THE TICKET IS NOT THE PROBLEM

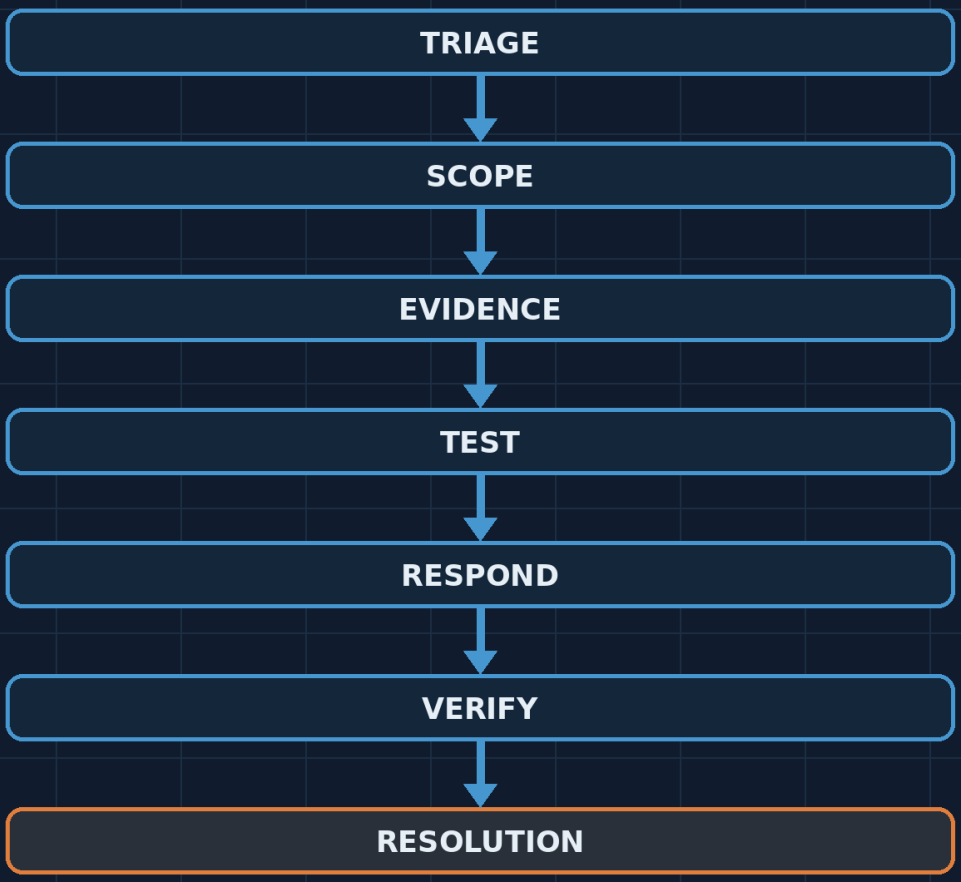

DANIEL LANG

Scenario 2 — “It Is DNS”: A Company-Wide Problem That Grows in Front of You

This scenario uses the same seven decisions, but the facts force a different response. A junior engineer should not be expected to solve a core-DNS outage alone. A careful early comparison detects that the problem is wider than a normal ticket, triggers fast communication and escalation, and prevents a risky local fix from wasting the first useful minutes.

Step 1 — Define the First Symptom Without Shrinking It

At 09:07, an accounts user reports that the payroll portal will not open. The browser says the site cannot be reached. A second user reports that Teams links and the intranet are also failing. Email is still working on some laptops. The first report is not “a payroll issue”; it is two users reporting failed name-based access to several services.

The junior records the exact messages, users, locations and time. They also check impact: payroll approval and customer-service systems may be affected, while email behaviour is inconsistent. There is no safe basis yet for restarting a browser, resetting an account or declaring a vendor outage.

Step 2 — Compare Working and Failing Paths, Then Widen Scope

A known internal application can be reached by IP address but not by its normal hostname. On one affected device, the default gateway responds, but the configured DNS resolver does not answer a lookup. Similar reports arrive from another floor and from a VPN user. That comparison moves the likely boundary from one user, laptop or application toward shared name resolution.

