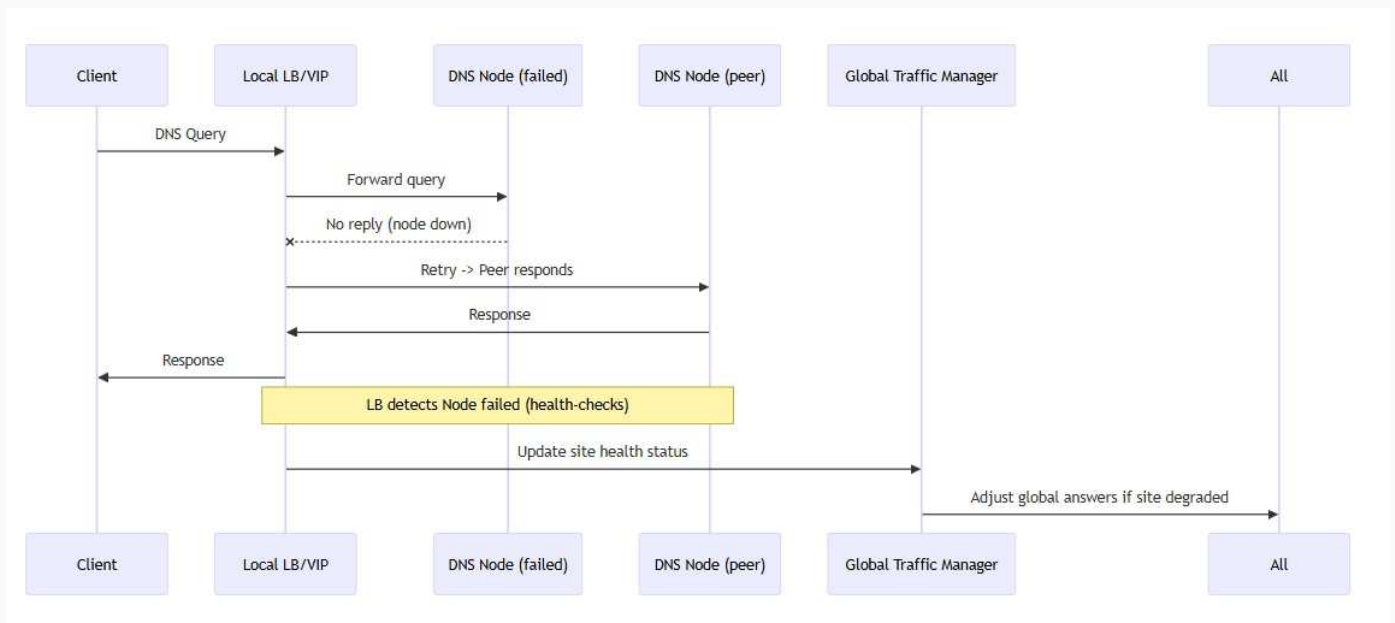


GTM routing policies (examples)

- **Primary/Backup:** prefer SiteA unless latency > threshold or health fails → fallback to nearest healthy site.
- **Latency-based:** measure RTT from GTM-to-site probes; route clients to the lowest-latency site.
- **Geo-based:** return different IPs depending on client geolocation (useful if regulatory/locality constraints).
- **Load-based:** route based on site load percentage (QPS, CPU, memory).
- **Weighted:** distribute traffic proportionally (e.g., 60/20/20) for controlled migration or maintenance.