Observation	What it suggests	What it does not establish
Default gateway responds	Local network reachability exists on this device	Every network path is healthy
Internal service works by IP but not hostname	Name resolution is a strong candidate	The DNS server itself is necessarily the cause
Multiple offices and VPN users see the pattern	Impact is broader than one segment	Every employee and service is affected

The junior now treats this as a potentially company-wide service degradation. They notify the duty lead using the approved channel and update the ticket with the observed scope. That is escalation based on impact and evidence, not on the label someone used.

Step 3 — Keep the DNS Theory Honest

“It is DNS” becomes the leading hypothesis because multiple name lookups fail while basic IP reachability still works. It is not yet a confirmed root cause. A firewall policy, routing path to resolvers, DHCP configuration, resolver forwarding path or an upstream dependency could produce a similar symptom.

The evidence register is deliberately short: successful gateway ping; failed lookup against the configured resolver; affected users in separate locations; timestamped reports; and any monitoring alert. The junior avoids turning an unverified observation such as “the DNS dashboard looks red” into proof without checking its time and scope.

Step 4 — Choose the Smallest Safe Test That Separates Explanations

The useful test is a bounded query against each configured resolver from an approved diagnostic host, alongside a comparison query to a known external resolver where policy permits. The test is read-only. It distinguishes “this workstation is misconfigured” from “our resolver path is failing”; it does not repair anything.

The primary and secondary internal resolvers both time out from the diagnostic host. The external comparison resolves a public domain, while internal names remain unavailable. This strengthens the case for an internal DNS service or its reachable dependency. The junior does not change workstation DNS settings across the fleet; doing so would create inconsistent state and might bypass names that should remain internal.

Step 5 — Stabilise Service Without Confusing Recovery with Diagnosis

The network or platform owner reports that a failed maintenance change interrupted both internal resolvers’ access to directory-integrated zone data. They roll back the approved change and confirm resolver responses return. The junior’s role is to keep impact and evidence current, communicate the approved workaround if one exists, and verify user tasks once recovery is declared.

At 09:32, hostname lookups succeed from the diagnostic host. At 09:35, users in accounts and on VPN can open the payroll portal and intranet. That establishes recovery of the user-facing symptom. It does not by itself prove every causal detail of the maintenance failure.

Step 6 — Escalate with the Useful Record, Not a Pile of Screenshots

The escalation and incident update contains: start time; impacted business tasks; affected locations; IP-versus-name comparison; resolver query results; actions deliberately not taken; owner of the rollback; and user verification after recovery. It names the next decision: investigate why a maintenance change removed both resolver paths or their common dependency.

The junior did not need to be an expert DNS administrator. They needed to recognise that the safe next step had moved beyond their authority, preserve a small set of discriminating evidence, and prevent five different people from “fixing” their own workstation at the same time. That gives the incident owner useful evidence and a clear next decision.

Step 7 — Verify, Analyse and Prevent Recurrence

Because the issue was broad, business-impacting and linked to a change, it needs more than “DNS restarted and people are happy again.” A subsequent post-incident review should test whether the triggering maintenance change, lack of staged validation, shared resolver dependency, alerting coverage, and rollback checkpoint contributed to the impact. Those are review questions, not findings established by the recovery alone.

Corrective actions might include a pre-change resolver query check, staged rollout with a stop condition, independent-path health monitoring, and a tested rollback. Each action needs an owner and a verification method. Root cause analysis identifies causes and contributing conditions; the corrective actions change the conditions that allowed one change to produce company-wide impact.

Scenario 2 Wrap-Up — The Response Widened as the Evidence Changed

The first user report looked like an application ticket. The comparisons changed that: gateway reachable, names failing, resolver unavailable, multiple locations affected. Each result made a broader response more justified. The outcome was not “the junior fixed DNS”; it was a faster transition from a local report to a managed outage response, supported by evidence.

The full guide will be announced at ticketnotproblem.com.