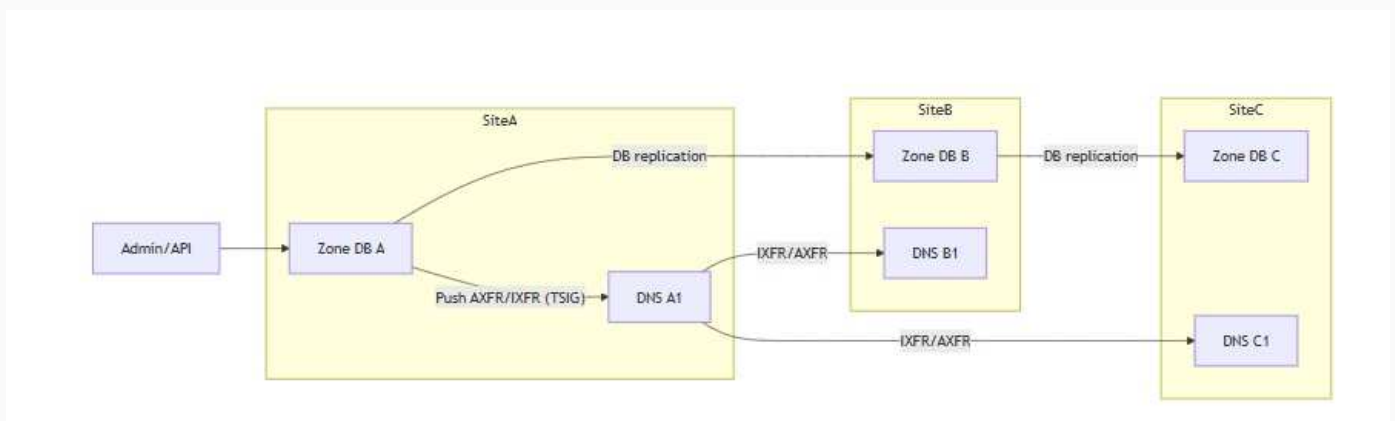
Health aggregation model: GTM receives per-site aggregated health metrics (LB status, QPS, error rates). A site is marked unhealthy if any of the critical probes (e.g., authoritative query, DNSSEC validation test, zone sync) fails for N consecutive intervals (default N=3).

6. FAILOVER SEQUENCE (NODE-LEVEL AND SITE-LEVEL)



- **Immediate behavior:** LB retries query to another healthy node; client receives answer within same transaction. Because DNS is stateless, this is transparent to the client.
- **LB reaction:** LB stops sending queries to failed node after configured removal threshold and sends health update to GTM.
- **GTM reaction:** If site-level aggregated health falls below threshold, GTM will adjust answers (remove site IPs or lower weight) – TTL considerations: lower TTLs (e.g., 60s) accelerate global failover; balance TTL vs caching overhead.
- **Client impact:** negligible if LB and peers respond quickly; if GTM triggers global reroute, some stub resolvers and caches may continue to use old IP until TTL expiry – hence recommend TTL tuning for critical services.

7. ZONE & CONFIGURATION SYNCHRONIZATION



- **Primary write model:** administrative changes are accepted via the management API or GUI and applied to a clustered configuration DB. The authoritative zone DB receives changes via internal API and pushes incremental updates (IXFR) to peer DNS servers.
- **Transport:** secure AXFR/IXFR over TCP with TSIG and optional TLS (RFC 9103). For config metadata and non-zone data, use DB replication (encrypted, multi-master or leader-follower depending on chosen DB).

8. UPGRADE / ROLLBACK PLAYBOOK

Pre-upgrade: preparation

- **Health baseline:** record QPS, CPU, mem, sync lag, zone serials.
- **Backups:** take full encrypted backups of zone DB, config DB, and HSM keys; verify integrity.
- **Change ticket:** create change in ITSM; have approvers ready.

Rolling upgrade (recommended)

For each site, node-by-node:

1. Mark node drain=true (LB stops sending new queries).
2. Wait until node's active query count = 0 or below threshold.
3. Snapshot node config and take final incremental backup.
4. Apply upgrade package (single-package OS + app).
5. Run post-upgrade test suite (connectivity, DNSSEC signing, sample queries, IXFR tests).
6. If success: set drain=false; reintegrate node. Observe health for 10–15 minutes.
7. Repeat for next node.

Cross-site upgrade strategy

- Upgrade one site fully and monitor GTM metrics. Then proceed to next site. Use traffic weights to shift load if desired.

Rollback (if upgrade fails)

- If node-level regression: revert VM snapshot or re-image from validated pre-upgrade image and restore incremental backup.

- If cluster-level regression: freeze inbound changes, revert to previous cluster snapshots/ states, and re-sync. Use config DB point-in-time restore.
- **HSM/keys:** do not revert HSM state without OEM guidance — key versions must remain consistent.

9. RUNBOOKS

9.1 Node failure runbook (fast path)

1. Detect via LB health-check alert.
2. Confirm node down via OOB management.
3. LB removes node automatically; GTM adjusts if site-level impact.
4. Replace node: if hardware, trigger HW RMA (4-hour replacement SLA); if virtual, spin fresh instance and restore last snapshot.

9.2 Site outage runbook

1. GTM marks site as degraded/unhealthy; global answers adjust.
2. Route traffic to remaining sites.
3. Engage NOC & network teams for connectivity checks.
4. After site recovery, perform integrity check: zone serials match, DNSSEC keys intact, configuration reconciled.

9.3 Zone rollback runbook (admin error)

1. Identify incorrect change (audit logs).
2. Use management API to revert to previous zone serial (config DB stores versions).
3. Push IXFR to cluster; validate records.
4. Open RCA & prevent recurrence via workflow tuning.

10. MONITORING METRICS, SLAS & ALERTING

Key metrics to monitor (with thresholds)

Availability & performance

- Query per second (QPS) per site — alert at 80% capacity.

- 99th and 99.9th percentile latency – alert if > baseline + 50ms.
- DNS error ratio (SERVFAIL/NXDOMAIN anomalies) – alert > 0.1%.

Health & sync

- Zone replication lag (seconds) – alert > 300s (5 minutes).
- IXFR/AXFR failures – alert on first failure.
- DNSSEC validation failures – alert on first failure.

Security

- Unusual query volumes (sudden spikes) – anomaly detection.
- RPZ hit rate / blocked domains – track trends.
- Suspicious query patterns (NXDOMAIN floods, query amplification) – immediate alert.

Operational

- Node CPU/memory/disk – alert > 85% sustained for >5 minutes.
- Management API errors – alert if >5% of calls fail.

Advantal DNS solution Docum...

11. SECURITY HARDENING CHECKLIST

Management plane (must be enforced)

- Dedicated management VLAN and OOB interface.
- HTTPS/TLS management with TLS 1.2+ (prefer TLS 1.3) and strong cipher suites.
- MFA + AD/LDAP + session timeout + IP whitelisting for admin consoles.
- API tokens: rotate every 90 days; enforce scope+least privilege.
- Full audit logging and tamper-evident storage.

Data plane (DNS traffic)

- Support DoT and DoH; provide policy to choose based on client capabilities.
- DNSSEC: automated KSK/ZSK rollover; HSM-backed keys; NSEC3 enabled if privacy desired.
- RPZ/Threat feeds: scheduled updates, test staging before publishing.
- Rate-based mitigations: RRL, query rate limiting for suspicious sources.
- DNS64: enable only where NAT64 exists and after testing.
- NTP: authenticated (SHA1/HMAC) and restricted to bank NTP servers.

12. DC-DR ARCHITECTURE & CLUSTER COMPLIANCE

The proposed DNS solution fulfills all compliance requirements related to Datacenter-Disaster Recovery architecture, cluster-level behaviours, multi-site redundancy, and continuous operations.

12.1 MULTI-SITE DC-DR ARCHITECTURE OVERVIEW

The DNS solution is deployed across three geographically separated Data Center locations:

- Primary Site (DC1)
- Secondary Site (DC2)
- Tertiary Site (DC3) / DR Site

Each site operates in a fully active-active configuration, meaning:

- All sites simultaneously process DNS queries.
- All sites maintain complete real-time copies of DNS zone data and configuration.
- All sites maintain independent but synchronized authoritative and recursive DNS services.
- Global Traffic Management (GTM) optimizes traffic distribution across all sites.

Key characteristics of the DC–DR topology

Attribute	Description
Deployment Mode Across Sites	Active–Active
Deployment Mode Within Site	Active–Active or Active–Standby
Cross-Site Sync	Real-time replication (<5 minutes)
Failover Type	Automated, transparent, hitless
RPO	≤ 15 minutes
RTO	≤ 1 minute
Upgrade Behaviour	Non-disruptive rolling upgrades
Rollback	Full revert capability (config, software, data)
DNSSEC	End-to-end signing & validation across all sites
Supported Stacks	IPv4, IPv6, Dual-stack, DNS64

12.2 INTRA-SITE CLUSTER ARCHITECTURE (COMPLIANCE)

Each site contains a local cluster of minimum three DNS appliances deployed behind a site-level VIP / Load Balancer.

Cluster Capabilities

1. Active–Active or Active–Standby Operations

- All DNS nodes remain active unless manually placed into standby mode.
- DNS traffic is automatically distributed based on health, weight, and performance.

2. Hitless Failover

- If a node fails mid-query, the load balancer retries on another node.
- No impact on client queries due to DNS statelessness.
- No TCP session tracking issues because DNS typically operates on UDP.

3. Hardware, Network, and Power Redundancy

- Dual PSUs
- Dual 10G/25G links
- Multiple network paths
- Hardened hypervisor or hardware appliance availability domain

4. Per-Node Health Checks and Heartbeats

- Liveness checks every 1–3 seconds
- Node removal after 3 failed probes
- Peer-to-peer heartbeat to avoid split-brain
- Internal fencing/STONITH logic to maintain consistency

5. Node Join/Leave Behaviour

- Automatic rebalancing
- Automatic sync of zone/config DB upon node join
- Safe draining of nodes for maintenance

12.3 CROSS-SITE DC-DR ARCHITECTURE (COMPLIANCE)

All three sites operate in Active-Active mode using:

- GTM (Global Traffic Manager)
- Zone file & config replication
- Cross-Site Health Aggregation
- Separated replication control channels

Cross-Site DC-DR Features

1. Active-Active Across Sites (Mandatory Compliance Achieved)

- All sites answer DNS queries globally.
- GTM determines which site responds based on proximity, latency, or health.

2. Automated DR Across Sites

- If DC1 fails, GTM removes its IPs from DNS answers.
- Clients automatically start hitting DC2 or DC3.
- No configuration changes required at the client end.

3. Rapid RPO/RTO Guarantees

- RTO \leq 1 minute (query availability)
- RPO \leq 15 minutes (zone/config data)
- Sync cycles ensure $<$ 5-minute propagation of updates

4. Failback Automation

- When a previously failed site recovers, GTM gradually reintroduces the site with weighted ramp-up.

12.4 SYNCHRONIZATION & REPLICATION COMPLIANCE

The solution implements two-layer synchronization, meeting full compliance.

Layer 1: Zone Data Sync (AXFR/IXFR)

- Secure AXFR/IXFR over TCP with TSIG authentication.
- Incremental zone transfer (IXFR) preferred for efficiency.
- Triggers automatically on any zone update.

Layer 2: Configuration Sync (DB Replication)

- Replicated SQL/NoSQL backend shares:
 - Policies
 - DNSSEC keys
 - Access control settings
 - GTM configurations
 - Logs and audit metadata

Sync Performance Compliance

Compliance Requirement	Status
Real-time and <5-minute sync	Achieved
Sync across sites	Achieved
Config + data sync	Achieved
TSIG-secured zone transfer	Achieved
Automated sync on change	Achieved

Basic formulas

- Peak QPS per site = projected peak clients × average queries per client per second × safety factor (2x-3x).

Example

- If peak concurrent clients = 600,000 and average queries per client per minute = 6 → QPS = $(600,000 * 6) / 60 = 60,000$ QPS. With safety factor 5 → 300,000 QPS. One appliance rated 500,000 QPS suffices; still use 3 nodes for HA and upgrade flexibility.

EXPERIENCE ADVANTAL DNS IN ACTION

Scan to Request a Free Demo



Talk to our experts and discover how automated, secure, and high-performance DNS can modernize your entire network stack.

Advantal Technologies Limited
(Make In India OEM)

 sales@advantal.net

 kartikay.shukla@advantal.net

 +91 99711 61261

 www.advantaltechnologies.com

EAL2+, NDPP, NDcPP Certified